

Sensorimotor Direct Realism: How We Enact Our World

Michael Beaton • University of the Basque Country, Spain • mjsbeaton/at/gmail.com

> Context • Direct realism is a non-reductive, anti-representationalist theory of perception lying at the heart of mainstream analytic philosophy, where it is currently generating a lot of interest. For all that, it is widely held to be both controversial and anti-scientific. On the other hand, the sensorimotor theory of perception (which is a specific development of Gibsonian approaches to perception) initially generated a lot of interest within enactive philosophy of cognitive science, but has arguably not yet delivered on its initial promise. **> Problem** • I aim to show that the sensorimotor theory and direct realism complement each other, and that the result is a philosophically radical, but fully scientifically realised, theory of perception. **> Method** • The article uses (non-reductive) philosophical analysis and discussion. It also draws on empirical evidence from the relevant cognitive sciences. **> Results** • Direct realism can be augmented by sensorimotor theory to become a scientifically tractable alternative to the mainstream, representationalist research programme within cognitive science. **> Implications** • The article aims to further clarify the philosophical importance of the sensorimotor approach to perception. It also aims to show that the apparently radical claim that we perceive objects themselves is amenable to normal scientific study. **> Constructivist content** • Objects are analysed as a kind of collaboration between the world and the perceiver. On this account, we can never perceive outside the categories of our own understanding, but we do perceive genuinely outside our own heads. Thus, the approach here is not exactly constructivism, though it shares many goals and results with constructivism. **> Key words** • Sensorimotor theory of perception, direct realism, phenomenology, non-reductive ontology, consciousness.

Introduction

« 1 » Constructivists are opposed to the naïve claim that the categories of external reality are unproblematically present, independent of our minds. One might call this ultra-naïve claim “naïve realism.” In this article, I am arguing for a different position, which is also sometimes called naïve realism (Martin 2002, 2004), though it is more commonly, especially more recently, called direct realism (Sedivy 2008). What, then, is the “naïve,” or direct, realism that I am endorsing? If one has not been exposed to too much philosophy, one might, apparently naïvely, take it to be the case that perception puts us in direct contact with reality. Empiricist philosophers since David Hume have supposed that “the slightest philosophy” (Hume 1748: ch. 12, Part I) shows us that this “naïve” understanding of perception must be wrong. Indeed, the current, broadly empiricist, near-consensus in cognitive science still agrees with Hume on this. The direct realist position I am arguing for here, in contrast, argues that this latter “naïve” claim is literally right: that we can and do directly perceive reality. As we will see, however, in

endorsing this second claim (about directness), I will end up very much disagreeing with the original naïve claim, that perceived reality is mind-independent. How does all this relate to a radical constructivist view? Am I agreeing or disagreeing with such a view? I will return to this.

« 2 » The current, near-consensus, broadly empiricist position in cognitive science is called representationalism. According to representationalism, when we have a sensory experience, the content of our experience (i.e., the information about what we are perceiving, or seeming to perceive) is carried by a representation. A representation is supposed to be an internal physical state (for instance, a neural state, in our brains) with some mapping between the physical structure of the internal state and the perceived structure of what we experience.¹ According to representationalism, when a sensory representation is caused, in the right way, by the external objects that it

maps to, then we are (indirectly, but as best we can, according to this theory) experiencing the external world. If a sensory representation playing the relevant role arises for some other reason, then we are not actually perceiving the world, but are instead imagining or dreaming or having a hallucination or subject to some kind of illusion. How to plausibly account for such non-veridical experiences, which are only “as of” rather than “of” what they seem to be of, is a major challenge for any competitor to representationalism.

« 3 » Direct realism is a mainstream theory within analytic philosophy (Haddock & Macpherson 2008; Byrne & Logue 2009)² that is explicitly opposed to representationalism. Nevertheless, direct realism is considered controversial and, by many, anti-scientific. The direct realism that I will argue

1 | According to most versions of the theory, not just the representation but the rest of the brain, with the representation playing the right role in it, is required for experience to occur.

2 | Disjunctivism and direct realism are very closely related positions (Sedivy 2008), and I will not explicitly distinguish between them in this article. The position I am presenting is both disjunctivist (it treats object-involving experience as fundamentally distinct from non-object-involving experience) and direct realist.

for here is (in most essential details) John McDowell's direct realism (McDowell 1996). McDowell's direct realism draws much from Immanuel Kant's approach to the nature of reality (Kant 1996), and Kant's approach was developed as a reaction against Hume's empiricism (Hume 1748). Nevertheless, direct realism, as it stands, remains a philosophical rather than a scientific claim about the nature of perception. By this I mean that Kantian rationalism has never been fleshed out as a scientific theory, in the way in which the empiricism of John Locke (1689) and Hume has. However, I will claim in this article that there is a non-representationalist, fully scientific theory of perception that fits perfectly with direct realism (even though it has not previously been presented as such). This is the sensorimotor (or sometimes sensorimotor contingency (SMC)) theory of perception (O'Regan & Noë 2001).

« 4 » The version of sensorimotor contingency theory that I will present here is (in most essential details) Alva Noë's version of the theory (Noë 2004).³ I will present the sensorimotor theory of perception first, and then I will present direct realism. As we proceed, it should become clear that, according to both the theories I am endorsing, the structure of external reality, as we perceive it, depends very much (but not entirely) on the structure of our embodied minds. As Kant himself argued, what was naïve was to have supposed that things could ever have been otherwise. However (and I will return to this at the end), this is not a slight on the independence of reality itself (nor did Kant think it was). It follows, from this particular non-naïve understanding of reality, that there really is an external world, which we perceive, albeit that we can only perceive it from within the structures of our understanding.

« 5 » It can be seen that the approach I will present is not radical constructivism. But, I assure the reader, it is definitely not

representationalism, either. I will return to discuss in more detail the relation between the approaches that I endorse and radical constructivism in the final section of this article.

The sensorimotor theory of perception

« 6 » The sensorimotor theory aims to show that perception is directly related to understanding. It argues, firstly, that we can perceive only what we can understand, and secondly, and more strongly, that perception is an active process of the understanding engaging with the world. Note that, in this approach, there is no intermediate thing called "experience" that understanding engages with; experience is the active engagement of the understanding with the world. Note also that the "understanding" in question is not an abstract, intellectual understanding but rather a more fundamental, practical knowing-how-to-do, including, crucially, knowing how to act and interact. This approach is closely related to the Gibsonian (Gibson 1979) analysis of perception in terms of affordances. Thus, for instance, I can only perceive a chair as a chair, in virtue of understanding which actions chairs afford (sitting on; pushing and pulling in and out from under desks and tables; and so on). The sensorimotor theory of perception aims to continue this approach "all the way down." Thus even the most "basic" percepts, such as line, colour and curve (but also tone of sound, smell, touch and so on) are to be understood in terms of a practical understanding, of the possibilities for interaction that such stimuli afford. It is certainly right to say that perceiving *x* (colour, say) requires having the correct sensory apparatus to be able to pick up on *x*. But, more than that, it requires practical understanding ("mastery"⁴ is the word often used in the context) of the types of sensorimotor interaction that are afforded by external properties of the relevant type. Without such understanding, so the sensorimotor theory goes, there can be no perception.

4| In the sense "having mastery over," not "learning to master"; though, of course, we do need an account of learning, too (Beaton 2014; Di Paolo et al. 2014).

Shape and space

« 7 » So how does the sensorimotor theory work? The clearest example is the visual perception of the shape of objects in three-dimensional (3D) space. According to the sensorimotor theory, perceiving the shape of a 3D object consists in understanding how to interact with that object in space. Thus, for instance, in perceiving a nearby object I know (implicitly, not as any kind of explicit, propositional knowledge) how to reach out and touch the various parts of the object that I can see. Even for more distant objects, I know how to point towards the various parts of the object. Moreover, in both cases, I know (in the same, practical sense) how to direct my gaze towards the various parts of the object. I know more than this, too. I know how these characteristic patterns of interaction will change as I move about the object or as it moves about me. The central claim of the sensorimotor theory is that mastery of these actual and available actions is the same thing as knowing how to perceive shape. A closely related claim, which is perhaps less explicit in the work of Noë, perhaps more so in the work of Kevin O'Regan, but that I would wish to make, is that this action structure, of which I have mastery, is the structure of my experience of the shape of the object.

