The Philosophical Quarterly Vol. 00, No. 0
ISSN 0031-8094 doi: 10.1093/pq/pqv097

TYING THE KNOT: WHY REPRESENTATIONALISTS SHOULD ENDORSE THE SENSORIMOTOR THEORY OF CONSCIOUS FEEL

By David Silverman

The sensorimotor theory of perception and consciousness is frequently presented as a variety of anti-representationalist cognitive science, and there is thus a temptation to suppose that those who take representation as bedrock should reject the approach. This paper argues that the sensorimotor approach is compatible with representationalism, and moreover that representationalism about phenomenal qualities, such as that advocated by Tye, would be more complete and less vulnerable to criticism if it incorporated the sensorimotor account of conscious feel. The paper concludes by arguing that the project of naturalizing phenomenal qualities would nonetheless be better served by abandoning 'representation' talk altogether, a move that would require only a small modification of existing representationalist accounts.

Keywords: sensorimotor theory, *enactivism*, phenomenal qualities, qualia, representationalism, consciousness.

I. INTRODUCTION

Scientifically oriented accounts of visual consciousness almost always appeal, in one way or another, to the idea that perceivers or their brains internally represent the outside world. The idea, known as 'representationalism', is a basic tenet of orthodox vision science, which following Marr (1982) claims that visual consciousness is realized by the brain's construction and deployment of representations. A more demanding kind of representationalism, which will be my focus in this piece, accepts the orthodox approach from vision science while also claiming that the phenomenal qualities that comprise visual consciousness—for example, the conscious feel of seeing the colour red—can be exhaustively explained by appeal to the content of the representations that realize the experience. One leading proponent of this claim is Tye (2000).

The sensorimotor theory[1] (O'Regan and Noë 2001) is an increasingly influential account of perception and phenomenal qualities which explicitly rejects orthodox vision science. Instead of appealing primarily to representation, it emphasizes the roles played by bodily interaction with the environment. The theory is often glossed as a variety of *enactivism*, an approach to cognitive science which typically rejects internal representation entirely, claiming that perception and other forms of cognition depend on nothing more or less than interactions between the brain, non-neural body and environment (Buhrmann, Di Paolo and Barandiaran 2013; Myin and Degenaar 2014). There is therefore a natural temptation to suppose that if one is committed to representationalism, one must reject the sensorimotor approach.

This paper aims to persuade those who take representation as bedrock to endorse the sensorimotor approach to conscious feel anyway. It will show that a combined representationalist and sensorimotor account is a viable position in logical space, and more importantly that such a position offers an improvement on standard representationalism about phenomenal qualities. After having, I hope, persuaded the reader of this headline claim, I will finish by arguing that abandoning talk of internal representation is a relatively small further step, and that there are good reasons to make such a move. In the process of offering these arguments, I aim to sharpen up the sensorimotor theory's sometimes unclear core commitments.

In the next section, I introduce the sensorimotor theory's accounts of perception and phenomenal qualities. In Section III, I explain why the theory faces apparent pressure to make explanatory reference to internal representation, before showing in Section IV how this invites certain prominent lines of criticism. In Section V, I show that representationalism about phenomenal character does not actively preclude the sensorimotor account and, in Section VI, show that representationalists have positive reason to endorse the sensorimotor account of phenomenal qualities. In Section VII, I propose a modification of the representationalist approach to phenomenal qualities, such that it talks about 'attunement' rather than 'representation'.

[1] By 'sensorimotor theory', I mean the theory set out by O'Regan & Noë (2001), building on earlier work by Hurley (1998), and developed in more recent work by those authors (e.g., Noë 2004, 2012; O'Regan 2011) and others who explicitly address their work (e.g., Burhmann, Di Paolo and Barandiaran 2013; Seth 2014). The theory can be situated amongst a complex web of theoretical ancestors and cousins that also consider perception or consciousness to be 'sensorimotor', but I will not address those here.

II. TWO ROLES FOR EMBODIED INTERACTION IN SENSORIMOTOR THEORY

To properly understand the sensorimotor theory, we must appreciate an important distinction, sometimes overlooked, between two separate projects in which the theory is engaged. The first is giving a scientifically tractable, naturalistic account of perception. This involves saying what sort of activity perception is, and what sorts of physical or functional processes enable it to take place. The second is explaining why perceptual consciousness feels the way it does: for instance, why seeing the colour red has the qualitative feel it does. This requires us to identify non-arbitrary principles by which descriptions of conscious perception couched in phenomenal and non-phenomenal physicalistic vocabulary can be bridged. We could think of the first project as the project of explaining perception, and the second as the project of explaining phenomenal character.

Although these projects are connected, they are not, on the face of it, identical. One big reason for thinking they are not identical is offered, for instance, by Block (1996), who argues that the intentionally directed component of perception is not sufficient for the phenomenal component, and therefore that accounting for the conscious feel of seeing the colour red requires us to appeal to facts distinct from the facts that explain our perceiving red objects. The sensorimotor theory, as I will understand it, endorses a version of the opposing view, intentionalism, which claims that there is nothing more to having an experience with a particular phenomenal quality than being intentionally directed in a particular way. The advantage of intentionalism, if it can be defended, is that it makes it easier to eliminate the explanatory gap between qualia and the physical, since it is often thought that explaining perception naturalistically is within our grasp, while we have few ideas about how we would naturalize qualia without appeal to intentional directedness.

If we assume, for the sake of argument, that intentionalism is correct, there is nonetheless a sense in which the projects of explaining perception and phenomenal qualities come apart. This is because there is more than one way to make the conceptual link between the physical or functional processes that yield perceptual consciousness and the phenomenal character of the consciousness so yielded. For example, Tye (2000) and Clark (2008) each endorse the generic claim that perception depends, subpersonally, on the activation of neural representations with an appropriate kind of content and functional role. Tye thinks that the contents of the representations are identical to the phenomenal qualities that feature in the conscious experience. Clark, however, endorsing a skill-based account (see Pettit 2003), claims that being in a perceptual state with a particular phenomenal character is identical to exercising a particular personal-level capacity for skilful discrimination, which

the neural representations merely help enable. Clark's and Tye's positions are, at least on the face of it, quite different.

