

Bodily skill and internal representation in sensorimotor perception

David Silverman¹ 

© The Author(s) 2017. This article is published with open access at Springerlink.com

Abstract The sensorimotor theory of perceptual experience claims that perception is constituted by bodily interaction with the environment, drawing on practical knowledge of the systematic ways that sensory inputs are disposed to change as a result of movement. Despite the theory's associations with enactivism, it is sometimes claimed that the appeal to 'knowledge' means that the theory is committed to giving an essential theoretical role to internal representation, and therefore to a form of orthodox cognitive science. This paper defends the role ascribed to knowledge by the theory, but argues that this knowledge can and should be identified with bodily skill rather than representation. Making the further argument that the notion of 'representation hunger' can be replaced with 'prima facie representation hunger', it concludes that although the theory could optionally be developed scientifically in part by reference to internal representation, it makes a strong and natural fit with anti-representationalist embodied or enactive cognitive science.

Keywords Enactivism · Sensorimotor theory · Sensorimotor contingencies · Perception · Internal representation

1

The sensorimotor theory (Hurley 1998; O'Regan and Noë 2001; Noë 2004; O'Regan 2011) is an unorthodox but increasingly influential approach to understanding perceptual consciousness. One important aspect of the theory is its account of phenomenal character. It claims that the phenomenal qualities of perceptual experience are determined by and explicable by reference to sensorimotor contingencies (SMCs), the ways that sensory inputs are disposed to systematically change in line with movement. SMCs

✉ David Silverman
davsil@gmail.com

¹ H440, Laboratoire Psychologie de La Perception, Université Paris Descartes, 45 rue des Saints-Peres, 75006 Paris, France

are partly determined by sense organs, and O'Regan and Noë (2001, henceforth O&N) propose that differences in the way e.g. eyes and ears determine SMCs accounts for inter-modal phenomenal differences, e.g. between vision and audition. SMCs are also determined by objects, and O&N propose that differences in the ways that objects with differing e.g. shape or reflectance profile determine SMCs account for intra-modal differences, e.g. between the looks of different colours.

Noë (2004) expands the notion of SMCs to include the ways that perspectival or 'P'-properties' change in line with movement. A P-property is the way an object appears from a given perspective, for example the elliptical aspect of a circular dinner plate when faced at an angle. To see the plate as objectively circular is to see it as having the tendency to appear as a more or less eccentric ellipse as the angle you face it from changes. This explains how we experience objective properties even though we always perceive the environment from a particular vantage point.

While SMCs help explain why an experience has one phenomenal character rather than another, an account must also be given of how we come in the first place to have perceptual experiences characterised by SMCs. This task is addressed by the sensorimotor theory's accompanying account of perception, which is the focus of this paper. The theory purports to reject the orthodox scientific view, which following Marr (1982), characterises perception as a process of internal representation by the brain of the outside environment. O&N and Noë (2004) state that perception is instead the possession and exercise of an embodied 'mastery' or highly implicit, practical (and not propositional) knowledge of SMCs.

This view of perception is open to two quite distinct interpretations. On one interpretation, a perceiver has sensorimotor knowledge by virtue of nothing other than the possession of bodily abilities, and perception is nothing other than the possession and exercise of these abilities. On this reading, the sensorimotor theory rejects the orthodox view's basic claim that perception is a process of internal representation.

On another interpretation, sensorimotor knowledge is an internal representation of SMCs, and perception is the deployment of this representation, an activity that is embodied only inasmuch as it is geared toward and causally enabled by bodily interactions. On this reading, the sensorimotor theory accepts that perception is internal representation, but rejects certain lines of thinking sometimes endorsed by orthodox theories, for example that the representations in question are pictorial (e.g. Kosslyn 1994), or represent space as a set of co-ordinates (a view countered by Terekhov and O'Regan 2013). On this latter reading, the theory is actively incompatible with 'radical' embodied or enactivist theories of perception and action (e.g. Gangopadhyay and Kiverstein 2009; Chemero 2009, 2016; Hutto and Myin 2013) which set out to dispense altogether with theoretical appeals to internal representation.

This uncertainty about the essential nature of sensorimotor perception is significant, in part, because of the bearing it has on attempts to develop the sensorimotor approach scientifically. A recent proposal by Seth (2014) offers a rendering of the sensorimotor theory along very much orthodox lines, suggesting that perception consists of the internal representation of SMCs by a predictive model. The theory has also been developed in a representationalist register by Maye and Engel (e.g. 2016) who use a representational implementation of sensorimotor knowledge in robot control architectures. Others, however, such as Flament-Fultot (2016) and Buhrmann et al. (2013), have developed the theory by appeal to entirely non-representational dynamical systems.

While these differences reflect differing ideas about the most empirically fruitful way to understand how sensorimotor perception is implemented, in so doing they also track different understandings of the sensorimotor theory's core commitments.

This paper will mainly focus on what McDowell (1994b) calls perception's 'constitutive' features, i.e. the features that are an essential part of what it is to perceive. If the exercise of sensorimotor knowledge is the exercise of a bodily skill (i.e. ability), it is not a process of internal representation, just as riding a bike is not a process of internal representation, even though internal representations could play a role in enabling bike-riding to take place. Conversely, if the exercise of sensorimotor knowledge is the deployment of an internal representation, it is not the possession and exercise of a bodily skill, even though it may enable and be enabled by various kinds of skillful bodily interaction.

I will advocate a skill- rather than internal representation-based constitutive account of sensorimotor knowledge, and so perception. The first part of the paper will explain why we should be motivated in the first place to construe perception as the possession and exercise of bodily skills. The next part of the paper will explain how sensorimotor knowledge can be intelligibly construed as a bodily skill rather than a kind of internal representation, and hence why the sensorimotor theory is compatible with radical anti-representationalism. Finally, I will attempt to clarify the uncertain role of internal representation in the sensorimotor theory by setting out limited respects in which theoretical appeals to internal representation are admissible, although not essential, in developing the sensorimotor approach.