« 8 » This theory sheds light on certain facts about perception that are sometimes seen as puzzling in the philosophical literature. For instance, when speaking loosely (or so it is alleged), one might say that a distant tree "looks" as if it is smaller than a closer tree. But, some philosophers object, when one understands the layout of the world, distant trees do not look smaller than nearby trees at all, they look to be exactly the size they are. In fact, both points are correct, and both can be accommodated at the same time in sensorimotor theory. In seeing a distant tree to be the size that it is, one is understanding (practically) that if one moved closer to it, then the reaching and looking movements necessary to delimit its shape would be exactly the same as those required to delimit the shape of a closer tree of its size. Nevertheless, even when I fully understand this (and hence see the distant tree as the size it is), there remains a literal sense in which the distant tree looks smaller than the closer one. For right now, the distant tree

3| I would have reservations about some of the philosophical additions to the theory that O'Regan (2011) has recently proposed, in particular around the correct treatment of the self, though space precludes adequate discussion. Nevertheless, O'Regan continues to develop very successfully the science that will be required if something like the basic sensorimotor theory that I am presenting here is correct.

genuinely does subtend a relatively smaller visual angle than if it were nearby. In fact, it subtends the same visual angle that a smaller tree, close up, would have done. Thus there is something objectively similar between the interactions I have to take to interact successfully with a small, close tree and with a distant, large one. My (practical) awareness of this can be my practical awareness of the fact (which is, on this account, more than just a subjective appearance) that a distant tree looks like a small tree.

« 9 » This latter is a special case of a general feature of the sensorimotor account, which is that it analyses the structure of experience as something that is not ineffably private. The structure of my experience is the structure of the actions that my perceptual coupling with the world makes available to me. This structure is perfectly well amenable to objective behavioural⁵ tests (Beaton 2013). I take it as an advantage of this theory that it shows that experience is not something intrinsically private. Indeed, I would wish to argue that there is no aspect of experience that fundamentally cannot be studied from the third person, despite the strong claim by authors such as David Chalmers (1996) that to say so is not to take experience seriously. Does this somehow mean that my experience is not actually mine? That anyone can have exactly the same access to it that I can? No. For I can, and normally do, know perfectly well what I am seeing, and know (practically) what this enables me to do, without your necessarily knowing any of this. That is, I have first person access to my own experience (on which, more below). My point is not that first and third person access to experience are the same. They are not. The point is that *what* I have access to is not something fundamentally different from what you may gain access to through careful behavioural examination, even though I have access to it in a different way.⁶

5| “Actional” would be more correct, but awkward. Experience, in the sensorimotor account, is about norm-filled action, and it is not reducible to meaningless (norm-free) movements, as the behaviourists claimed.

6| I adopt these points from Shoemaker (1996), and would thoroughly recommend his deeply insightful analysis of the nature of introspection; although, for all that, I would disagree

« 10 » I have emphasized that the understanding on which sensorimotor theory depends is practical, not theoretical. This practical understanding is required in order to perceive, according to the theory. Theoretical understanding of what one practically knows how to do is also, of course, possible; but it is certainly not required, simply in order to perceive. However, I would argue that the fact that the relevant, practical sensorimotor knowledge is implicit, not explicit, is not just some quirk of how we are “built,” but is a fundamental requirement for any account of practical action. As Ludwig Wittgenstein (1953: §§185–238) and Lewis Carroll (1895) have both emphasized, if I know explicitly how to do something, then such explicit knowledge must be grounded in practical steps that I “just can” take; that do not depend for their justification on further explicit knowledge.

« 11 » However, it should not be supposed from this that it is either wrong or misleading, or both, to describe such base-level practical skills as knowledge, or as understanding. Our perceptual skills are *not* just hard-wired, automatic, not amenable to change, for all that many philosophers write as if they were. Consider the experimental evidence regarding adaptation to prisms that slightly offset the direction of a target (Fernández-Ruiz & Díaz 1999) or, indeed, consider the very rich first-person account given by Ivo Kohler (1964) of his experiments (largely on himself) of adaptation to inverting prisms, worn for long periods of time. The fact of the matter is that our basic perceptual skills can adapt and change. Moreover – and as Noë has emphasized – this adaptation requires attention, purpose, and voluntary interaction on the part of the agent: long term passive exposure to such disturbed sensory stimuli does not suffice for readaptation. The agent has to be trying to learn to see again; trying to understand their new situation. When they do so, their “basic” perceptual skills can and do readjust. Inter alia, this shows that these skills are very, very far from basic, in the sense of “simple” or “trivial.” It also shows that such skills are, indeed, a species of knowledge or understanding, properly said; for they are

with him as to the correct treatment of qualia (Beaton 2009a).

norm-governed, fundamentally flexible and responsive, and furthermore are deeply integrated with the structure of our more explicit knowledge.

Colour

« 12 » Sensorimotor theory can be extended to all aspects of our sensory world (sound, touch, taste, colour, etc.). Colour is often considered a hard example for the theory. What on earth are we perceiving when we perceive colour? Indeed, is colour not some purely internal response to external stimuli, which do not, themselves, have any such property? Many, since Locke, have believed in such an “error theory” of perception (a theory of perception in which some fundamental aspect of how the world seems to be, is not how the world is). But, as David Philipona and O’Regan (2006, 2008) have emphasized, there is, indeed, a very rich structure of interaction present in colour vision. In the case of colour, the interactions in question are to do with the interactions between lights, surfaces and our own visual system. We understand (very implicitly) what discriminations we can make in different lighting: which surfaces will look the same as each other, and which different; how all this will vary as the lighting, reflection, surface orientation, and so on, change.

« 13 » In this theory, colour constancy is not something that we achieve by adjusting to variations in our raw sensations of colour, for there are no raw sensations of colour. To perceive a colour is to perceive (to pick out, to master the existence of) the constancy in all this change (change in actual and available interactions).⁷ In this theory, it is correct to say that colours are not (purely) a feature of the external world, for the structure of the colours that we perceive is fundamentally

7| I am choosing forms of words here that are noticeably different from those that Noë chooses at this point in his discussion (Noë 2004: ch. 4). I have never been able to become entirely clear as to whether or not there are raw sensations of some sort, prior to understanding, in Noë’s theory, even though there is certainly no perception of the external world as such, prior to understanding, in Noë’s theory. In any case, I mean there to be no raw sensations of any sort, prior to understanding, in the version of sensorimotor theory I present here.

shaped by the visual system that we have. On the other hand, colours are not purely internal, either. For example, the physics of colour determines that the saturated colours – such as the bright reds and yellows found as warning colours in nature – will appear saturated to essentially any colour visual system that covers the same wavelengths. Thus it is an objective property of these objects that they will appear highly saturated in colour to any perceiver whose visual system covers the wavelengths that they reflect.⁸ Evan Thompson (1994) discusses extensively the inevitable failure of any account of colour that does not recognise colour as involving properties both of the perceiver and of the world. But for all that, in this theory, colours are objective properties in one important sense: whilst they are not perceiver independent, they are nevertheless an objective property of the available *interactions* between a certain class of perceivers (with certain visual receptors) and a certain class of features in the world (including surface reflectivity).

Sensorimotor direct realism?

«14» Thus we have a theory of perception that argues that perceiving is the same thing as engaging in (or being poised to engage in) meaning-filled, physical action in the world; specifically, in the characteristic patterns of world-involving engagement that define the proper objects of the various senses.⁹ What does this theory of perception have to do with the apparently quite differ-

ent philosophical claims of direct realism, which was developed separately, and which states that when we perceive normally, we directly perceive publicly accessible objects and properties themselves? I wish to argue that the sensorimotor theory is exactly what direct realism needs to turn it from a purely philosophical theory into a viable scientific theory of perception. If this succeeds, then it will also shed additional light on sensorimotor theory since this is not normally presented as being a direct realist theory (even though it is, quite normally, presented as being anti-representationalist in its own way). So next, I will next present direct realism; showing as I go how it fits with the sensorimotor theory of perception. Then, I will respond to objections to the theories I have been describing, which will provide further opportunities for showing how the two theories work together.