The import of this point is that we must take care not to confuse an account of the processes that constitute or enable perception with an account of the principles that bridge descriptions of these processes and descriptions of the phenomenal character of perceptual experience. The sensorimotor theory, for its part, makes a relatively unorthodox appeal to bodily interactions to explain both perception and phenomenal character. However, it appeals to them differently in each case.

II.1 Bodily interaction and phenomenal qualities

Let's first see how the sensorimotor theory de-emphasizes representation and emphasizes bodily interaction in its account of phenomenal character. Phenomenal qualities are first-person properties of conscious experience, for example the look of red or the smell of coffee, and there is a well-worn debate about whether they can be re-described using non-phenomenal physicalistic vocabulary, and indeed whether they are physical at all (e.g., Jackson 1982). As physicalists, we might be happy to take it as an inviolable principle that everything that exists is physical, and therefore that phenomenal qualities cannot imply the existence of anything non-physical. Nonetheless, it will be unsatisfying, especially if we want to persuade those swayed by anti-physicalist intuitions, if we cannot show how the qualities can be made naturalistically intelligible.

One popular view is that attempts to naturalize phenomenal qualities have made no progress at all so far. Block, for instance, states that

> 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: 1113)

Sensorimotor theorists agree with a version of this claim: *if* we take consciousness to depend only on neural processes, we have no way to explain why there is consciousness or why particular experiences have one qualitative feel rather than another. O'Regan & Noë (2001) observe that we might, for instance, try to explain the phenomenal difference between audition and vision by pointing to different areas of the brain that are active respectively when we hear and see, or different kinds of neural activity, for example oscillations, that occur in those areas. The problem is that this kind of approach necessarily leaves unanswered the question of why the relevant neural and phenomenal events are correlated.

The sensorimotor theory departs from Block's perspective by denying that the phenomenal character of consciousness, especially perceptual

consciousness, is solely determined by facts about the brain. To explain why a perceptual experience has one feel rather than another, the theory refers instead to the character of loop-like interactions between the brain, non-neural body and outside environment. More precisely, it appeals to *sensorimotor contingencies* (SMCs), the pattern-like ways that inputs from the sense organs are prone to change in line with movements by the agent or objects in the environment, properties to which the theory claims we exhibit a skilful sensitivity when we perceive.

The appeal to SMCs promises to explain the phenomenal quality associated in general with each sense modality (e.g., vision), and the individual phenomenal qualities within each modality (e.g., the look of red). In the former case, the relevant SMCs are those determined by the physical characteristics of the sense organs. Consider an example used by O'Regan and Noë. A horizontal straight line, meeting your eyes at their equator, will be projected onto the retina as a great arc, such that it would appear in a flattened version of the retinal image as a horizontal straight line. The eyes' curvature means, however, that moving your eyes up or down would result in the image of the line becoming increasingly curved.

Eye-related SMCs like these do not apply to the ears, which follow their own distinctive patterns: rotating your head to the left and right, for instance, typically produces increases and decreases in the amplitude of the signal as the ears get further or closer from a sound source. Thus, the different sense organs are associated with different kinds of experience because of the differing ways they modulate patterns of movement-dependent change in sense input.

The differing phenomenal qualities that make up a given sense modality are accounted for by patterns of SMC determined by characteristics of the object. This is easiest to grasp in the case of shape. For instance, a circular object such as a coin, faced at an acute angle, is projected onto the retina as an ellipse, and the projected shape will be more or less elliptical depending on the angle you face it from. SMCs like these determine the subjective character of shape experience.

Crucially, SMCs are supposed to account in a similar fashion for the more challenging case of colour experience. Here, the theory is concerned with pattern-like changes to the ways photoreceptor cones are activated. The theory claims that the character of colour experience is determined by a relation between changes in the composition of the light that meets the object before being reflected onto the retina and changes in cone activation (Philipona and O'Regan 2006; O'Regan 2011). The composition of the light meeting the object before being reflected onto the retina varies in line with changes to ambient lighting conditions and also, in a more finely sensitive way, with changes to the perceiver's precise bodily alignment to the object, i.e. in line with movement. These patterns of movement-related and ambient light-related

change in retinal stimulation are different for differently coloured objects, and it is these patterns that the sensorimotor theory uses to explain the phenomenal qualities of colour experience.

The sensorimotor account claims that phenomenal qualities can be fully accounted for by identifying them with patterns of SMC. The authors attempt to make this plausible by casting doubt on the 'qualia' conception, the idea that a 'qualitative residue' (O'Regan and Noë 2001: 1014) necessarily remains even after we have re-described phenomenal qualities with reference to relevant physical events or properties. To this effect, Noë (2004) claims that careful introspection reveals that it is impossible to ever isolate in one's experience a simple quale, and that visual experience is instead always an experience of structure, which the account cashes out in terms of SMCs.

The SMC approach to colour qualities is given support by empirical work from Philipona & O'Regan (2006), which looked at changes in how photoreceptors are stimulated in the presence of a range of surfaces and lighting conditions. There are three kinds of cone photoreceptor, which respond respectively to light of long, medium and short wavelengths. To describe cone activation in the presence of a given surface, in a given lighting condition, you must specify a value for each of the three photoreceptors. To specify how cone activation varies in different lighting conditions, you must therefore, ordinarily, specify a value for each of three variables in each lighting condition.

The study found, however, that for surfaces with certain reflectance profiles, the explanation can be simplified, since the activations of the three photoreceptors do not vary independently of one another in line with changes to lighting. In the case of red, yellow, blue and green surfaces, it was possible to account for changes in retinal stimulation across varied lighting conditions by specifying values in each lighting condition for only one or two variables.

Philipona and O'Regan observed that this correlates with anthropological data showing that red, yellow, blue and green are the four most universally adopted colour classifications (Berlin and Kay 1969). The idea that colour experience depends on patterns of ambient light-related change in cone activation offers a way to explain these data, supporting the idea that colour qualities are identical to SMCs.

A more general advantage of the sensorimotor approach is its ability to explain and in principle quantify degrees of phenomenal difference and similarity. Consider the improvement it makes on other accounts that address such differences and similarities. Palmer (1999) uses verbally reported similarity and difference judgements to avoid Locke's (1689/1975) suggestion that two individuals could give the same colour terms to the objects around them but undergo colour experiences that are inverted relative to one another. Palmer argues that colour space is plausibly asymmetrical, meaning that should one person's colour phenomenology be inverted relative to another's, their verbal

reports of qualitative difference and similarity between the colours of objects will sometimes differ, and thus the inversion can be detected.