2

Let's begin by seeing what motivations there are for construing perception as the possession and exercise of bodily skills, and so endorsing what I will refer to as the 'skill-based' view of perception.

2.1

One important reason to endorse the skill-based view is that it offers the best way to explain the visual feeling of presence of a richly and uniformly detailed scene. The retinal image does not at any given instant contain rich and uniform detail, and work on phenomena such as change blindness (e.g. O'Regan et al. 1999) has suggested that the brain probably does not build up a richly detailed inner model to compensate. However, as O'Regan (1992) among others has emphasised, perceivers access detail directly from the environment by making constant use of exploratory bodily movements such as saccades, and the body is poised to re-orient the eyes toward sudden changes in stimulation occurring at the low-resolution periphery of the retina (this is called 'grabiness' e.g. by O&N). Thus he makes the compelling suggestion that we access detail directly from the environment, when required, in a serial fashion.

Dennett (2002) argues that visual experience does not really present rich detail, and that the ability to interrogate the environment in the way described causes perceivers to

be standardly under the grip of a ‘grand illusion’, falsely believing their own visual experiences to be much richer in detail than they really are. This is analogous to thinking the refrigerator light is always on because it is on whenever you open the door to look (O&N). On this proposal, bodily skill accounts for the false belief that perception is rich, while perception itself may be constituted by nothing other than internal representation.

Noë (2002) endorses O’Regan’s (1992) suggestion that bodily skill explains why visual experience seems to present rich and uniform detail. Instead of claiming that this is an illusion, however, he argues that visual experience really does present rich detail. He claims that the experience of richness is not constituted by internal representation, but by the detail’s accessibility, i.e. its readiness to be accessed using skillful movements. This proposal has a significant advantage, namely that it is more parsimonious, since it does not require us to suppose that perceivers are as standard radically deceived about the nature of their own experiences. Importantly, however, it requires us to conceive of the sensorimotor skills we use to access detail as constitutive of visual experience rather than mere causal scaffolding.

The phenomenology of richness therefore offers a powerful reason for thinking that perception is at least partly constituted by bodily skills. This does not preclude the possibility that sensorimotor knowledge is an internal representation, and perception the deployment of this representation plus the exercise of bodily skills. The reasons I give next for endorsing the skill-based view will call on us to take sensorimotor knowledge itself to consist of bodily ability.

2.2

Another important reason to endorse the skill-based view is to help secure the philosophical thesis of direct realism. Direct realism claims that to perceive is to come into direct and unmediated epistemic contact with the objects in the environment rather than with representations of them (e.g. McDowell 1994a, b). Although it is contentious, direct realism has the crucial advantage of according with perceptual phenomenology, which *feels* direct and unmediated. This is an important reason to endorse direct realism, provided the thesis can be shown in other respects to be sufficiently plausible.

Beaton (2016) has compellingly argued that the skill-based view of perception can be considered a scientifically-tractable accompaniment to direct realism, given that exercising a bodily skill is itself a way of interacting in a direct and unmediated way with the outside environment. Whereas the orthodox scientific view of perception does nothing to support direct realism, the skill-based view, Beaton notes, lends direct realism scientific credentials and has the advantage of doing justice to the phenomenology of directness which motivates it.

The sensorimotor theory also offers some reasons of its own to endorse direct realism. As I explained in the last passage, the skill-based view is partly motivated by the phenomenology of richness. To further enrich the story, we can observe that this phenomenology also helps motivate direct realism. The idea that there is a grand illusion is persuasive if you subscribe to the indirect realist thesis that one’s perceptual access to the world is mediated by a representation, since such a representation would

probably be sparse in detail. If you are an indirect realist, you therefore come under pressure to accept the unparsimonious thesis that our beliefs about what our perceptual experiences are like accord with what the world is like, but not with what our perceptual experiences are really like.¹

Noë's account of richness is compelling because it hints at a simpler explanation: we have the visual experience of being situated within a richly detailed environment because being in an appropriate way situated is just what vision consists of. This is a direct realist claim. Hence the skill-based account of visual richness supports direct realism, and direct realism in turn supports a skill-based account of the entirety of visual experience (as argued by Beaton 2013, 2016).

A final reason to endorse direct realism is derived from the sensorimotor theory's broad claim that perceptual experience depends on mastering patterns of sensorimotor contingency. Noë (2008) grants that the theory must account for the fact that perceptual experience seems to present what Campbell describes as the "ordinary world, there independent of us" (Campbell 2008, p. 667) as opposed merely to counterfactual consequences of possible movement. To this end, Noë denies that sensorimotor contingencies are the objects of perception, claiming instead that we perceive the ordinary world *by* mastering SMCs.

To pursue this solution, a version of the sensorimotor theory which claimed that perceptual experience is mediated by a representation would need, I suggest, to propose that we use sensorimotor knowledge (whether construed as representation or skill) to construct a separate representation of the ordinary world, since it could only explain how we perceptually experience the world as a world of ordinary objects by appealing to the deployment of such a representation. According to Noë's more elegant proposal, experiencing the ordinary world does not require this extra step, since exercising sensorimotor knowledge is enough by itself to put us in direct perceptual contact with the ordinary world. We can explain how sensorimotor knowledge does this by positing, as just suggested, that sensorimotor knowledge is a bodily skill used to interact in an unmediated way with the outside environment.

2.3

One final reason to adopt the skill-based view of perception is that, on account of avoiding a serious shortcoming with representationalist accounts of phenomenal quality, it offers a better way to naturalise the phenomenal qualities of perceptual consciousness.