Direct realism

«15» Direct realism is a movement within analytic philosophy that is explicitly offered as an alternative to representationalism. It was first introduced, in the modern literature, by John Michael Hinton (1973), and has been most famously developed by McDowell (1996). McDowell's direct realism draws on the rationalism of Kant, both directly and via Wilfrid Sellars's reading of Kant. It should be noted, however, that McDowell also draws strongly on the writings of various phenomenologists (McDowell 2009). Direct realism is the claim that when I see, I see things themselves. That when I see an apple, for instance, my experience is directly of that apple itself, with no intervening mental image or representation. Furthermore (or so the claim goes in the version of direct realism I am arguing for), the external, publicly accessible apple itself plays an essential, ineliminable role in my first-person, phenomenal experience of it.

«16» What is the motivation for believing in direct realism? It might be acknowledged that it would be something of a relief, philosophically, if one could believe in it. After all, if one sits in the park on a sunny day and does not overthink things, one's perception certainly seems to reach out to, and embrace, the leaves on the trees, their motion as the slight wind moves them and the blue sky beyond. Unfortunately, science (or even "the

slightest philosophy" as Hume put it when discussing exactly this point, Hume 1748: ch.12, Part I) would seem to tell us that perception cannot reach out to the world in this way, that perception experience is an internal state that can, but need not, be caused by whatever objects we see, or only seem to see. Or so those of us brought up on a diet of empiricism and standard cognitive science have been led to believe. Is there an alternative? Can we find a viable way to defend and embrace the direct realist approach if we wish to take the relevant science seriously? I will argue that we can.

«17» For most people, their immediate objections to direct realism will centre around cases where the subject has an experience that is as if there were an apple (for example) present when there is not; for instance, dreams, illusion, hallucination, even imagination. We know that such cases occur (agreed). So we know that I can have an experience as of an apple without the apple being present (agreed). Does this not tell us that there must be an aspect of my experience of the "real apple" that it has in common with these imaginary experiences and which we can and should think of as being "the experience itself," this therefore being something that can occur with or without an apple being present? Furthermore, do we not also know – these days – that we could create this very same internal experience (at least in principle exactly the same, as far as the subject can ever know by introspection) by direct brain stimulation? And does this latter not show us, even more clearly, that the experience itself is something that occurs inside subjects, that it does not directly involve the world at all?

«18» All these objections try to show that veridical experience cannot involve the apple, because the very same type of experience (it is claimed) can occur without any apple being present. I will respond to such objections directly in the penultimate section of this article. But for now, in order to bypass such apparently strong objections, think of it this way. Imagine that we were allowed to treat successful, veridical, perceptual experience of actually present objects as the primary case. Imagine that we wanted to work out the best way to think about and understand this most central case first. And imagine that only afterwards would we try

8] Even so, need they "look the same" to different perceivers? I discuss this in Beaton (2009a) (see also my Author's Response §§26–31), wherein I try to give an externalist account of this qualitative aspect of sensation. (This involves affect and association, which can also be objectively studied.)

9] I do not think that our various senses are any more or less fundamentally separate than are the functions of the internal organs in our bodies, such as the liver, heart and lungs. We do have specialised sense organs for the detection of specific types of features of the world. But the understanding that is brought to bear in vision, for instance, is the same understanding, by the same agent, of the same world, that is brought to bear in hearing, touching and so on. Indeed, for example, the perspectival structure of space – the very same regularity – can be detected by sight, hearing and touch (Beaton 2013).

to understand illusion, hallucination, imagination, even direct brain stimulation derivatively, as less central cases that can best be explained in terms of their relation to the most central case. Imagine that we were allowed to do all this. Why might we even then think that perception, or even less plausibly (it might seem), first-person phenomenal experience, directly involves things in the external world? Here are some reasons.

Perceptual coupling

« 19 » There is a large movement within artificial life and evolutionary robotics that is dedicated to showing that the best way to understand perceptually guided action, in simple natural and artificial systems, is to think about perception as an ongoing process of interaction with the world (Harvey et al. 2005). Authors in this area would argue, for example, that insects do not perceive by converting stimuli into internal representations then processing those representations and deciding what to do next. Rather, perception in these simple cases involves a brain-body-world loop. An example that Inman Harvey has used is that of an insect trying to extract nectar from a flower. The flower is formed into a funnel shape, and the shape has been honed by evolution such that it guides the insect into the flower. This is not semiotic guidance: the plant is not just signalling to the insect which way to go. Rather, this is physical guidance. If the insect pushes in the vague direction of the nectar, it will find it. The flower is structured to make the insect's life easier. This means that part of the insect's sensorimotor job, of finding the nectar, is effectively done for it, in the world in which it finds itself.

« 20 » Andy Clark, for his part, likes to use the example of Barbara Webb's crickets, which perceive the sound of their mates not by taking in the full spectrum of external sound and then processing it but by having a tracheal tube of a given length running between the sound receptors on their legs, which makes this system physically tuned to the sound of mates of their own species, but not others (Clark 2001: 104–106). Equally, as the thought experiments of Valentin von Braitenberg (1984) and the actual experiments of many others (Beer 2003; Izquierdo & Di Paolo 2005; Harvey et al. 2005) have shown, very simple direct links between sen-

sors and actuators can sometimes achieve complex, highly interesting behaviour. Once again, these studies show that aspects of these interesting behaviours are achieved by "offloading" some of the work onto the world. You cannot give an explanation of the relevant behaviours in terms of processing or computation over internal states and symbols. Instead, physical aspects of the task (for instance, the agent's changing physical location in the world) are part of the correct explanation of how the task is performed.

« 21 » Now, once one has realised that perception can be like this, it becomes at least an open possibility that, at the lowest level, all perception, including our own, is like this. We are certainly capable of abstract, "off-line" thought, and imagination. But that still leaves open the possibility that such thought and imagination is built on, and fundamentally dependent upon, the possibility for real-time, on-line interaction of which we are also capable.

Direct realism and phenomenology

« 22 » Even if one accepts that perception, defined behaviourally, can be world-involving as just discussed, one might still balk at the claim that phenomenal, first-person experience is world-involving. Do we not *know* that experience itself happens in our heads?¹⁰ What possible reasons could there be to believe otherwise?

« 23 » Now, in Beaton (2013) I argued that (despite claims to the contrary) enactivist neurophenomenologists themselves remain dedicated to looking for the immediate correlates of experience in brain dynamics.¹¹ I suggested that this was the last remaining internalism in enactivism (and that it was actually quite deeply rooted there), but that it can and should be removed if we wish to progress with developing a non-representationalist, non-internalist science of mind.

« 24 » Here, I briefly repeat two arguments made in that earlier paper. I will claim

10 | I return to, and respond to, a particularly pointed version of this objection in the section "Introspection," below.

11 | Which is not to deny that such researchers acknowledge (indeed, emphasize) that brain dynamics are often tightly linked to, and influenced by, the dynamics of the external objects and properties that are being perceived.

that, in at least the two ways I am about to discuss, an externalist account of perceptual experience actually matches our first-person phenomenology better than any non-externalist account ever could. This is relevant because I am comparing direct realism to representationalism and to radical constructivism, each of which, in its own way, denies that external objects can play any constitutive role in experience.

Directness

« 25 » The position under discussion is direct realism, and I remind the reader of the earlier claim that the naïve phenomenology of perception is that we *are* directly in contact with the world.

« 26 » Consider the following quote from Maurice Merleau-Ponty:

“When we come back to phenomena we find, as a basic layer of experience, [...] not sensations [...], but the features, the layout of a landscape or a word.” (Merleau-Ponty 1962: 25)

« 27 » And this from Martin Heidegger:

“We never [...] originally and really perceive a throng of sensations, e.g., tones and noises, in the appearance of things [...]; rather, we hear the storm whistling in the chimney, we hear the three-engine aeroplane, we hear the Mercedes in immediate distinction from the Volkswagen. Much closer to us than any sensations are the things themselves. We hear the door slam in the house, and never hear acoustic sensations or mere sounds.” (Heidegger 1977: 136; quoted in Crane 2006)

« 28 » Indeed, representationalists do not deny this phenomenon. But they talk about the “transparency” of representational mental states (Speaks 2009) in a sense in which it is supposed that we “see through” our perceptual states, as if we were seeing the world directly. But in the representationalist view, it is only ever “as if.” Direct realism acknowledges – indeed embraces – the same phenomenon, but it claims that this seeming directness is more than mere seeming. If the direct realist account is right, there is no need to invoke “transparency” to explain an illusion of directness. Perception “seems” to be direct because it is. Of course, no internalist account can claim that.