Clark (2000) takes this idea a step further, arguing, contra Palmer, that accounting for the structure of colour space in terms of phenomenal difference and similarity suffices to explain the phenomenal character of colour qualities. This is not, however, the same as explaining how physical events in the brain or elsewhere give rise to phenomenal qualities, and Clark concedes that to naturalize phenomenal qualities it would be necessary also to identify neural features with which the phenomenal qualities are correlated.

The sensorimotor theory attempts something more ambitious than both these authors. It is compatible with Clark's claim that accounting for the structure of phenomenal quality space is sufficient to account for the nature of phenomenal qualities themselves. However, instead of relying on verbally reported similarity and difference judgements, it offers a candidate for the physical features that actually constitute phenomenal qualities, namely SMCs, and it uses these to quantify degrees of difference and similarity.

The advantage of this approach is that it will result, if successful, in an account not just of the phenomenal character of colour qualities, but of how they are physically instantiated. Moreover, the sensorimotor theory encompasses much more than just colour, offering a way to explain phenomenal qualities across modalities, including differences and similarities both within and between them, while also doing justice to the natural thought that phenomenal qualities are determined at least in part by properties of objects in the environment, and not only by the brain.

II.2. Perception as bodily exploration

SMCs are one way the sensorimotor theory appeals to bodily interaction; it does so in a different way in its account of perception, as we will see now. The sensorimotor theory claims that perception consists primarily of bodily exploration rather than inner modelling of the outside environment, as orthodox accounts claim. On its most moderate reading, the theory suggests that internal modelling or representation is not as pivotal as traditionally suggested. On some formulations, the theory is allied to a literature that denies we should appeal to representation at all when accounting for perception and other cognitive phenomena (Buhrmann, Di Paolo and Barandiaran 2013; Myin and Degenaar 2014).

The theory's claim that perception is mainly a matter of bodily exploration rather than representation is derived from O'Regan's earlier idea of the 'world as an outside memory' (O'Regan 1992). This idea was motivated by the observation that orthodox accounts of vision have not made convincing inroads into addressing what should be considered a central question, namely how people see so well given the limitations of the human visual apparatus.

It is well known that the human retina is, from one perspective, very poorly designed for the job of vision. The photoreceptors are distributed unevenly throughout the retina and get increasingly sparse as they get further towards the outer perimeter. The retinal image is obscured in many places by a web of blood vessels, and at the point where the optic nerve meets the eye there is a blind spot completely obscuring a significantly sized portion of the visual field. This means that much of the retinal image is blurry, missing key detail and short on the information about wavelengths required to detect colour.

In spite of these limitations, visual experience seems to ordinary visual perceivers to present a highly resolved and colourful environment. To explain the disparity, orthodox accounts must claim that typical visual perceivers are radically mistaken about the character of their own visual experiences (Dennett 2002) or alternatively that the visual system internally reconstructs a richly detailed visual scene, a view which is implausible given empirical evidence that visual systems do not process large amounts of detail all at once (see O'Regan and Noë 2001 for a review of some of this evidence).

The outside memory thesis rejects the widespread view, following the tradition of Marr (1982), that perceptual consciousness is enabled solely by the retinal image or representations derived from it. The thesis instead claims that perception is a temporally extended process analogous to remembering. In O'Regan's (1992) examples, you might be reminiscing about your grandmother, or remembering what you had for breakfast. Long-term memories typically come to be consciously entertained after being deliberately sought out, and do not arrive fully formed in all their detail: instead, you engage in a process of asking yourself questions, and answering them, and this way accessing facts about the past in a piecemeal way.

The outside memory view contends that vision is an analogous process, although the role of the memory store is played by the outside world and not the brain.[2] Perceivers make use of shifts in attention and especially bodily movements, including saccadic eye movements, to probe the environment. This involves accessing visual detail serially, a little at a time, and in accordance with the perceiver's or visual system's present interests, i.e. the facts they want to find out and the locations they want to explore. Rather than subserving the construction of a detailed internal model or representation, this activity of temporally extended probing takes the notional inner model's place. The fact that perceiving is constituted by a temporally extended process of accessing rich environmental detail in this way indicates that perceivers are not in the grip of an illusion when they take it that visual experience presents them with rich detail (Noë 2002).

An analogy can also be made between vision and touch (see e.g. O'Regan and Noë 2001; Noë 2004; O'Regan 2011). While vision is prone to invite an

[2] Similar ideas have been earlier offered by Minsky (1986) and Brooks (1991).

intuitive commitment to the idea that the experience is realized by nothing more than a detailed neural world model, it is very natural to think that perceiving the world by touch is a process of temporally extended bodily probing: in an example suggested by Noë (2004), we might think of a blind person who uses a cane to sample the environment one detail at a time.

Importantly, O'Regan (1992) claims that it does not typically feel to blind people that they are embedded within anything other than a richly detailed environment. He suggests that the experiential presence of the richly detailed environment is accounted for in the case of a blind person by her touch-based probing of the environment, and concurrent appreciation of the movements required at a given moment to bring further detail into view. The suggestion is that vision works in the same way as touch, and that the eye is, in a relevant but generally unacknowledged respect, similar to the hand.[3]

Perceivers must, we should note, somehow grasp how the detail they are presently accessing forms part of the larger visual scene. O'Regan's early account claimed that retinal stimulation is accompanied by a 'non-metric awareness' (1992: 474) of the locations of specific objects in the visual scene, which incorporates an implicit understanding of the eye movements or other bodily movements needed to access more detailed information about them by displacing the retina. O'Regan claimed that you may have this kind of implicit awareness even of objects you have your back to, implying that the awareness is neither pictorial in its phenomenal character, nor properly speaking visual at all. The early paper also suggested that it is enabled by neural representations which are symbolic rather than iconic, and sparse in detail. The representations do not themselves realize visual awareness, but help enable the temporally extended probing of which vision consists.