¹ I am taking it that visual experiences do indeed seem to present a large and uniformly detailed visual scene, i.e. that we believe they present such a scene. This is supported by the surprise expressed by subjects in change blindness and inattentive blindness experiments upon discovering that they have missed unusual occurrences taking place in full view (see Dennett 2002). Dennett claims that this belief is illusory. The sensorimotor theory, on a skill-based interpretation, can claim that the belief is correct, since presence is not representation but accessibility. An indirect realist could avoid the claim that there is a grand illusion by denying that we have any such belief. To account for the surprise, she could say that subjects correctly believe that detail is usually accessible using bodily skill, but deny that this accessibility is a kind of perceptual presence, hence a constitutive part of visual experience. However it is, I believe, phenomenologically more apt to claim that momentary perceptual experiences do *seem* like they present a large and uniformly detailed scene.

The problem with internal representation, in short, is the absence of a satisfactory way to naturalise content i.e. truth conditions.² Hutto and Myin (2013, henceforth H&M) observe that the leading attempts to naturalise content (e.g. Dretske 1981; Millikan 1989) depend on the idea that standing in a co-variance relation with something is identical to bearing content about it, at least when certain conditions are met. And as they comment, co-variance by default does not constitute content - for instance, a tree's rings do not strictly speaking bear truth conditions about the tree's age - and the burden of proof is therefore on the theorist intending to show that in certain cases it does.

The problem, as H&M argue, is that nothing has yet been done to adequately meet this burden, either by appealing solely to particular types of co-variance or to co-variance with the addition of further theoretical apparatuses. Millikan (1989), for example, appeals to the proper function brains have of possessing and making appropriate use of co-variance relations. However, this merely shows that brains have the function of exploiting patterns of co-variance with the environment and does not establish that they bear content about it. H&M conclude that cognitive scientific accounts of action and perception should dispense altogether with the idea of content and consequently representation.

I will not argue for the radical conclusion that H&M do, since even though truth conditions have not so far been naturalised, many proponents of what H&M call 'content involving cognition' have plausible rationales for appealing to representation or content that do not depend on doing so. One approach is simply to stipulate that covariance of a certain kind constitutes content (e.g. Rupert 2011). This entails deflating the notion of content so that all we mean by content is such a relation. It can be granted that the brain makes use of certain covariance relations for producing behaviour.

An alternative approach accepts that content consists of truth conditions, but treats these as a kind of explanatory fiction (e.g. McDowell 1994b; Sprevak 2013; Egan 2014). For the sake of argument, we can grant that physical states sometimes interact with one another to produce behaviour in a way that can best be predicted or made intelligible by reference to the fact that they are interacting as if they were contentful representations.

Representationalism about phenomenal qualities (e.g. Tye 2000), however, has a distinctive explanatory aim, namely to render unmysterious the phenomenal properties often considered to be irreducible or even non-physical 'qualia'. Borrowing from orthodox perception science the assumption that perception consists of internal (e.g. neural) representation, this variety of representationalism proposes that phenomenal qualities such as the looks of colours can be reduced to the content of the representations that yield the experience. In so doing, representationalists like Tye take themselves to be appealing to the genuine existence of internally (e.g. neurally) borne truth conditions.

The problem is that this kind of representationalism, though it may adequately reduce phenomenal properties to content, makes no useful progress toward convincingly identifying truth conditions with types of physical relation. While naturalising truth conditions is not necessarily important for the scientific project of explaining behaviour (including perceptually-guided behaviour, reports of visual experience, etc.) in representational terms, it is essential to the more philosophical project of using representation to naturalise qualia.

² In this passage, 'truth conditions' will be used for the sake of brevity to mean 'truth and/or correctness conditions'.

We might entertain that representationalism about phenomenal qualities could make use of the stipulative or fictionalist views of content just described. The problem is that attempts to account naturalistically for the intentional directedness and phenomenal character of perceptual experience are subject to greater constraints than attempts to explain behaviour. Identifying phenomenal qualities directly with features of the brain, representationalists agree, would not render phenomenal qualities unmysterious, so stipulating that content is merely a kind of brain state will not help.

Moreover, we cannot reduce phenomenal qualities to content if that content is fictional, since this would amount to reducing the qualities to something non-existent. One might point out that fictional content does exist, in the sense in which Sherlock Holmes exists. Even if we grant this, it will only entail that content exists relative to a theoretical model, and not that it exists from the point of view of a conscious subject. And it is precisely the qualities featured in the first-person experience of the conscious subject that we are trying to identify with something physically-realizable. So representational content will not help account for perceptual phenomenology unless it is 'real' in a more robust sense than has yet been convincingly established.

Bodily skills, by contrast, are not explanatory fictions. Their exercise can be straightforwardly observed in the behaviour of perceiving subjects. The shortcomings of attempts to naturalise subpersonal representational content therefore offer another reason to endorse the view that sensorimotor knowledge consists of bodily skills.

3

The considerations just outlined give us good reason to construe perception as the possession and exercise of bodily abilities. Since we are taking it that perception is the possession and exercise of sensorimotor knowledge, we therefore have good reason to construe sensorimotor knowledge as bodily ability. O'Regan and Noë (2001) accordingly describe sensorimotor knowledge as a variety of practical knowledge (i.e. knowledge-how) and, expanding upon this point, Noë (2004, 2005, 2014) argues, contra Stanley and Williamson (2001), that practical knowledge in general consists of ability (being able to perform an action) and not propositional knowledge (knowing that such-and-such a movement is the way an action is performed). In what follows, I will assume that Noë's arguments against Stanley and Williamson are correct.