«29» Additionally, in the version of direct realism presented here, we can offer more detail than is normally given about what this alleged directness might mean, scientifically. It means that the structure of my introspectible, first-person, phenomenal experience literally involves the external objects in the world. More specifically, the detailed structure of my experience is equated to the detailed structure of the actual and available norm-involving actions that my sensorimotor coupling with the world makes available.

Richness

«30» In this section, I will argue that internal brain dynamics can also never be a good match for the richness of experience. It turns out that the richness of experience (and the closely related notion of fine-grainedness) have been used by nonconceptualists to argue against the claim that I am making here, that experience is grounded in the direct interaction of the understanding with the world (Peacocke 2001). Effectively, I am trying to reverse that nonconceptualist argument, and to use richness to argue for conceptualism.^{12, 13}

«31» Experience presents the world to us as transcending what we know. Experience outruns us, surprises us. There is always more to find. To take a very simple example, consider the visual experience of a plate of salad in a restaurant. To start with, there is just a salad there. But if we look more carefully, we may see the particular leaves. If we

look more carefully again, we may see the veins on the leaves and the whorls of their edges. We may start to see colours that we had not noticed before. We can start to see the way the light and shade interact with those colours. If we look carefully enough, we will surely start to see types of things we have never noticed before.

«32» How best should we explain this? Is all the detail we could potentially attend to (at least from one particular position of the head and eyes) already copied into our heads somehow, ready for us to attend to it? Or is it the salad itself that contains all this rich detail, and our experience – which is the right type of skilful, involved interaction with the world – that enables us to encounter new aspects of this worldly richness as we attend to it, as it is, out there? According to the position argued for here, it is the latter which is correct. This claim is fully compatible with psychophysical experiments on covert attention (Wright & Ward 2008). When a subject fixates on one place, without moving their head and eyes, they can nevertheless change what they are attending to. It is empirically incorrect to suppose that the centre of your visual attention has to coincide with the area of the world that is focussed on your fovea (nor, indeed, is there a well-defined edge to “the fovea” anyway). Your attention can be focused on any part of the visual field, without changing where your eyes are looking. Such covert changes of attention are quite measurable objectively. For instance, a subject is quicker to spot changes in what they are covertly attending to, and can pick out more detail in it, compared to a comparable target at a non-attended location in the visual field. According to the present theory, shifting your attention amounts to more tightly coupling your (actual or potential) action to this or that part of the external world. Thus it is not compulsory to account for attention in terms of scanning an internal image, nor is it required that the correlates of the attended details are in the brain when non-attended; nor, indeed, is it required that the non-attended details be, in any other way (for instance, “non-conceptually”) present in experience ready to attend to (except, of course, in the sense that they are modes of interaction with the world that are, and that I implicitly understand to be, accessible from my current mode of interaction).

Objections

«33» Direct realism is prey to one objection above all others, the argument from illusion, which amounts to this: how do you account for experience as of an object when the object is not there (e.g., the experience as of a red triangle, with no red triangle present)? Sensorimotor theory, on the other hand, is, or has been, prey to one objection above all others, which is this: how do you account for experience when there is no ability to act (as, for instance, in locked-in syndrome)? I respond to both of these, and to some further important objections, in this section. Inter alia, I hope to thereby show that direct realism, when combined with the sensorimotor theory, is able to say something positive at points where standard direct realism has had to remain silent.

The argument from illusion

«34» The formal version of the argument from illusion (and its variants) aims to show that the existence of illusion (hallucination, etc.) demonstrates that perception cannot be direct. It is actually notoriously hard to fill out a watertight, logical version of such an argument, and I will not attempt to do so here. Indeed, I think that the correct exegetical reading of the argument from illusion is not as a watertight argument but rather as a standard rivalry between two competing theories, as follows.

«35» In a representational theory, it is trivial to account for an experience as of x when x is not present. We simply postulate a representation of x that is not caused by an instance of x but that otherwise plays the same role (or a sufficiently similar role) as the role that would be played by a representation of x that is caused by x in the normal way. Direct realism, on the other hand, talks about the role of objects themselves in experience, yet we can have experiences that seem to be exactly the same subjectively, with no objects present. Does this not show that the central aspect of an experience (the internal, introspectible, phenomenological aspect) cannot depend on external objects, since it can be had without them? And does this, in turn, not show that the representational theory of experience fits the facts far better than the direct realist theory?

12| I mention the terms conceptualism and nonconceptualism here only, simply in order to situate the argument and to help refer the interested reader to other relevant literature.

13| McDowell is a conceptualist, as is Noë (1999, 2004: ch. 6), as am I; but “conceptualist” is a misleading word. It is used to describe those who think that all aspects of experience constitutively involve *understanding*. However, this understanding need not be verbal, despite the name of the position. Indeed, it can be practical, as Noë suggests. I would also argue (space precludes adequate discussion) that McDowell has been wrong to say that literal understanding requires the possibility of explicit self-reflection; this is an aspect of McDowell’s conceptualism that seemingly rules out genuine experience in animals and infants, a result that I, and many others, find unacceptable.

« 36 » Many direct realists have been wary of the response to this objection which would claim that our experience has to be fundamentally different, from the first-person perspective, as between the situation where an object is present and the situation where it is not but seems to be. (Such theorists still certainly think that the external object plays a constitutive role in veridical experiences, but they are not so sure that this difference necessarily makes a difference to the first-person phenomenology.) I believe that direct realism is actually stronger if it endorses what I have already argued for above (in the sections on “Richness” and “Directness”): that the role of external objects in our experience makes a genuine, first-person phenomenological difference.

« 37 » It would be foolish to deny that we can sometimes have experiences that we take to show rich, direct contact with the world, even though they do not; but the direct realist should argue that these are cases of being misled. It is not compulsory to accept that full-blown experiential richness – and certainly not directness – is actually present there in such cases, even if we mistakenly think that it is.¹⁴ That is, as it were, the negative story: that it is not compulsory to ascribe certain features (richness, and genuine, world-involving directness) to non-veridical experience. This is all that traditional direct realism has been able to say, precisely because it has lacked a positive, scientific account of experience. But if the sensorimotor theory of perception is the right scientific theory to complement direct realism, then direct realism can and should adopt whatever positive story sensorimotor theory can give about the relation between illusion (and hallucination and visual imagination and so on) and veridical experience.

« 38 » I suggest that the best way for the sensorimotor theory to deal with experiences as of non-present objects is to discuss the relevant similarities between (even though not identical structure of) the actions that

an agent would take (if appropriately tested) when imagining (or hallucinating or having an illusion of or dreaming about) a given object (or property, etc.) and when actually perceiving such an object.

« 39 » For instance, if I visually imagine a horse, I would, if prompted, be able to sketch at least an approximate outline of the horse I am imagining; I could say which flank of the horse I am imagining seeing, and so on.¹⁵ If I visually imagine the front of my house then I can point out (physically point out, if asked) where the windows are, where the front door key goes and the like.¹⁶ Thus sensorimotor theory should say that when a subject is having an experience as of an object (property, etc.), with no such object present, then there is a currently instantiated structure to the actions that the subject is poised to take that is relevantly similar to the structure of the sensorimotor action that the subject would be poised to take when actually perceiving such an object.¹⁷ When I talk about actions a subject is poised to take, I mean to include actions a subject may (depending on further aspects of the situation and of the subject's goals) actually be taking. But I also mean to imply that all such actions are actions that the subject actually would take, given some (perhaps counterfactual – see next section)

15| I do not wish to claim that visual imagination of a horse will have all the perceptual properties of veridical perception of a horse – indeed not; but I am claiming that to the extent that it is visual imagination, it will have at least some of them.

16| Of course, unless I am having an illusion or hallucination (and not necessarily even then) I will not take the apparent object of my sensory experience actually exist. In illusion and hallucination themselves, however, the experience continues to seem real (despite any knowledge that I may have that indicates that it is not). To give a sensorimotor account of this sense of verisimilitude, specifically in illusions and hallucinations, I would argue that there must be a behaviourally detectable tendency to act as if the apparent objects of illusions and hallucinations actually exist, even when we know that they do not.

17| Beaton (2013) responds to the apparent problem for this response posed by cases of imagining non-existent or impossible objects, such as ghosts, unicorns or square circles.