The sensorimotor theory, as it has emerged in more recent work, endorses an updated version of this view (O'Regan and Noë 2001; Noë 2004; O'Regan 2011). It now claims that there is no fundamental distinction between one's perceptual awareness of what is present and one's non-perceptual understanding of the movements required to find out more. Both result from one and the same faculty, known as sensorimotor mastery—an implicit understanding of the sensory results of possible movements (i.e. SMCs). Thus, vision is the same as visual exploration. It takes place for the purpose of improving the perceiver's sensorimotor mastery, while being constituted by her exercise of the same sensorimotor mastery. And it is by reference to sensorimotor mastery that the conceptual link is created between the outside memory thesis and the SMC account of phenomenal character, since phenomenal character is determined by the SMCs over which the perceiver is presently exercising mastery.

[3] O'Regan credits the description of the eye as a 'giant hand' to a lecture delivered to high-school students by Donald Mackay.

III. WHY SENSORIMOTOR MASTERY MAY REQUIRE INTERNAL REPRESENTATION

We have seen that the sensorimotor theory claims that perception should be understood as exploration rather than representation, and that it thereby downplays the theoretical role, if any, that representation can play.

I underline the difference between SMCs and mastery of SMCs because the concepts are not always distinguished as clearly as they should be. Buhrmann, Di Paolo & Barandiaran (2013), despite making useful progress on formalizing the concept of SMCs, appear to conflate it with sensorimotor mastery when they say: 'According to the sensorimotor approach, perception is a form of embodied know-how, constituted by lawful regularities in the sensorimotor flow or in SMCs in an active and situated agent' (Buhrmann, Di Paolo and Barandiaran 2013). This is out of line with what I take to be a proper reading of the theory. SMCs do not constitute perception, nor are they properties by virtue of which the agent perceives. What they do is explain why conscious perception has the phenomenal character it does, assuming it takes place. When perception takes place, this depends on the agent's exhibiting some appropriate type of skilful sensitivity to the SMCs that will be in evidence if and when she moves.

The distinction between sensorimotor mastery and SMCs should be respected on account of two significant challenges the theory faces, each a different side of the same coin. First, SMCs can be manifest in an agent's skilful bodily interactions without any attendant phenomenal consciousness taking place. For instance, a guided missile engages in a basic kind of systematic coupling with the environment, but we do not suppose it thereby has phenomenal consciousness. Even a mature human, who is in general conscious, can engage in bodily interactions of the sort described by the sensorimotor theory without these interactions being felt in consciousness—as appears to be the case when you are engaged unreflectively in an activity such as driving (O'Regan and Noë 2001).

Secondly, perceptual consciousness may take place, and have a particular phenomenal character, even when the SMCs that individuate the phenomenal qualities are not at this moment evident in the subject's bodily interactions with her environment. This is demonstrated, in part, by evidence that completely paralyzed subjects have visual experiences (Whitham *et al.* 2011). Indeed, it is a basic tenet of the sensorimotor theory that subjects, due to their possession of sensorimotor mastery, visually experience the felt phenomenal 'presence' of rich environmental detail which goes beyond the environmental detail that is presently perturbing their sense inputs (e.g., Noë 2004, 2012).

These cases show that ongoing bodily interaction with the environment is neither necessary nor sufficient for perceptual consciousness. The sensorimotor

theory must, therefore, appeal not only to SMCs, but also to the agent's possessing an appropriate kind of mastery of SMCs. This entails that the perceiver is skilfully sensitive, in some appropriate way, to how her sense inputs *would* vary in line with merely *possible* movements.

The stipulation that the agent must have an appropriate kind of sensorimotor mastery thus serves a dual purpose. A guided missile, although it possesses a variety of sensorimotor skill, does not possess sensorimotor mastery of a sophisticated enough kind—it is able to act in a way that is (somewhat) skilfully sensitive to the ways that sense inputs will change in line with movement, but it lacks various more sophisticated capacities. For instance, it does not exhibit a particularly flexible bodily responsiveness to the environment. Moreover, it is not able to think about those same regularities, or plan its actions in accordance with them (O'Regan and Noë 2001; O'Regan 2011). A conscious human, who is not currently moving, experiences the presence of environmental detail because she implicitly knows what would happen if she did move (Noë 2004).

The literature on sensorimotor theory does not yet offer a well-developed account of sensorimotor mastery and the subpersonal processes that enable it. The theory's core principles do not offer any reason to actively reject the orthodox claim that perception depends on the deployment of neurally-realised representations. And in the absence of a well-developed non-representational account of sensorimotor mastery, commentators such as Roberts (2010) have been free to claim that sensorimotor mastery itself requires internal representation.

IV. WOULD REQUIRING INTERNAL REPRESENTATION COMPROMISE THE SENSORIMOTOR THEORY?

Prominent critics of the sensorimotor theory such as Clark (2008) and (see O'Regan and Block 2012) have argued not only that the sensorimotor theory faces pressure to concede that there is internal representation, but that such a concession significantly undermines the theory. One problem is that it may now seem to be neurally realized representations, rather than SMCs, that are doing the explanatory heavy lifting when it comes to accounting for phenomenal character. Consider this objection from Block, who writes

> In having a cognitive appreciation of a law involving inputs and outputs, one has to think of or represent those inputs and outputs in some WAY. A machine or a creature from outer space might be able to think of human inputs and outputs in WAYS that do not involve any conscious experience. Alternatively, the WAY might itself be phenomenal—say if our cognitive appreciation is coded in imagery. Given that cognition cannot grasp

> anything without grasping it in some WAY, the appeal to COGNIZING in explaining sensory qualities smuggles in the very notion that is supposed to be explained. (O'Regan and Block 2012)

The quote is suggestive of a few specific problems. Let's consider two, briefly, before considering one that identifies representation, in particular, as a threat to the sensorimotor account of phenomenal character.

One potential problem with the appeal to cognitive appreciation (for which we can, at present, read representation) is that such an appreciation is insufficient to account for the existence, in general, of consciousness. This need not concern us a great deal here, since the theory actually introduces further apparatus to account in general for consciousness. O'Regan (2011) appeals, for instance, to the existence of a 'self', an aspect of the account Block does not address, while the sensorimotor theory might otherwise be allied, for instance, to the biological enactivist tradition (e.g., Thompson 2007), which appeals to the self-organizing properties of living organisms.