Here, however, we meet a problem: namely that this characterisation of sensorimotor knowledge is in tension with the sensorimotor theory's compelling account of phenomenal qualities. O'Regan and Noë claim that phenomenal qualities of perceptual experience, e.g. the looks of colours, are determined by properties which can be specified by counterfactual conditionals, namely the sensory changes that would occur if certain movements took place. To explain how we come to have perceptual experience determined by SMCs, the authors claim that we exercise "practical knowledge of the effect movement will have on nervous influx" (pp. 970–971). But this claim seems incoherent. The object of knowledge-how must be a doing (e.g. knowing how to ride a bike) and the sensory changes that would occur as a result of possible movements are not doings. The upshot, on the face of it, is that we can either construe sensorimotor knowledge as practical but not counterfactual, or counterfactual but not practical, but not both. And we will now see that neither option is satisfactory.

Suppose that sensorimotor knowledge is practical but not counterfactual. This would mean that perception is constitutively dependent solely on abilities to perform bodily actions of one kind or another. In consequence, there would be no obvious way to make it intelligible that the phenomenal character of perception is determined by sensory changes that would occur as a result of merely possible movements, and we would need an alternative way to explain how bodily interactions determine the phenomenal character of perceptual experience. Alternative principles are available, but have important shortcomings.

According to one available alternative, phenomenal qualities are determined entirely by the sensory changes that are now occurring as a result of presently ongoing movements. However, this suggestion fails to accommodate one of the sensorimotor theory's most prominent explananda, namely the perceptually experienced presence 'in absence' of hidden or occluded features of the perceived scene (e.g. Noë 2004; O'Regan 2011). Noë observes, for instance, that only one side of a tomato or other opaque object can be seen at any one time, and yet that it does not typically feel as if you are looking at a mere tomato facade. The tomato's hidden side is perceptually present *as absent*, he suggests, because you appreciate that certain movements would result in the hidden side coming into view. Importantly, this experience only requires that you look at the tomato: you do not have to actually interact with the tomato's hidden side. The experience therefore cannot be accounted for by the sensory changes taking place as a result of presently occurring movements.

According to another available principle, the phenomenal character of perceptual experience is determined by possibilities for action, e.g. affordances of the sort described by Gibsonian accounts (for recent examples see Chemero 2009, Rietveld and Kiverstein 2014). By adopting this principle we could explain perceptual presence, including presence in absence, by appeal to your knowing how to make use of action possibilities, including the possibility of interacting with the tomato's hidden side. However, while it is beyond the scope of this paper to give a detailed comparison of SMCs-based and affordance-based accounts of perceptual phenomenology, I wish to highlight one important apparent shortcoming with this proposal.

Subjects in experiments using left-right inverting goggles, upon first putting the goggles on, are unable to see and behave as normal, but after an extended period of adaptation are able once again to effectively co-ordinate their bodily movements. Reports vary, but many suggest that visual phenomenology remains inverted even after adaptation (see, e.g., Klein 2007, Degenaar 2014).³ The sensorimotor theory, appropriately understood, can accommodate this straightforwardly by noting that putting on the goggles changes the SMCs that presently obtain, and that they remain changed whether or not the subject is adapted to them. The affordance-based approach cannot accommodate the inverted phenomenology, because the adapted goggle-wearer makes use of the same possibilities for action she made use of while not wearing the goggles.

³ Degenaar's view, based on his own first-hand experience with inverting goggles, is that perceptual experience has an intentional component which re-inverts after adaptation to the goggles, and a distinct phenomenal component which does not.

Given that there are important drawbacks to supposing that sensorimotor knowledge is practical but not counterfactual, let us suppose that it is counterfactual but not practical. The task, if we adopt this construal, is to explain what sensorimotor knowledge consists of. The most obvious proposal, consistent with orthodox perception science, is that it is a variety of propositional knowledge and consists of internal representation. Conversely, it is hard to see how it could consist of ability, because an ability, if a kind of knowledge, must be knowledge-how. The problem, of course, is that we are attempting to avoid a conception in which perception is constitutively dependent on internal representation. This makes it hard to give any positive characterisation of sensorimotor knowledge.

This dilemma will seem even more acute if we endorse Noë's claim that sensorimotor knowledge is "the ground of your possession of dispositions to respond to the presented object" (Noë 2004, p. 88). If H&M's interpretation is correct, this positively rules out the possibility that sensorimotor knowledge is itself a disposition (or ability). They say:

Whereas the endorsement of representationalism is coherent but costly, without that commitment there appears to be no way to specify in positive terms what kind of thing ... [sensorimotor] knowledge is meant to be, given that it is not merely embodied in dispositions of organisms but is supposed to form the ground for such dispositions. (H&M, p. 26)

Thus it is hard to give a positive characterisation of sensorimotor knowledge as counterfactual but not practical without construing it as an internal representation, and even harder if we also accept the suggestion that sensorimotor knowledge, being a ground, is not embodied in the dispositions of organisms. I believe that it is right to view sensorimotor knowledge, being counterfactual, as a ground of bodily abilities, because various bodily abilities require you to appreciate the sensory consequences of possible movement. However this need not, as it first appears, require denying that sensorimotor knowledge is a kind of practical knowledge, embodied in the dispositions (or, more aptly, abilities) of perceiving subjects. The next section will explain how it can after all be so-understood.

4

Sensorimotor knowledge consists, I propose, of the ability to act in a way that is sensitive to the SMCs that presently obtain. A subject could not logically be said to have sensorimotor knowledge, on this view, if she did not have abilities to perform various bodily actions. However, sensorimotor knowledge is not identical to such abilities either. While knowledge of how to perform a particular action, such as ride a bike, is individuated by a behavioural intention, sensorimotor knowledge is individuated by SMCs to which the subject is sensitive when she acts.

Sensorimotor knowledge is a condition of possibility for most action abilities: for example, you would not be able to catch a ball if you were not able to act in a way that is sensitive to the changes in sensory stimulation that movement by the ball or changes to your own bodily positioning will produce. In this sense, sensorimotor knowledge can

rightly be said to ‘ground’ many action abilities. This does not mean, however, that sensorimotor knowledge is something other than a bodily ability, such as an internal representation: it is, rather, *an ability which grounds other abilities*.