14| Perhaps, in the case of experiences caused by some hallucinogenic drugs, there is a different kind of richness; but assuming that it is internally caused, then one can certainly argue that it cannot be as inexhaustibly rich as experience involving ongoing interaction with the world.

behavioural test. Note, finally, that even though the action-structure of illusion and hallucination is similar to the action-structure of actual perception, it is not identical; it cannot be, because the external object is not actually there to guide the subject's actions in detail.

The challenge of locked-in syndrome

« 40 » We now turn to the problem posed by locked-in syndrome for the sensorimotor theory. As is widely known, locked-in syndrome refers to actual cases of patients who have no (or almost no) ability to act, but still have fully (or largely) preserved consciousness and perceptual abilities (Laureys 2005). Such patients can perceive and imagine, and yet they cannot act at all. So much the worse, it would seem, for any theory that treats perceiving and imagining as being all about the actions that an experiencing subject takes (or would take, if appropriately tested). I present a fuller version of what I believe is the correct response to this in Beaton (2013), but my response to this objection is as follows.

« 41 » We should firstly note that counterfactual reasoning (thinking about what would happen if only some aspect of the world were not as it actually is) is ubiquitous in normal science. It is also required here. Given that we are allowed to use counterfactual reasoning, I believe that sensorimotor theory needs to argue that when a locked-in patient is perceiving, or imagining, that they still know how to act in exactly the ways that sensorimotor theory describes, and that they would so act if only they were re-enabled to do so. One way to enable them to interact (if only very minimally) is by using eye-blinks, either communicating directly with another person or via an eye-tracker attached to a computer. In this manner, such patients can indicate that they do indeed know how to interact in more complex ways; they just cannot do so. Another way to re-enable them to interact with the world would be via a brain computer interface (which is already close to the realm of science fact). Yet another (definitely counterfactual and currently science fictional) way to re-enable them to interact would be to repair the localised damage to their brainstem that has caused their

locked-in syndrome in the first place. But in all these cases, it is our evidence about the current (actual) state of the patient that gives us good reason (the very best of reasons) to believe that they would indeed interact, in exactly the ways described by sensorimotor theory, if only, counterfactually, they were enabled to do so.

« 42 » Indeed, this response to the challenge posed by locked-in syndrome is not so fundamentally different from the account that sensorimotor theory should give anyway of the case where a subject is seeing but simply choosing not to act. We always have to allow for the full structure of the space of sensorimotor actions that the subject would or might take (for instance, if appropriately tested) if we want to get at the full phenomenal structure of experience that sensorimotor theory can capture.

The scientific and manifest images

« 43 » Since I am writing an article that aims to introduce, and make plausible, direct realism, I would imagine that I am addressing an audience at least a part of whom are likely to be looking at other introductory sources on this topic, such as Wikipedia. The current, short, Wikipedia article on direct, or naïve, realism¹⁸ mentions two objections to the position that I wish to address here. The first is that naïve realism “propose[s] no physical theory of experience.” Indeed, this may have been correct to date. It is also, arguably, a glaring omission, since there is a leading near-consensus scientific theory about perception, namely representationalism, that is incompatible with direct realism. In this article, building on Beaton (2013), I have attempted to address this lack by arguing that sensorimotor theory is exactly the right scientific theory to provide the positive story of experience that direct realism has been lacking until now.

« 44 » The second complaint is that naïve realism is somehow incompatible with quantum mechanics. To address this second objection might seem to threaten to take us on a left-field excursion, far from the main journey through this topic. But I think that this is actually just a special case

of the more general claim, familiar since Sellars (1963),¹⁹ that the world of the senses (the “manifest image”) cannot be the same world as the world of science (the “scientific image”) since (so Sellars and many others have thought) science tells us that the world is not like that: it does not contain tables or chairs or colours, but rather contains protons, neutrons, electrons, electromagnetic radiation and so on. But, I submit, science says nothing of the sort as regards tables, chairs and colours. Science, read correctly, might well say that there are no such things as essentially private, unverifiable, internal qualia. That’s fine. But if tables, chairs and colours are ways of interacting, particularly suited to certain creatures with certain bodies and sensory apparatus, then science certainly does not (or should not) say that these things do not exist. Such ways of interacting are surely a part of the universe – the very same universe that includes (at a different scale) fundamental particles and fields.

« 45 » In the same vein, it is rather generally supposed that since quantum mechanics is even more strange than the classical physics of particles and fields, this simply makes it even more “clear” that the world as we perceive it is not the world as it is. I contend that this, too, is an incorrect conclusion. Indeed, most interestingly, the scientists centrally involved in developing the modern interpretation²⁰ of quantum mechanics have made a point of specifically emphasising that quantum mechanics allows us to recover something recognisably like our naïve picture of the world at the macroscopic level (Omnès 1999). That is not to argue that these authors advocate direct realism as such – they do not – but simply to argue that direct realism does not have the *prima facie* incompatibility with modern quantum mechanics (nor with any aspect of the “scientific image”) that it is alleged to have. Nothing in direct realism says that our senses must give us direct access to every aspect of the universe. But direct real-

ism, properly construed, does say that our senses give us one partial, but in the main true,²¹ view of things in the world, with such things construed as affordances for²² interaction for creatures like us. Even the quantum world *does* afford these ways of interacting at our size and scale. All this, I think, sits very well with the Kantian point (Kant 1996) that our direct perceptual access to the world is not something outside of which we can ever step but is, rather, where all of our explorations of the world (including our scientific explorations) must be based.

The distant stars

« 46 » There is one further, more or less scientifically based, objection to direct realism, which is to do with the distant stars and which I believe is worth responding to briefly, since I suspect something like it will be occurring to many readers. How – the objection goes – can I be directly perceiving the distant stars since we know that one can only see stars as they were hundreds, thousands or millions of years ago? And if I do not directly perceive the stars, then why believe that I directly perceive anything? The right response to this, I submit, is to argue that I *do* directly perceive the distant stars. The position I am arguing for is not called instantaneous realism. Perception is a process. I am, indeed, coupled to the distant stars as they were hundreds, thousands or millions of years ago; it is only their state that long ago that can affect me. Nevertheless, because of this coupling, I can (for instance) point at where the stars are (or rather were); I can also change my own relation to them relative to other things (for instance, I can move such that a certain star is, or is not, behind a tree or behind my hand); I can open or close my eyes (to see, or not see, them). Thus, more generally,

- I can act in respect of them, and
- I am indeed coupled to them, whilst viewing them (within the constraints of physics).

19| Note that McDowell draws on Sellars; and I certainly agree with McDowell that Sellars got many important things right. But not (I respectfully submit) this.

20| The details of which are far too little known in modern philosophy of science.

21| I choose to use the word “true” intentionally: on McDowell’s account normal, veridical perception is not merely “correct,” it is also epistemically fundamental.

22| More correctly, but more stiltedly, “that-which-affords.”

18| https://en.wikipedia.org/wiki/Naïve_realism, accessed 13 August 2015.

Thus, I would argue, there is nothing about perception of the distant stars (nor the more general fact that perception is a physical process that takes time) that amounts to a valid objection to the form of direct realism presented here.

Introspection

« 47 » I have been arguing for an externalist account of first-person experience. As such, I believe that one final objection will have been at the back, if not already at the front, of many readers' minds.²³ Experience is at-least-in-principle accessible to introspection, is it not? (Yes, I agree that it is.) But direct realism, of the form described here, claims that *experience itself* includes objects that are outside our heads. (Also agreed.) But that means that your position implies that a subject can *introspect* things that are literally outside of their heads. Is that not ridiculous?

« 48 » I agree that this would be ridiculous if introspection were some kind of inner perception, as it is very often supposed to be (with perhaps too little thought applied to the implications of this assumption). It is far less ridiculous if introspection involves rational transitions from mental states²⁴ to self-ascription of those mental states, as argued extensively and eloquently by Sydney Shoemaker (1996) (for some further discussion, see Beaton 2009b). To explain, the original state (which may or may not become introspectively known) might be that of seeing an apple. In the account I am arguing for here, this makes the apple itself available as (part of) the perceiver's reasons for action. For instance, the agent (animal; human; perhaps, at some

future date, artificial) might reach out and take the apple that it sees because the apple is there, because the apple is food and because it, the agent, is hungry. None of this needs to be at all reflective. Nevertheless, let us suppose that, in addition, the agent is in possession of the concepts necessary to ponder the truth, or falsity, of the self-ascription "I am (now) seeing an apple." Then, according to Shoemaker's persuasive line of argument, no kind of inner perception is needed in order to make the transition from seeing the apple to knowing that you are seeing it, precisely because of the intimate logical link between seeing an apple and knowing that you are.