Another reading of the WAY objection notes that to count as a representation, a neural state must have non-representational properties which help make this the case. It then says that it must be non-representational properties of the neural states involved in appreciating sensorimotor laws, rather than the representational properties of those states, that determine phenomenal character. However, it is not obvious that an appeal to such properties can do the explanatory work needed to show why the neural state is correlated with a particular phenomenal state, and the appeal to representation holds promise to make the relation between phenomenal and physical properties naturalistically intelligible. In Section VI, we will see that the sensorimotor theory has an advantage compared to mainstream representationalist accounts in the debate with those taking Block's position.

For the time being, let's focus on a version of the WAY objection which pits the sensorimotor theory against those who would account for phenomenal character by appeal to the internal representation of properties other than SMCs. The problem is that should we find ourselves obliged to talk about representations of SMCs rather than SMCs themselves, we may no longer have reason to appeal to SMCs at all. Suppose it is allowed that a red object is identical to an object with which an observer, embodied in a particular way, stands in a particular sensorimotor relation, this being a relation that characterizes the SMCs that will be manifest when the perceiver and object interact. The representations, since they carry meaning, must have intensions as well as extensions, and it appears that a red object could be represented under an intension that does not include SMCs. In light of this possibility, admitting internal representation may seem to undermine the theoretical role the sensorimotor theory ascribes to SMCs.

V. WHY THE SENSORIMOTOR THEORY AND REPRESENTATIONALISM ARE COMPATIBLE

In fact, if the phenomenal character of perceptual consciousness is determined by the contents of a set of neural representations, the contents cannot, I argue, be such that they are not representations of SMCs. Suppose that a perceiver, Bertie, sees a green apple. The SMCs that feature in Bertie's sensorimotor interactions with the apple can be thought of as a relation between Bertie and the apple. For the sake of brevity, let's call this relation R. When Bertie consciously sees the apple, his experience is partly characterized by the phenomenal quality associated with the colour green.

Representationalists like Tye (2000) claim this is because Bertie's brain represents some property of the apple. The property represented is the particular property that causes the apple to be green, although neither we nor Bertie need to know what that property is. Tye offers that the property is an in-general disposition to emit light of certain wavelengths. This is perhaps meant to be an intrinsic property of the apple, and is at any rate a property other than R. Let's use X to refer to the intrinsic properties that green apples possess that contribute to their being green. What we will see now is that a representation of X is, in line with one leading version of representationalism, also necessarily a representation of R (or, to put it differently, every representation with X as its content also has R as its content). This suggests that there is nothing to actively prevent representationalists from appealing to SMCs to individuate phenomenal qualities if they so choose.

Tye claims that a neural state represents a green apple if, under optimal conditions, the neural state is present if and only if a green apple is present.[4] Content is individuated externalistically, i.e. by properties of the apple, rather than intrinsic properties of the neural state, for example its topological properties. When Bertie's brain activates in the presence of the apple, let's suppose it represents X. Whenever X is present, R is also present. On the face of it, we might think, therefore, that a representation of X is also a representation of R.

The account, however, is designed to avoid the need to stipulate that a neural representation necessarily bears content about every feature of its extension. Tye imagines a world in which every purple object is poisonous and vice versa. In this world, the expressions 'purple object' and 'poisonous object' are co-extensive, meaning every time a neural representation activates in the presence of a purple object, it is also doing so in the presence of a poisonous object. Tye claims, however, that the brain may represent such an object under the intension 'purple object' and not 'poisonous object'. This works because

[4] Tye adds that the neural state must also be present *because* the target state is present. This does not affect my argument.

we could coherently hypothesize that the representation would become active if there was a purple object that was not poisonous and would not if the converse were the case. Perhaps, then, Bertie's neural representation of X could become active in the presence of X even if R did not obtain. If so, a representation deployed by the visual system might represent a green apple without representing the associated SMCs.

The advantage may yet be returned to the sensorimotor theory. The notion of intension under consideration only allows two representations with the same extension to differ in intension if there is a possible world in which the representations pick out different objects. The expressions 'Morning Star' and 'Evening Star' differ in intension but share an extension, because they both pick out Venus. However, there is a possible world in which they pick out two different planets. If there were no such possible world, the two expressions would not just be co-extensional, but also co-intensional. 'H_2O' and 'water', for example, are co-intensional because there is no possible world in which one is present without the other.

To show why the claim that 'H_2O is water' is informative, we must appeal to an aspect of meaning called 'hyperintension' (that being the label given to whatever aspect of meaning can play this role). The neural representations we are considering cannot, at Tye notes, possess hyperintension, because the only means we have to determine what they represent is by constructing hypotheses about which sorts of property they would and would not become active in the presence of. Since we cannot coherently consider how a neural state would respond to H_2O that was not water, there is no way that the neural representations we are considering could represent H_2O without also representing water. If R obtains by metaphysical necessity whenever X is present, then every neural representation of X must also be a representation of R.

The problem is that, in contrast with water and H_2O, R does not obtain by metaphysical necessity every time X is present. For one thing, there are possible worlds in which X exists but the perceiver, Bertie, does not. This is not a problem, as the question we are considering is whether Bertie's brain can represent X without representing R. The only possible world in which Bertie's brain can represent X is one in which he exists, so we can disregard possible worlds in which Bertie does not exist. More seriously, although Bertie and the apple necessarily stand in relation R to one another if all other things remain equal, this relation might be disturbed. If R broke down in a completely haphazard way rather than being systematically altered, Bertie would probably not be able to perceive at all. This would save the sensorimotor theory, because we are only considering cases in which Bertie successfully experiences the apple. However, R could perhaps be modified systematically such that the Bertie's brain is still able to perceive the object by representing X.

Arguments from Hurley (1998) and Myin (2001) show that there are constraints, in terms of physical possibility, on the extent to which the sensorimotor relation R could be systematically altered while keeping X and the perceiver constant. For at least certain modifications of R, the required complexity of a device, such as a neural implant, that could cause R to vary while the perceiver's body and the object stay the same would be such as to make the device a physical impossibility. Nonetheless, for some modifications of R, a relatively simple device, such as goggles that invert the visual field from left to right, is a physical possibility. Moreover, in more challenging cases, systematically modifying R while keeping the subject and object the same is at least a metaphysical possibility.