It is necessary to distinguish between sensorimotor knowledge and the differently, usually more narrowly, individuated action abilities it grounds because it is sensorimotor knowledge that accounts for perceptual phenomenology. Two subjects could have differing sets of action abilities and goals while sharing identical sensorimotor knowledge, and the sensorimotor theory, as presently understood, predicts that sameness of perceptual phenomenology requires sameness of sensorimotor knowledge and not sameness of ability or inclination to carry out specific goal-directed actions.

If we adopt this characterisation of sensorimotor knowledge, the sensorimotor theory agrees with H&M and Chemero (2016) that perceptual ability is not constituted by internal representation. Those authors also object to the very idea that perception requires ‘knowledge’, deeming this unacceptably intellectualist. This issue, though not trivial, does not have a material bearing on the points made so far: those persuaded, for instance, by Noë (2004, 2012) remarks defending the idea that perceptual abilities are a kind of practical knowledge or understanding can read the present treatment as a gloss on the faculty he describes. Those moved by anti-intellectualist considerations to deny that perception requires any kind of knowledge might prefer to expunge all references to knowledge and give the ability in question an alternative label such as sensorimotor ‘attunement’.

My position, for what it’s worth, is that the ability in question deserves to be described as ‘knowledge’ in adult humans and any other subjects that are able to have thoughts which express knowledge about what they perceive. Qua knowledge, the ability is both knowing *how* to act with the relevant sensitivity, and, in a highly implicit sense, knowing *that* certain SMCs apply. To know that you are seeing, for example, a cube, you must know, at least implicitly, that your relationship with the environment is governed by the SMCs characteristic of seeing cubes, since your belief that you are seeing a cube would otherwise lack justification. Someone could be said to have this knowledge by virtue of being able to recognise and report cases in which an object’s appearance does not vary in line with movement in the way expected.

Supposing, as seems plausible, that knowing what you perceive is an essential feature of the kind of perception enacted by adult humans, this kind of perception itself requires knowledge of SMCs. Appealing here to thought does not mean appealing to internal representation as long as we consider thought itself to be constituted by a relation a practically-skilled subject stands in to the environment, as Noë (2012) for instance proposes, rather than the disembodied manipulation of mental representations. On the other hand, sensorimotor ‘attunement’ better describes perceptual ability in subjects that fail to integrate their perceptual skills with capacities for thought, including (conceivably) machines, infants, animals and even more basic organisms, all of whom may nonetheless perform actions (construed, in an undemanding sense, as purposeful doings) that exhibit a sensitivity to SMCs.

Our account of sensorimotor knowledge or attunement will need to appropriately accommodate the phenomenological claim, mentioned in the last section, that hidden or occluded features of the environment can enjoy perceptual presence *in absence*. If the phenomenological description given by Noë (2004) is correct, merely looking at a tomato is enough to visually experience the presence in absence of its hidden side. On a

representationalist construal of sensorimotor knowledge, one could account for this by suggesting that the perceiver's brain represents the sensory consequences that would occur if she were, say, to turn the tomato around. On the present construal of sensorimotor knowledge, the case requires some further explanation, bearing in mind that a perceiver does not necessarily overtly exhibit a sensitivity to the tomato's hidden side when she looks at the tomato.

Consider O'Regan and Noë's (2001) example of a guided-missile pursuing a plane. Suppose that the missile's goal is to keep the plane aligned in the centre of its sensor and that when the target instead appears toward one edge of the sensor, the missile makes the movements required to re-align the target in the centre.⁴ Given a suitable assumption about what the missile's goal is, one could observe in this behaviour an overt sensitivity to the sensory consequence that will occur if it makes this movement. When the plane appears in the centre of the missile's sensor, the missile makes no corrective movements. Considering this event in isolation, there is nothing to distinguish the missile from a different missile that has been programmed to maintain its course regardless of its sensor state. Yet the original missile is nonetheless exercising its sensorimotor attunement by not changing course, because it is making use of its sensitivity to the fact that changing its course would result in it no longer achieving its goal of keeping the plane aligned in the centre of the sensor.

Similarly, you may look at a tomato but lack any intention to pick it up, turn it around, and so forth. By analogy with the missile that merely maintains its course, it would not be incoherent to take it that you are nonetheless making use, in an entirely implicit way, of your full repertoire of sensorimotor knowledge. This could be captured by broad principles such as "picking up the tomato would not result in my achieving my present behavioural goal" but also narrower principles describing, for instance, the ways the stimulation you will receive from the tomato will vary as you turn it around. The proposal that you are always making implicit use of all your sensorimotor knowledge allows us to adequately account for the phenomenology of presence in absence.

5

I have now established, I hope, that sensorimotor knowledge can and should be understood as a variety of bodily skill. Let's now consider what theoretical role for internal representation may remain, given the version of the sensorimotor theory I have advocated. The thesis that perception is wholly or partly constituted by internal representation - call this 'constitutive representationalism' - is featured in the indirect realist view of perception (see, Snowdon 1992) and also in cognitive science as the view, tacit or otherwise, that perception and other cognitive processes depend on internal representation as a matter of conceptual necessity or even definition. The cognitive scientific version is identified and rejected, for instance, by Ramsey (2015) and Noë (2004). In the characterisation I have given, bodily abilities are sufficient for sensorimotor knowledge, so we have rejected an important motivation for constitutive representationalism.

⁴ It is worth noting, as an anonymous reviewer points out, that in reality a guided-missile should aim toward where the target is going to be, based on its present trajectory, rather than where it is now. But a simplified description of a guided missile's behaviour will do for present purposes.