« 49 » To draw a parallel: if one understands conjunction, then no kind of inner perception is needed to go from believing "A" and believing "B" to believing "A and B" – one simply has not fully understood what "and" means in this application, if one is not prepared to make this transition when and because it is true. So also with the self ascription of mental states, in Shoemaker's account: if one is not prepared to make the self-ascription of seeing (for instance) when and because it is true, then one simply has not fully understood what seeing means, in self-application, in this context. For much more on this see Shoemaker (1996); for a little more, see Beaton (2009a, 2009b).

« 50 » Now, if all this is right, then it can easily additionally be argued that the introspective state (of knowing that one is seeing an apple, say) naturally inherits as much, or as little, world-involvingness as the original, introspected state (of seeing an apple). I endorse Shoemaker's account of introspection (note that Shoemaker is not a direct realist), and I would endorse this latter line of argument. Thus, I am indeed saying that introspection (when understood correctly: not as internal perception, but as an ineliminably self-reflective mental transition) can and does reach out to include objects in the world, precisely because perception already does so.

« 51 » I should make a further important clarification. Unfortunately, or otherwise, there remains a place for the "inner" in the theory proposed here, in an important, albeit metaphorical, sense. Consider that some aspects and features of the

world are my mental states (my thoughts, feelings, experiences, etc.). Other aspects and features of the world are your mental states. Still other aspects of the world are public objects and properties (which can themselves be essential parts of either my or your mental states, while we are perceiving them, according to this theory).²⁵ Now, there is a long history of using the word "introspection" (misleadingly implying some kind of inner-directed perception) to refer to my special access to my own mental states. I certainly accept that there is such a thing as introspection (although, as discussed, I do not accept that it involves inner perception). There is also a long history of metaphorically talking of such introspectible states as if they were internal to me. I am afraid that in this article, I have not broken free of this metaphor. But where representationalism (and perhaps even some enactive theories) suppose that my mental states, as such, depend only on aspects of the world that are literally internal to me (states of my brain and/or body), the picture presented here argues that my introspectible, phenomenal mental states can constitutively involve not just my brain and body, but also (even more radically) the public, external objects and properties in the world. It is normal for enactive cognitive science to emphasize that a full explanation of action must depend on brain, body and world; but it is not common for enactive cognitive science to take this final radical step and argue that phenomenal introspectible mental states themselves literally and directly involve public objects.

23 | This is the objection to which I promised to return earlier at the start of the section "Direct realism and phenomenology."

24 | By "state," I mean something like "current position in what may in fact be a continuous, fluid process." Physicists, and many other hard scientists (and perhaps even some well-informed analytic philosophers), freely use the word "state" in this way. Because of this, it would fly in the face of the principle of academic charity to presuppose, without further evidence, that an author using this word is necessarily tied to an outmoded, static view of the mental.

25 | The meaning-filled interactions between me and the world are my mental states. These very same meaning-filled interactions can also be your public object of study. You would then be accessing what is first-person for me in a different, third-person, way. Furthermore, you can certainly also interact with my mental states in what has been called a second-person way (that is, by directly responding to their meaning, rather than studying them as objects, Thompson 2001); notably, McDowell emphasized this in his very earliest work on direct perception, which specifically concerned the direct perception of emotion (McDowell 1982).

Comparison with other approaches

« 52 » Before concluding, I will provide some necessarily brief pointers as to how the approach that I have laid out here relates to other strands of work within enactive cognitive science and, finally, to radical constructivism.

Enactive cognitive science

« 53 » I have stated that the position described here should not be misread as behaviourist. Norms and meaning are at its heart. All the actions with which sensorimotor theory is concerned are done by creatures for reasons. It is correct to say that sensorimotor theory tries to abstract away; for instance, to show that there is a common structure, which can be captured mathematically, to all those actions that would count as interaction in 3D space with 3D objects. Indeed, this approach abstracts away so successfully that no particular norm or motivation is mentioned in the structures that are described. All the same, there would be no meaningful action at all if there were no norms or motivation. Any given action, by any given agent, with a public 3D object (for instance) is done with some reason or other lying behind it (indeed, this is definitional of “action,” in the sense in which I am using it).

« 54 » The theory presented here does not attempt to explain the origin or provenance of such norms. It simply presupposes that they exist. It therefore definitely requires some further theory in order to be anything like a complete account of agency. For the same reason, it is quite compatible with those other strands of enactivism that are centrally concerned with accounting for the origin of norms, for instance, in biological autonomy.

« 55 » There have been claims, and counter-claims, as to whether the autopoiesis that is at the centre of much enactive cognitive science leads to some kind of internalism (Di Paolo 2009). Indeed, I myself have made the claim that some strands of enactivism read to me as quite internalist, as regards the correlates of first-person experience (Beaton 2013). It certainly seems tempting (and it seems to me that it has tempted some) to equate the supposed inner-directedness of experience more or less

directly with the self-directed, homeostatic nature of biological autopoiesis (autopoiesis being a theoretical term for active, self-referential self-maintenance). Be that as it may, I am perfectly happy to accept that the need to self-maintain is central to the origin of norms. However, if the theory presented here is correct, phenomenal experience is an emergent property of the necessary *interactions* between any self-maintaining agent and its world, and not a property just of the inner-directed self-maintenance itself.

« 56 » On another, perhaps related, point, I have tried to emphasize that when I say that I am relating perception to understanding, I am talking about practical, embodied understanding, “knowing how to do,” and not intellectual understanding, “knowing that.” Indeed, I believe that the latter is fundamentally built upon the former. For this reason, I believe that it is possible to equate the understanding that I am talking about here to the term “sense-making,” as more commonly used in enactive cognitive science (Thompson 2007; Di Paolo 2009). I do not believe that any creature has norms and purposes until it is engaging in at least basic sense-making interactions with the world. (Even the most minimal autopoietic cell does, necessarily, engage in such interactions.) Furthermore, once an agent is engaging in such interactions, it follows – on the view here – that the agent has at least minimal experience, and at least a minimal *umwelt* (von Uexküll 1957).

Radical constructivism

« 57 » I have persisted in talking about public, physical objects, and it may sound as if I am in serious danger of returning to the ultra-naïve position that I rejected at the start: of speaking as if the public world of objects is unproblematically present to different thinkers and unproblematically mind-independent. Actually, one goal of this article has been to show that some (but not all) of the important parts of this naïve picture can be successfully recovered.

« 58 » Consider the case of an apple. An apple, for me, can only ever be apple-as-I-understand-it. My own understanding of “apple” will always depend on my own idiosyncratic learning history. Thus someone else may well not agree with me on exactly what counts as an apple; perhaps they in-

clude quinces, where I do not; or do not include crab-apples, where I do. Indeed, I think, for almost any type of public object, there will be borderline cases that show that none of us ever exactly agree on anything. But exact agreement is not required. What is required is sufficient agreement, in commonly encountered cases, for shared reference to be possible. This, I think, is perfectly achievable. You and I can successfully discuss whether or not quinces should be counted as apples. Similarly, many non-human higher animals spend their entire lives engaging with each other in extremely complex, meaning-filled interactions concerning shared, public items in their worlds (food, prey, potential mates, other conspecifics, other resources and so on).

« 59 » This raises two technical points, which I can only address very briefly. Firstly, for reasons that have received a lot of attention within analytic philosophy (see, e.g., Davidson 1974), I doubt that it is possible for us to recognise something else as having a mind without thereby recognising at least some meaningful overlap between its public world and ours. If we cannot find such an overlap, then we will have nothing to tell us that we have found another mind; and once we have found such an overlap, we will have exactly what we need to tell us that we have found another mind. I would argue (if I had more space), that this follows directly from what we mean by mind (whether from the first-, second- or third-person perspective). The reader may perhaps discern why all this (which I can only mention here all too briefly) helps to explain why I find the concept of a shared, public world less problematic than I otherwise would.