This is where Tye's claim that content is determined by covariance relations that occur under 'optimal conditions' becomes relevant. Tye introduces the optimal conditions clause to defuse an argument levelled by Block against phenomenal quality representationalism. The Inverted Earth thought experiment (Block 1990) features a planet which is identical in all respects to our own, but for the fact that the colours of objects are complementary to the colours they possess on Earth, meaning the sky is orange, the grass red, and so forth. The inhabitants of Inverted Earth correctly describe the sky as 'blue', and the grass as 'green', since the meanings of those words on Inverted Earth differ appropriately from their meanings on Earth. Block asks us to imagine that we put on a pair of colour-inverting goggles—whatever they entail—and travel to Inverted Earth. The inverting effect of the glasses means that things appear phenomenally to us on Inverted Earth exactly as they do on Earth without the glasses. Yet, on standard externalist accounts of representational content, the content our experiences possess must now differ in line with the differing properties of the objects we are now in causal contact with. This appears to show that phenomenal qualities cannot be adequately accounted for by representational content, so long as the content is individuated externalistically.

To avoid this line of attack, which depends as before on the metaphysical possibility of goggles or neural implants that systematically modify the character of retinal input in an appropriate way, Tye asserts that representational content is only determined by environment–brain covariance occurring under 'normal' or 'optimal conditions', which Tye claims colour-changing goggles (etc.) violate. Whatever these conditions comprise—more on which below—if they are enough to defend Tye's account, they are also enough to defend the sensorimotor theory.

Where optimal conditions are appealed to by Tye to ensure that representations of X are accompanied by the appropriate phenomenal character, the sensorimotor theory may observe that R is also present when Bertie and the apple are present and those optimal conditions obtain. From Tye's perspective,

one might claim that R is a necessary enabling condition for Bertie to correctly represent X, since it forms part of the optimal conditions which allow us to tell that the neural representation has the content X. There is, however, no principled reason to regard the obtaining of R merely as a prerequisite for the representation of X, and not as one of the properties represented.

This lesson can also be applied to the critique of the sensorimotor approach offered by Clark (2008). With reference to empirical work by Milner & Goodale (1995), Clark claims that the neural system primarily responsible for visual consciousness, the ventral stream, is responsible for reasoning and long-range action planning, and is not intimately involved with ongoing bodily movement. For this reason, he argues, the ventral stream probably does not represent SMCs—the consequences of possible movements—but coarser-grained 'categories, types and relative locations' (Clark 2008: 192). In addition to appealing to the ventral stream's function, Clark appeals to its architecture. He states that the system is not routinely sensitive to efferent copy, information about activity in the brain's motor regions which ought to be a key source of information about the SMCs that presently apply. He suggests instead that the ventral stream responds to information about SMCs during occasional learning phases, in which it uses information about SMCs to help it associate certain kinds of sense input with certain types, categories and locations.

Clark does not say that phenomenal qualities should be individuated directly by the contents of the representations, but by the personal-level capacities for discrimination they enable. For example, one shade of red looks darker than another, in this view, just in case you are able to judge that it is darker, or plan your actions in accordance with the fact it is darker. Clark does not, however, make explicit what these acts of discrimination entail. If judgement and action planning are themselves constituted by the deployment of internal representations, then phenomenal qualities will be individuated by the contents of the representations anyway.

Note, now, that Clark's critique is silent on the question of content. A covariance-based account of content such as Tye's is indifferent to facts about architecture and functional role. On the question of content, there is thus no relevant distinction to be made between the ventral stream's making routine, constitutive use of efferent copy and only using it instrumentally during learning phases. Clark agrees that information about SMCs is important during learning, but the present account of content does not admit any distinction between reliable sensitivity to SMCs exhibited because of present efferent copy or because of learning. The representations of types, categories and relative locations Clark discusses may also therefore be representations of SMCs.

VI. WHY REPRESENTATIONALISTS SHOULD APPEAL TO SMCS

We have seen that the sensorimotor theory can avoid the charge that when we perceive, say, a green apple, it is representational content about the apple alone and not about the associated SMCs that determines the phenomenal character of the experience. However, the theory remains vulnerable to the objection that SMCs are not the most interesting aspects of the representational content to appeal to. It may be more appropriate, for instance, to appeal to representations of types, categories and locations, as Clark suggests.

Let's consider some positive reasons a representationalist about phenomenal character has to appeal to representations of SMCs rather than some other property. According to Tye, the phenomenal quality of red reduces to representations of whatever property it is that constitutes being red. His loose suggestion is that such properties include being 'disposed to reflect such-and-such percentages of light of so-and-so wavelengths' (2000: 55). In fact, the sensorimotor theory offers a better developed version of this claim. The property of being disposed in general to do something might be an intrinsic property; however, having a disposition to do something implies that this disposition is activated in particular circumstances, and Tye's undeveloped account says nothing about what these circumstances are. As we saw in Section II.1, objects of differing colours reflect light of different wavelengths under different illumination conditions, and there is evidence that the ways they change across lighting conditions are theoretically interesting features that distinguish the colours associated with basic colour terms like red, green, yellow and blue from other colours.

Once we specify ways that retinal stimulation from objects of various colours changes in line with illumination, and so movement, it becomes compelling to view the disposition as a relation between the perceiver and object: one characterized by SMCs. Thus, one reason for a supporter of Tye's approach to endorse the sensorimotor theory is that it makes better developed, empirically supported moves towards an account of the properties that feature as the content of the visual representations.

Next, recall from the last section Tye's claim that content is individuated by the features the brain covaries with under optimal conditions. The clause is introduced to avoid being problematically committed to the view that on Inverted Earth, your brain represents the features you are now in causal contact with, despite the unchanged phenomenology. As Tye understands it, 'optimal conditions' are the conditions that obtain in the ancestral environment in which the organism's brain first evolved to covary with the property purportedly represented. As Tye notes, this threatens to entail that a 'Swamp Man'—an atom-for-atom duplicate of a human subject, lacking any evolutionary history—represents nothing, and therefore lacks phenomenal consciousness,

and this would be an unpalatable consequence of the theory. To avoid this, Tye claims that for perceivers without an evolutionary history, neural states represent those features they covary with under 'normal' conditions, normal conditions being those in play when the creature originated. This allows us to suppose that if Tye's physical duplicate, Swamp Tye, materialized on Earth, put on colour-inverting goggles, and travelled to Inverted Earth, he would misrepresent the Inverted Earth grass as being green, as it was on Earth, rather than correctly representing it as being red.