Recall that McDowell (1994b) contrasts constitutive features with ‘enabling’ features. To ‘enable’ perception is to play a causal role in perception, in particular one where the cause takes place simultaneous to the effect but at a different explanatory level. ‘Enabling representationalism’, as we might call it, claims that perception is enabled by the subpersonal deployment of internal representations. You can deny that perception is a kind of internal representation without necessarily denying that it so-enabled. So it would not compromise the skill-based view to allow that perception might involve subpersonal representations. Notice how the skill-based view, by this light, inverts the usual order of priority, in which cognition is thought to be constituted by internal representation and merely enabled by bodily interactions (Adams and Aizawa 2001, cf. Hurley 2010).

Here, however, things get a little more complicated. McDowell would reject constitutive representationalism even if representational explanation at the subpersonal level were indispensable, because he takes perception necessarily to be a personal-level (or agent-level) activity. But this move requires signing up to McDowell’s own particular view of the mind, and sensorimotor theorists are not required to accept it (Noë 2004, for instance, doubts that there is a useful personal/subpersonal distinction to be made). If we cannot rely on the personal/subpersonal distinction to distinguish constitutive representationalism from enabling representationalism, then we must rely on the distinction between the necessary and the contingent, where constitutive representationalism takes perception to be necessarily representational and enabling representationalism takes it to only be contingently representational.

The case of presence in absence, illustrated by the tomato example, shows that perception involves responsiveness to absent features. Clark and Toribio (1994) argue that cases of perception and action like this are subject to ‘representation hunger’, meaning they could not intelligibly occur unless there was internal representation. If perception is subject to representation hunger, constitutive representationalism is correct.

Constitutive representationalism in turn compromises the skill-based view, since if the bodily skills deployed in perception require internal representation, then we must understand sensorimotor knowledge as an internal representation, and this gives us reason to think that the bodily skills that the internal representation enables are contingent rather than constitutive features of perception. As we have seen, there are independent reasons (e.g. direct realism) for endorsing a skill-based view of perception which do not derive from the nature of sensorimotor knowledge. But if our conception of sensorimotor knowledge entails that perception is subject to representation hunger, it makes the skill-based view less plausible rather than lending it support.

I wish to maintain that sensorimotor knowledge is not constituted by internal representation. To establish this, we need not rely on McDowell’s claim that perception is necessarily an agent-level phenomenon. Nor need we deny that perceptual experience has characteristics, such as presence in absence, that make it appear subject to representation hunger. Instead of taking these characteristics to imply representation hunger, we should take them merely to imply what we could think of as *prima facie* representation hunger. *Prima facie* representation hunger does not entail that internal representation is theoretically indispensable, merely that there is theoretical work to be done that could be done by internal representation.⁵

⁵ For a rebuttal of the notion of representation hunger, see also Degenaar and Myin (2014).

To see that sensorimotor knowledge merely entails *prima facie* representation hunger, and not representation hunger proper, consider that there are non-representational ways that sensorimotor knowledge could be implemented. These are explored by Flament-Fultot (2016), who proposes that we appeal to synergies found in complex systems. Crucially, his proposal does not abandon the idea that features with which we are not in sensory contact can enjoy presence in absence. Instead, it accounts for presence in absence by appeal to behavioural intentions we have toward such features, which can be explicated without appeal to internal representation.

Flament-Fultot gives the analogy of a coiled spring. The spring's future behaviour is determined by presently instantiated properties such as its stiffness, and changes in its present stiffness mean changes in its behavioural propensities. His suggestion is that one's behavioural intentions toward the back of a tomato depend on properties analogous to stiffness, which do not require representational content. Although the version of the sensorimotor theory I have recommended does not account for the phenomenal character of perceptual experience by reference to particular behavioural intentions, e.g. the intention to interact in a particular way with a tomato, it does do so by reference to an ability to act with sensitivity of obtaining SMCs embodied *in* our intentional behaviour. So Flament-Fultot's appeal to properties analogous to stiffness has the potential to show how sensorimotor knowledge is implemented.

The point I want to emphasise at present is that the sensorimotor theory rejects *constitutive* representationalism. To this end, it does not in principle matter whether perception does or does not happen to be enabled by the deployment of internal representations. So long as it is true that perception as construed by the sensorimotor theory does not *necessarily* draw on internal representation, we can be assured that the sensorimotor theory actively rejects constitutive representationalism.

One reason this matters is that it secures the sensorimotor theory's credentials as a scientific complement to direct realism. McDowell's argument shows that direct realism is not ruled out by the suggestion that perception makes use of internal representations, since internal representations may be recruited by the processes we use to directly encounter the outside environment. However, indirect realism says that perception is necessarily mediated by an internal representation. Once we understand the sensorimotor theory to actively reject constitutive representationalism, we see that, in contrast with the orthodox view of perception, the sensorimotor theory is actively incompatible with indirect realism, and this secures its status as a scientifically-tractable version of direct realism. Rejecting constitutive representationalism also means that the sensorimotor theory is compatible with 'radical' non-representationalism and therefore could be scientifically developed in a non-representationalist framework.

6

In the introduction, I mentioned that the constitutive account we give of sensorimotor perception will have a bearing on how the sensorimotor approach should be developed scientifically. One consequence of adopting a wholly skill-based constitutive account, as I have aimed to show in the last section, is that sensorimotor perception could be implemented without internal representation, though it need not be. This might be considered of little consequence for accounts that do refer to internal representation to

explain how sensorimotor perception occurs. Let me explain now how it does have some bearing, using Seth's (2014) treatment of the sensorimotor approach as a case study to illustrate some points that are also potentially applicable to other representationalist renderings of the theory.