« 60 » Secondly, there is the issue of the privileged (or otherwise) status of the scientific world-view – something that I believe is certainly of interest to constructivists. It seems to me that any agent or group of agents (for instance, as a thought experiment, an alien race) that we recognised as engaging in anything like science would eventually find the same fundamental components of the universe that we have found: electromagnetic and other fields; quarks, leptons and other sub-atomic particles; the fundamental quantum nature of matter; etc. That is my bet, though perhaps it is just a failure of my imagination to think so. However, even if I

am right about this, I do not think that this somehow entails that the scientific mode of enquiry is the only type of access to the only type of “truth,” or anything similar. Indeed, I fully endorse the claim that doing science is a very abstract, and – in a certain sense – very *non*-fundamental way of interacting with the world.

« 61 » How, then, should we explicitly address the question of whether or not the current approach is compatible with radical constructivism? It seems to me that radical constructivism (at least in as much as it has appealed to enactive cognitive scientists) features two key claims, one of which I can happily agree with, and the other of which I must reject.

« 62 » The first key claim, with which I agree, is that knowledge should be seen as coordination, not representation:

“what matters is not to match the world, but to fit into it [...] ‘to know’ is not to possess ‘true representations’ of reality, but rather to possess ways and means of acting and thinking that allow one to attain the goals one happens to have chosen.” (Glaserfeld 1991: 4)

« 63 » This is a central insight of enactive cognitive science, not just of radical constructivism. Even if I disagree with some aspects of radical constructivism, I certainly do not mean to disagree with this.

« 64 » However, the second radical constructivist claim, with which I must disagree, is the claim that meaning always is, and only can be, about the construction of arrangements of private experience:

“cognition produces conceptual structures by reflective abstraction [only] from material that is available within the system.” (Glaserfeld 1991: 4)

“all it means [to see a book, for instance,] is that in some part of our experiential field there is the kind of raw material which, if coordinated in a particular way, is sufficiently close to what our concept of book demands, so that we accept it as an instantiation of that concept.” (Glaserfeld 1991: 5)

« 65 » It should be clear that I must reject this second key claim of radical constructivism, for the whole point of my argument has

been to claim that we can and normally do have direct (non-representational – I agree with the first claim) access to the world. Thus, seeing a book is not representing it (that does not work), but it is nevertheless an interaction with the public, shared book, with the book itself forming an integral part of the experience. The object of such interactions is not, *per impossibile*, the “thing in itself” lying behind the book, it is just ... the book.

« 66 » As in radical constructivism, an agent’s experiential reality, in the view I have presented here, depends on, and is enacted through, the agent’s cognitive structures. Also as in radical constructivism, reality is here described in terms of fit, not in terms of representation. But radical constructivism describes the fit as being between private cognitive structures and private experience, whereas I describe the fit as being between action (always to be seen as at least counterfactually public²⁶) and a cognitively permeated (but publicly shared) world. Experience, I claim, is not something to which we could fit our cognitive structures; rather, experience is the ongoing process of an agent fitting itself to its world (presumably in order to maintain and extend the viability of what it cares about).

« 67 » But, someone might object, to what is an agent fitting its cognitive structures, if not to its experience? My answer is: to the world. That answer cannot be sufficient if we assume that we could only access the external world (if we could access it at all) via experiences that represent that world. In that case, the world “itself” could not be what we are actually fitting our cognitive structures to, only the (inner) experience could. It then might well be appropriate to argue, as Ernst von Glasersfeld does, that we do not access an external world at all, that all we do is find patterns in our own

26| By which I mean that the relevant actions always actually would be made, visibly and publicly, in at least some possible, even if not actual, state of affairs. (For what it is worth, I do not deny the reality of mental actions; but I do deny the reality of any supposed mental actions that could not be cashed out, not even counterfactually, in at least potentially public, visible actions. How else could a mental action be part of fitting a creature to its world?)

experience. But von Glasersfeld only argues against a representationalist account of our access to the external world. I have presented a non-representationist view. That view cannot be subject to the skeptical objections to representationalism that von Glasersfeld highlights and endorses (Glaserfeld 1991).

« 68 » From the point of view argued for here, as from the radical constructivist point of view, it is correct to say that we have no way of accessing the world, except via our cognitive structures. Indeed, it is a Kantian insight that we could have hoped for no more (Kant 1996). But it was also Kant’s insight that this is enough; that this gets us all the way to reality, with reality now understood as it should have been all along. Thus, contrary to the second key claim of radical constructivism, I argue that the Kantian view of the fit between mind and world (McDowell 1996) gets us out of our heads (Noë 2009) and into a cognitively permeated, but shared, public world.

Conclusion

« 69 » I have presented direct realism: the claim that we directly perceive objects in the world, with no intervening mental image or representation. On the other hand, I have argued for this by also arguing for the Kantian point that we can never step outside the structures of our understanding. The objects we can perceive are the objects we can understand. Nevertheless, these are public objects; other, similar, perceivers can perceive them too, at least well enough to engage in verbal, or non-verbal, collaboration and alignment of understanding.

« 70 » In order to support this direct theory of perception, I have had to draw a clear distinction between cases of normal, veridical perception and all other cases (illusion, hallucination, imagination, dreaming, etc.). I have suggested that the non-veridical cases should be analysed derivatively, in terms of their relation to cases of veridical perception. It is not that the non-veridical cases are unimportant, it is just that we should not allow them to divert us, too early, from the correct analysis of veridical perception.

« 71 » The position I have argued for is quite radical, in that it claims that public ob-



MICHAEL BEATON

is a member of the IAS Research Group at the University of the Basque Country (UPV/EHU) and a Visiting Research Fellow in the COGS Research Group at the University of Sussex. He is particularly interested in the sensorimotor theory of perception, direct realism and the hard problem of consciousness. He is drawn to those approaches that argue that perceptual experience is or involves bodily understanding, and is also keen to develop approaches which can satisfactorily reframe the hard problem of consciousness in such a way as to make it soluble, without thereby denying the reality of conscious experience.

jects and properties play a direct role in our “internal” (first-person, phenomenal) mental lives. I have even argued that the external, public objects that are part of our experience can be accessed in introspection. This, I have suggested, is no more or less radical than the original claim that objects play a direct role in our experience of them, at least as long as we follow the correct analysis of introspection: not as a kind of internal perception, but as a self-reflexive transition within an agent’s understanding.

« 72 » There is a complaint that direct realism offers no scientific account of perception. This is justified, but I have tried to remedy that here. The sensorimotor theory of perception argues that perceiving constitutively involves practical understanding, that what I can see is what I know how to interact with. I have presented the sensorimotor theory in some detail, and then argued that even though it is not normally presented as a direct realist theory, it is exactly what direct realism needs to turn it from an interesting but controversial philosophical view

into a perfectly workable scientific theory of perception. Doing this sheds light on the correct interpretation of sensorimotor theory, as well as on the correct interpretation of direct realism.

« 73 » I have argued that there is good scientific evidence in favour of the direct realist view of perception. Additionally, I have argued that sensorimotor direct realism matches our first-person phenomenology better than any internalist theory could.

« 74 » I have responded to various objections to the views presented here. Perhaps most importantly, I have argued that the structure of actual and available action, in cases of illusion and hallucination, is relevantly similar to the structure of actual and available action in cases of veridical perception. But it is not identical, for the external objects are not there to play their guiding role.

« 75 » The sensorimotor, direct realist theory that I have presented here entails that we cannot access the world except via the structures of our own understanding.

But for all that, it endorses the existence of a shared, public world to which we all have access. Indeed, it shows us what perception is, such that it may give us access to the objects and properties of the external world, and what these objects and properties are, such that we may perceive them.

Acknowledgements

I would like to thank one of the editors, Tom Froese, for his repeated prompts to me to write a paper laying out this position in more detail. Thanks also to Inman Harvey, and to the anonymous reviewers for this journal, for valuable feedback and requests for clarification.

I gratefully acknowledge support by the Spanish Government MINECO Project, Reference FFI2014-52173-P.

RECEIVED: 16 AUGUST 2015

ACCEPTED: 7 JANUARY 2016

Open Peer Commentaries

on Michael Beaton's "Sensorimotor Direct Realism"

Realities in the Plural

John Stewart

CRED, Université de Technologie de
Compiègne, France
js4a271/at/gmail.com

> Upshot • Direct realism can be better distinguished from objectivism and naïve realism, by recognizing the radical plurality of the incommensurable realities that can be enacted by living organisms in coupling with their environment.