For various reasons, it would be better to abandon the optimal-or-normal-conditions clause. Macpherson (2002) has pointed out that it creates a mismatch between Tye's externalism about perceptual content and ordinary externalism about belief content, which claims that the content of a person's beliefs is fixed by the linguistic community they are presently embedded in. If both were true, it would mean that the Inverted Earth subject sees the grass on Inverted Earth as green, but believes that it is red, and also mistakenly believes that he sees it as red. Again, this in an unpalatable consequence.

A more general objection to the optimal-or-normal-conditions clause is that it has been heavily tailored to allow the theory to avoid the threats posed by the Inverted Earth and Swamp Man scenarios, and for this reason is ad hoc. A case could be made that introducing such a clause specifically to avoid the counterexamples is a perfectly legitimate move, since being internally consistent and free of problematic consequences is itself a good reason to endorse the account, even in the absence of independent motivation. If this is the aim, however, it would certainly be better to abandon the optimal-or-normal-conditions clause if a more parsimonious account could be found.

According to the sensorimotor theory, the only way colour-inverting goggles could cause a perceiver to have the colour experiences she had on Earth on Inverted Earth is if they somehow instantiated the SMCs—or at least, a relevant subset of the SMCs—that applied on Earth (see e.g. Hurley 1998 for a defence of this claim). In this event, the sensorimotor approach, on a representationalist reading, does not need to appeal to optimal or normal conditions to individuate content. Instead, it just appeals to covariance between neural states and the SMCs that obtain in the present environment. On the basis that the goggles on Inverted Earth must preserve appropriate SMCs from Earth to preserve the phenomenal character, we need not say that the subject's brain on Inverted Earth misrepresents a green apple as having the intrinsic properties that make an object red. Instead, we suppose that his brain correctly represents that he stands in a sensorimotor relation to an object characteristic of the colour red. This allows us, importantly, to abandon the mention of optimal or normal conditions.

VII. WHY WE SHOULD ABANDON 'REPRESENTATION' TALK ENTIRELY

This paper has set out to show that the sensorimotor theory makes useful progress in a well-worn debate about how phenomenal qualities should be naturalistically explained and individuated. In particular, I have been attempting to show that the sensorimotor approach to phenomenal character ought to have appeal even for those who are resistant to unorthodox anti-representationalist approaches to cognitive science. Moreover, since the sensorimotor approach faces apparent pressure to concede the existence of internal representation, it is worth noting that endorsing internal representation would not fundamentally endanger the theory. That said, appreciating the limited explanatory work that appeals to internal representation as presently conceived actually do leads us towards a more unorthodox conclusion.

Tye's account aims to naturalize phenomenal qualities by showing that they can be reduced to representational content. However, his account appears to just assume that naturalizing representational content itself presents no deep problem. In fact, there are strong reasons to doubt that representational content can be naturalized at all, least of all by means of a simple covariance relation. Consider the rings of a tree trunk. Hutto & Myin (2013) observe that the mere fact that the rings reliably indicate the tree's age (assuming they do) does not entail that they bear content (i.e. truth conditions) about the tree's age. Hutto and Myin's argument, in condensed form, is that the burden of proof is on endorsers of covariance-based accounts of content to show that covarying in some way with something is the same as bearing content about it, and that they have no apparent resources to meet this burden. The best they can do is *stipulate* that content is covariance, an option Hutto and Myin suggest violates naturalism.

It does not, at present, matter whether or not we oppose ourselves as a matter of metaphysical principle to 'content' and 'representation' talk. For some explanatory purposes, it might be OK simply to stipulate that a covariance relation of a particular kind is sufficient for content, for instance if doing so is a helpful way to make the behaviour of a system intelligible. However, the project of naturalizing phenomenal qualities is itself motivated by deep-seated metaphysical concerns—we are trying to find out what kinds of things phenomenal qualities really are, contending in particular with the background threat of anti-physicalism (Jackson 1982; Chalmers 1996). We cannot solve this problem by appealing to representational content where the content only exists because we say it does. To do the metaphysical work that representational content seems to do when it comes to naturalizing phenomenal qualities, the content must in a robust sense be real.

Add to this the fact that pure covariance-based accounts of content like Tye's (see also e.g. Rupert 2011) are extraordinarily thin. For Tye, the covariance relation is all that is needed, with no essential roles, for instance, for teleology or isomorphisms between the neural structures and the features represented. Considering how thin Tye's account of content is, and the fact that it is does not serve well as a metaphysically robust account of content, there is a strong case for avoiding the use of 'representation' and 'content' talk when attempting to explain phenomenal character. The remedy, I suggest, is not to thicken up the notion of content in play, but to be explicit about the metaphysically thin role that the alleged 'representations' play.

Suppose on this basis that we introduce a change of terminology. Instead of representing SMCs, let us say that brains are 'attuned' to SMCs, meaning they activate in a different way when different patterns of SMC apply, but do not bear content about them (for similar suggestions, see O'Regan and Block 2012; Myin and Degenaar 2014). Representationalists and anti-representationalists in vision science can agree that there are neural states that systematically covary with other physical states, and that such states play a causal role in perception and other forms of mindedness. The dispute does not turn, at least in an immediate way, on facts that can be discovered empirically, but on the concepts we use to describe the physical processes that take place.

The choice of terminology is important, as the shift from 'representation' talk to 'attunement' talk invites a significant shift of emphasis. In representationalist accounts of perception and phenomenal character such as Tye's, the important explanatory work is done almost entirely by the representations themselves, and the question of what the representations represent is barely discussed. This is unsatisfying considering that in a thin covariance-based account of content like Tye's, the content is individuated entirely by properties of the extension and not by any particular property of the vehicle. The sensorimotor account of phenomenal character appeals primarily to SMCs, which are the core of the theory. In addition, it observes that the brain must, as I am current proposing, be attuned to them. Because attunement is an overtly thin notion, we have no choice but to make central appeal to the properties to which the perceiver or her brain is attuned.