Seth takes perceptual phenomenology to be determined not only by present sensorimotor activity but, as I have indeed suggested here, by SMCs construed in counterfactual terms. Adopting a predictive coding framework, he operationalises sensorimotor knowledge by identifying it with generative models containing predictions about the sensory changes that would follow from movement. Since Seth does not make use of the distinction between constitutive and enabling features of perception, or anything closely related, it is hard to tell if he takes perception to be identical to the deployment of these models or if he would allow that perception is a kind of skillful interaction with the environment, which the generative models he describes merely help implement. What I shall claim, in either case, is that it would be a mistake to treat the generative models Seth describes as if they were the whole story, even if we set to one side the distinctively philosophical considerations concerning qualia and direct realism discussed in section 2.

One of Seth's main claims is that sensorimotor knowledge can usefully explain a feature of grapheme-colour synaesthesia in which seeing words induces a non-veridical experience of colour. The colour experience induced by this kind of synaesthesia lacks what Seth calls 'subjective veridicality', meaning that the colour does not appear to the subject to be a genuine property of the outside world. Subjective veridicality is the result, he proposes, of 'counterfactually-rich' generative models, i.e. models representing a sufficiently large range of SMCs, and it is the failure to represent a sufficiently large range that distinguishes synaesthesia-induced colour experience.

Notice that the difference between ordinary and synaesthesia-induced colour experience could be accounted in the way suggested by Seth whether we take perception to be constituted by a predictive model, or take it to be constituted by bodily abilities - themselves instantiating counterfactually-rich sensorimotor knowledge - that are merely enabled by generative models. There are, however, other features of perceptual phenomenology that are better addressed if we adopt the latter view. Some of these, such as the phenomenal difference between sensory perception and inner senses such as interoception are mentioned by O'Regan and Degenaar (2014).

Right now, let's return to the visual presence of an expansive and uniformly detailed environment, as discussed in section 2. It is notable that this aspect of visual phenomenology is itself distinguished, as Seth claims subjective veridicality is, by a high level of counterfactual richness. Though you are not presently in sensory contact with uniform detail from an expansive visual scene, you visually experience the presence of an environment containing a rich multitude of possibilities for movement and associated sensory consequences. It is doubtful, even within the predictive coding framework, that the brain makes uniformly detailed predictions about all of these, in particular those sensed far into the retinal periphery, so an account which takes perception to be constituted by generative models is likely to be confronted by a hard to explain disparity between the relative sparseness of the predictions and relative richness of the visual world that experience apparently presents us with. The skill-based approach, as we have seen, can explain the phenomenology in question by reference to relatively sparse sensorimotor knowledge applied within a rich and stable external environment.

Returning to Seth's account, it is conceivably true that counterfactual richness accounts for subjective veridicality. In this event, it should be recognised that this is in the first instance a fact about the phenomenology of perceptual experience - an experience feels veridical because it feels counterfactually rich. The question of how an experience so-characterised occurs is secondary, and is at least in some cases better answered by reference to bodily ability than generative models. Seth's treatment is a useful hypothesis about the implementational details of sensorimotor perception, but must be read in this light. The possibility of accounting for perceptual phenomena by reference to non-representational and/or extra-bodily features should not be closed off by conflating the treatment with a constitutive account of what it is to perceive.

7

The sensorimotor theory is sometimes presented as if it were essentially as a variety of orthodox cognitive science and at other times essentially as a 'radical' embodied or enactive theory. I claim that there are certain key principles that any version of the sensorimotor theory should be committed to and which delimit the approach as a substantive theory and not just a loose collection of intersecting claims. These principles are compatible with varying attempts to develop the theory scientifically, both with and without internal representation, but they do shape in important ways the manner in which the theory should be scientifically cashed out.

Considerations concerning the phenomenology of richness, direct realism, and phenomenal qualities suggest that perception is constituted by the possession and exercise of bodily abilities rather than a process of internal representation. Although we need the notion of 'sensorimotor knowledge', or at least something closely related such as 'sensorimotor attunement', we can understand this as bodily ability rather than an internal representation. On this construal the theory is not guilty, as radical enactivists worry, of being committed to internal representation, but in fact becomes a natural ally of anti-representationalist cognitive science.

A worry remains that the understanding of sensorimotor knowledge on offer nonetheless entails that perception is subject to representation hunger, meaning it depends subpersonally on internal representation. Since the sensorimotor theory does not necessarily endorse McDowell's use of the personal/subpersonal distinction, we cannot rely on this distinction to establish that perception is not constitutively representational and must instead rely on the necessary/contingent distinction. The danger is that representation hunger means that representation is necessary, and therefore that perception is constitutively representational. However, I suggest that perception is only subject to *prima facie* representation hunger, since there are non-representational (e.g. dynamical systems) accounts available that could do the explanatory work that would otherwise have been done by internal representation. The availability of such approaches is enough to show that perception as construed by the sensorimotor theory is not constitutively representational.

The sensorimotor theory may be developed in a way that appeals to internal representation, so long as it is agreed that internal representation is not essential but is used merely as an explanatory tool to show how the sensorimotor skills of which perception consists are contingently enabled. On this view, the epistemic access

perception gives us to the world and the phenomenal quality of perceptual experience must be accounted for by reference to the possession and exercise of bodily skills rather than internal representation. Moreover, the possibility remains that the sensorimotor theory be developed scientifically in an entirely non-representationalist register.

Acknowledgements Thanks to Jan Degenaar, Kevin O'Regan and two anonymous reviewers for useful feedback on earlier drafts of this piece.

Funding This work was funded by ERC grant 323674 'FEEL'.