In this paper, I have not addressed the conditions other than attunement to SMCs that are required for a perceiver to undergo perceptual experiences. Tye, for instance, claims that in addition to representing the outside world, the subject must be poised to use the representations for belief formation. O'Regan's (2011) version of the sensorimotor theory says something broadly similar, namely that the subject must be poised to make use of her attunement to SMCs for the purpose of discursive thought or highly reasoned behaviour. More should be said about the nature of these capacities, and such a discussion is beyond the scope of this paper: however, we should bear in mind that there

is no reason to presume that such capacities necessarily depend on neural representation.

VIII. CONCLUSION

This paper has argued that representationalist approaches to perception and phenomenal character are compatible with the sensorimotor theory. More importantly, I have argued that there is positive reason for representationalists about perception and phenomenal character to endorse the sensorimotor account of phenomenal character. The sensorimotor account of phenomenal qualities has various explanatory advantages, and SMCs are worth appealing to in particular for endorsers of representationalist accounts such as Tye's because they offer a parsimonious way to respond to the Inverted Earth problem identified by Block (1990) without the need for Tye's problematic 'optimal-or-normal-conditions' clause. The theme underlying all this is that you need not subscribe to an unorthodox anti-representationalist approach to cognitive science to find reason to endorse the insights of the sensorimotor account of phenomenal character, notwithstanding the radical tenor with which the sensorimotor theory is often presented. Nonetheless, once we realize that SMCs present in perceivers' bodily interactions with the environment are doing most of the interesting explanatory work, and that the notion of internal representation, in addition to being metaphysically tenuous, is doing relatively little, we will find that there is reason to reformulate the concept of internal representation so that we are not mislead about its proper explanatory role.[5]

REFERENCES

Berlin, B. and Kay, P. (1969) *Basic Color Terms: Their Universality and Evolution*. Berkeley and Los Angeles: University of California Press.
Block, N. (1990) 'Inverted Earth', *Philosophical Perspectives*, 4: 53.
—— (1996) 'Mental Paint and Mental Latex', *Philosophical Issues*, 7: 19.
—— (2009) 'Comparing the Major Theories of Consciousness', in M. Gazzaniga (ed.) *The Cognitive Neurosciences IV*. Cambridge, MA: The MIT Press.
Brooks, R. (1991) 'Intelligence Without Representation', *Artificial Intelligence*, 47 : 139–59.
Buhrmann, T., Di Paolo, E.A. and Barandiaran, X. (2013) 'A Dynamical Systems Account of Sensorimotor Contingencies', *Frontiers in Psychology*, 4.
Chalmers, D. J. (1996) *The Conscious Mind: In Search of a Fundamental Theory*. Oxford: OUP.
Clark, Andy (2008) *Supersizing the Mind: Embodiment, Action, and Cognitive Extension*. Oxford: OUP.
Clark, Austen (2000) *A Theory of Sentience*. Oxford: Clarendon Press.
Dennett, D.C. (2002) 'How Could I Be Wrong? How Wrong Could I Be?', *Journal of Consciousness Studies*, 9(5–6): 17–28.

[5] Thanks to Jan Degenaar, Kevin O'Regan, Simon Prosser, Mike Wheeler and an anonymous referee for helpful comments on earlier versions of this piece. This work was financed by ERC Advanced Grant 323674 'FEEL'.

Hurley, S. (1998) *Consciousness in Action*. Cambridge, MA: Harvard University Press.
Hutto, D.D. and Myin, E. (2013) *Radicalizing Enactivism: Basic Minds without Content*. Cambridge, MA: The MIT Press.
Jackson, F. (1982) 'Epiphenomenal Qualia', *Philosophical Quarterly*, 32(127): 127–36.
Locke, J. (1975) 'Of True and False Ideas', in P.H. Nidditch (ed.) *An Essay Concerning Human Understanding*. Oxford: OUP.
Macpherson, F. (2002) 'Review of M. Tye "Consciousness, Color and Content", (Cambridge MA: MIT Press, 2000)', *Philosophical Quarterly*, 53(213): 619–21.
Marr, D. (1982) *Vision*. Cambridge, MA: The MIT Press.
Milner, D. and Goodale, M. A. (1995) *The Visual Brain in Action*. Oxford: OUP.
Minsky, M. (1986) *The Society of Mind*. New York: Simon & Schuster.
Myin, E. (2001) 'Color and the Duplication Assumption', *Synthese*, 129(1): 61–77.
——Degenaar, J. (2014) 'Enactive Vision', in L. Shapiro (ed.) *The Routledge Handbook of Embodied Cognition*, 90–8. London: Routledge.
Noë, A. (2002) Is the Visual World a Grand Illusion? *Spec. Issue of Journal of Consciousness Studies*. 9(5–6): 1–12.
—— (2004) *Action in Perception*. Cambridge, MA: The MIT Press.
—— (2012) *Varieties of Presence*. Cambridge, MA: Harvard University 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): 461.
—— (2011) *Why Red Doesn't Sound Like a Bell: Understanding the Feel of Consciousness*. Oxford: OUP.
——Block, N. (2012) 'Discussion of J. Kevin O'Regan's "Why Red Doesn't Sound Like a Bell: Understanding the Feel of Consciousness"', *Review of Philosophy and Psychology*, 3(1): 89–108.
——Noë, A. (2001) 'A Sensorimotor Account of Vision and Visual Consciousness', *Behavioral and Brain Sciences*, 24(5): 883–917.
Palmer, S. E. (1999) 'Color, Consciousness, and the Isomorphism Constraint', *Behavioral and Brain Sciences*, 22(06): 923–43.
Pettit, P. (2003) 'Looks As Powers', *Philosophical Issues*, 13(1): 221–52.
Philipona, D.L. and O'Regan, J.K. (2006) 'Color Naming, Unique Hues, and Hue Cancellation Predicted From Singularities in Reflection Properties', *Visual Neuroscience*, 23(3–4): 331–9.
Roberts, T. (2010) 'Understanding "Sensorimotor Understanding"', *Phenomenology and the Cognitive Sciences*, 9(1): 101–111.
Rupert, R.D. (2011) 'Embodiment, Consciousness, and the Massively Representational Mind', *Philosophical Topics*, 39(1): 99–120.
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): 97–118.
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*. Cambridge, MA: The MIT Press.
Whitham, E.M. et al. (2011) 'Visual Experiences During Paralysis', *Frontiers in Human Neuroscience*, 5. http://doi.org/10.3389/fnhum.2011.00160.

Paris Descartes University, France