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

References

- Adams, F., & Aizawa, K. (2001). The bounds of cognition. *Philosophical Psychology*, *14*(1), 43–64. doi:10.1080/09515080120033571.
- Beaton, M. (2013). Phenomenology and embodied action. *Constructivist Foundations*, *8*(3), 298–313.
- Beaton, M. (2016). Sensorimotor direct realism: how we enact our world. *Constructivist Foundations*, *11*(2), 265–276.
- Buhrmann, T., Di Paolo, E. A., & Barandiaran, X. (2013). A dynamical systems account of sensorimotor contingencies. *Frontiers in Psychology*, *4*. <http://journal.frontiersin.org/article/10.3389/fpsyg.2013.00285/full>.
- Campbell, J. (2008). Sensorimotor knowledge and naïve realism. *Philosophy and Phenomenological Research*, *76*(3), 666–673.
- Chemero, A. (2009). *Radical embodied cognitive science*. Cambridge MA: MIT Press.
- Chemero, A. (2016). Sensorimotor empathy. *Journal of Consciousness Studies*, *23*(152), 138.
- Clark, A., & Toribio, J. (1994). Doing without representing? *Synthese*, *101*(3), 401–431. doi:10.1007/BF01063896.
- Degenaar, J. (2014). Through the inverting glass: first-person observations on spatial vision and imagery. *Phenomenology and the Cognitive Sciences*, *13*(2), 373–393. doi:10.1007/s11097-013-9305-3.
- Degenaar, J., & Myin, E. (2014). Representation-hunger reconsidered. *Synthese*, *191*(15), 3639–3648. doi:10.1007/s11229-014-0484-4.
- Dennett, D. C. (2002). How could I be wrong? How wrong could I be? *Is the Visual World a Grand Illusion? Spec. Issue of Journal of Consciousness Studies*, *9*(5–6), 17–28.
- Dretske, F. (1981). *Knowledge and the flow of information*. Cambridge MA: MIT Press.
- Egan, F. (2014). How to think about mental content. *Philosophical Studies*, *170*(1), 115–135. doi:10.1007/s11098-013-0172-0.
- Flament-Fultot, M. (2016). Counterfactuals vs constraints: towards an implementation theory of sensorimotor mastery. *Journal of Consciousness Studies*, *23*(5–6), 153–176.
- Gangopadhyay, N., & Kiverstein, J. (2009). Enactivism and the unity of perception and action. *Topoi*, *28*, 63–73.
- Hurley, S. (1998). *Consciousness in action*. Cambridge MA: Harvard University Press.
- Hurley, S. (2010). The varieties of externalism. In R. Menary (Ed.), *The Extended Mind* (pp. 100–153). Cambridge MA: MIT Press.
- Hutto, D. D., & Myin, E. (2013). *Radicalizing Enactivism: basic minds without content*. Cambridge MA: MIT Press.
- Klein, C. (2007). Kicking the Kohler habit. *Philosophical Psychology*, *20*(5), 609–619. doi:10.1080/09515080701549330.
- Kosslyn, S. M. (1994). *Image and brain: the resolution of the imagery debate*. Cambridge MA: MIT Press.
- Marr, D. (1982). *Vision*. Cambridge MA: MIT Press.

- Maye, A., & Engel, A. K. (2016). The sensorimotor account of sensory consciousness: implications for machine consciousness. *Journal of Consciousness Studies*, 23(5–6), 177–202.
- McDowell, J. (1994a). *Mind and world*. Cambridge MA: Harvard University Press.
- McDowell, J. (1994b). The content of perceptual experience. *The Philosophical Quarterly*, 44(175), 190–205.
- Millikan, R. G. (1989). Biosemantics. *Journal of Philosophy*, 86, 281–297.
- Noë, A. (2002). Is the visual world a grand illusion? *Spec. Issue of Journal of Consciousness Studies*, 9(5–6), 1–12.
- Noë, A. (2004). *Action in perception*. Cambridge MA: MIT Press.
- Noë, A. (2005). Against intellectualism. *Analysis*, 65(4), 278–290.
- Noë, A. (2008). Reply to Campbell, Martin, and Kelly. *Philosophy and Phenomenological Research*, 76(3), 691–706.
- Noë, A. (2012). *Varieties of Presence*. Cambridge, Mass.: Harvard University Press.
- Noë, A. (2014). *Concept pluralism, direct perception, and the fragility of presence*. MIND Group: Open MIND. Frankfurt am Main Retrieved from <http://open-mind.net/papers/concept-pluralism-direct-perception-and-the-fragility-of-presence/getAbstract>.
- 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.
- O'Regan, J. K. (2011). *Why red Doesn't sound like a bell: understanding the feel of consciousness*. Oxford: Oxford University Press.
- O'Regan, J. K., & Degenaar, J. (2014). Predictive processing, perceptual presence, and sensorimotor theory. *Cognitive Neuroscience*, 5(2), 130–131.
- O'Regan, J. K., & Noë, A. (2001). A sensorimotor account of vision and visual consciousness. *Behavioral and Brain Sciences*, 24(5), 883–917.
- O'Regan, J. K., Rensink, R. A., & Clark, J. J. (1999). Change-blindness as a result of “mudsplashes”. *Nature*, 398(6722), 34–34. doi:10.1038/17953.
- Ramsley, W. (2015). Must cognition be representational? *Synthese*. doi:10.1007/s11229-014-0644-6.
- Rietveld, E., & Kiverstein, J. (2014). A rich landscape of affordances. *Ecological Psychology*, 26(4), 325–352.
- 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.
- Snowdon, P. F. (1992). How to interpret direct perception. In T. Crane (Ed.), *The contents of experience* (pp. 48–78). Cambridge: Cambridge University Press.
- Sprevak, M. (2013). Fictionalism about neural representations. *The Monist*, 96(4), 539–560.
- Stanley, J., & Williamson, T. (2001). Knowing how. *Journal of Philosophy*, 98(8), 411–444.
- Terekhov, A. V., & O'Regan, J. K. (2013). Space as an invention of biological organisms. *arXiv:1308.2124*. Retrieved from <http://arxiv.org/abs/1308.2124>.
- Tye, M. (2000). *Consciousness, Color, and Content*. Cambridge MA: MIT Press.