« 1 » The target article by Michael Beaton makes an important contribution by articulating sensorimotor theory and direct realism, to their mutual benefit. However, to my mind, Beaton's position remains uncomfortably close to "naïve realism" or "objectivism." Objectivism is the position according to which, ontologically, there is a single well-defined reality that exists and is what it is independently of any relation to an observer, and epistemologically, that scientific knowledge can and should aim at providing a perfect representation of this reality. Naïve realism is the common-sense version of objectivism concerning everyday non-scientific knowledge. Constructivism is the antithesis of objectivism. It is therefore crucial to be clear how the direct realism proposed by Beaton is distinct from objectivism. In spite of his sophistication, it seems to me that Beaton does come uncomfortably close to objectivism, and in particular to making the assumption that there is indeed a single well-defined reality that exists and is what it is independently of any relation to an observer. Thus, in §65, he writes: "[...] seeing a book is [...] an interaction with the public, shared book, with the book itself forming an integral part of the experience. Not, per impossibile, the "thing in itself" ly-

ing behind the book, but just [...] the book." In a similar vein, in §58 he takes up his favourite example of "an apple":

"Consider the case of an apple. An apple, for me, can only ever be apple-as-I-understand-it. My own understanding of 'apple' will always depend on my own idiosyncratic learning history. Thus, someone else may well not agree with me on exactly what counts as an apple; perhaps they include quinces, which I do not; or do not include crab-apples, which I do. Indeed, I think, for almost any type of public object there will be borderline cases that show that none of us ever exactly agree on anything. But exact agreement is not required. What is required is sufficient agreement, in commonly encountered cases, for shared reference to be possible. This, I think, is perfectly achievable. You and I can successfully discuss whether or not quinces should be counted as apples."

In spite of Beaton's intelligent disclaimers and qualifications, I consider that there is here an undigested remnant of objectivism.

« 2 » As a possible antidote to this latent objectivism, I would like to make what is hopefully a constructive proposal. This consists of bringing out, more clearly and explicitly than Beaton himself, the radical *diversity* of the "realities" (in the plural) that can be enacted by living organisms in structural coupling with their environment. A classical prototype of the "lived-worlds" that can be brought about by living organisms actively engaged with their environment is provided by the "world of the tick" as described by Jakob von Uexküll (1992: 319). "Reality," for the tick, consists of three chained sensorimotor cycles: "butyric acid → drop; hairy surface → crawl; smooth surface → stick in proboscis and, if liquid is at 37°C → suck." By enacting this very simple lived-world, consisting of just three elements, the tick achieves the extraordinary feat (essential for

its trans-generational viability and therefore for its very existence) of catching a mammal (thousands of times larger, and able to run hundreds of times faster) and getting to suck its blood. Now the point I want to make is that, to a fair approximation, there are as many "lived-worlds" or "realities" as there are biological species – i.e., a countable but open-ended and unlimited infinity. Moreover, these various "realities" are quite radically incommensurable. Imagine, for example, the impossible dialogue between an oak-tree and a worm. The oak-tree: "The reality is, that you have to have a strong trunk, and branches and twigs with leaves to catch the sunshine, and strong roots to withstand winter storms." The worm: "My dear friend, you're just not with it. The reality is that you have to have a small, elongated, supple body so that you can burrow in the earth to find your food and escape from predators." Humans may be able to engage in a dialogue as to what properly counts as "an apple" (more on humans just below); but the worm and the oak-tree do not have a snowflake's chance in hell of even coming close to any sort of compromise agreement. Throw in the tick to make it a three-way discussion, and there is nothing but pandemonium.

« 3 » So what about humans? One of the particularities of humans is that we find a radical plurality of enactable but incommensurable realities *within a single species*. To kill several birds with one stone, imagine the hopeless "dialogue de sourds" between a scientist and an ordinary, common-sense human being; to make it more precise, between a classical physicist and an illiterate peasant. The peasant: "The reality is, that you first have to get a plot of land – better if you can own it, otherwise rent it from a landlord. Build a house for the wife and the kids. Then you have to prepare the land, plant and harvest the crops; and keep a herd

of cows (for milk and cheese), some pigs (for the meat), some chickens (for the meat and the eggs). And if you want to get fancy and philosophical, real experience is available through the senses: lights and colours, sounds, things you can touch and get a hold of." The physicist: "Land? Landlords? Crops? Cows and chickens? Lights and sounds and tangible objects? I'm sorry, but those are all mere superficial appearances. Reality, ultimately, consists of matter (atoms and molecules) and energy (mechanical and kinetic energy, heat, electromagnetic radiation); but those are not directly available to the senses, only to the intellectual mind equipped with scientific instruments." With the best will in the world, a consensual agreement here just is not on the cards. I am not trying here to *discredit* science; science, if done properly (but this holds for *all* constructable lived-worlds) does enact a perfectly legitimate reality. But what I *do* want to do is to put science back in its proper place: "scientific reality" is only one among a multitude of incommensurable realities, and it has no ontological privilege.

« 4 » This last point is worth elaborating, because a formidable stronghold of objectivism is "scientism": the view that science does have a privileged access to ultimate ontological reality. This is where radical constructivism can come to the rescue. In their book *Laboratory Life*, Bruno Latour and Steve Woolgar (1979) identify the processes of "splitting" and "inversion" by which a scientific hypothesis takes on the appearance (!) of being a "fact," a faithful representation of ontological "reality-in-itself." Every scientific fact starts out its career as a speculative hypothesis in the mind(s) of one or several scientists. At this stage, there is no doubt about its status: the hypothesis is a pure statement, neither more nor less; there is no question of it being a "reflection of reality." Now the overwhelming majority of these hypotheses die out: many die a rapid and painless death, because they are refuted by empirical observations and/or logico-theoretical considerations; and most of the rest die a more lingering death, because the scientists concerned do not take the time and trouble to perform experiments designed to refute/confirm the hypothesis, so it just withers away and is forgotten. However, in a small minority of cases, things work out

differently. It does happen, sometimes, that the community of scientists takes sufficient interest in a hypothesis to design and carry out experiments designed to test it. The longer the hypothesis avoids refutation by such experiments, the greater the interest, and the greater the motivation to design and perform even more powerful and exacting experiments. If this goes on for long enough, there are *two* distinct events – "splitting" and "inversion" – which generally follow so quickly on each other that a crucial sleight-of-hand goes un-noticed. First, then, a "splitting": the hypothesis projects a "double copy" of itself into "the real world out there" where it becomes "a real object." If one can catch the process at this stage – when the doubling has happened but the inversion has not yet followed – it is blatantly clear that the "real object" is indeed nothing other than a pure copy of the hypothesis (at most disguised by a simple paraphrase). At this point, it is not possible to say *anything* about the "real object out there" that is not simply a copy of the terms of the hypothesis. Thus, the status of "the real object" (the scare-quotes are deliberate) is that of a *construction* in the collective mind of the scientific community concerned. This "doubling" is followed by an "inversion": everyone starts to talk, rhetorically, as though the "real object" had been there all the time, patiently waiting to be "discovered" by the scientists in question. We arrive at the position that those who have participated in the process call "realist," whereas it would be more precise to call it "objectivist." The "referential impression" is strengthened by the rhetorical device of suppressing all the traces that identify the "real object" as a construction. Of course, this "inversion" process is itself a construction – an ultimate stage in the construction, but a construction nevertheless. If one does not catch the crucial stage when "doubling" has occurred, but has not yet been succeeded by an "inversion," it is virtually impossible to disclose the sleight-of-hand, and the "real object" no longer appears to be a construction. There is, however, a last resort if one takes the longer-term view of the history of science. The key cases are those when *after* the doubling-plus-inversion have occurred, some new evidence is produced that does now refute the hypothesis. In such cases, what happens is actually quite amusing: the

"real object" obediently dissolves and quietly goes away, reverting to its primary status as a hypothesis (which has now been refuted). Of course, if the "real object" really had corresponded to "reality-in-itself," correctly discovered by the scientists in question, this could not happen; any self-respecting "real object" that actually did correspond to ontological "reality-in-itself," would not meekly "go away" in this fashion.

« 5 » To be fair, in §60, Beaton himself does openly address the issue of "the privileged (or otherwise) status of the scientific world-view." In this paragraph he frankly admits that his latent physicalism may just be due to a failure of imagination. And even more importantly, he not only recognizes but fully endorses the claim that "doing science is a very abstract, and – in a certain sense – very non-fundamental way of interacting with the world"; in other words, that science is indeed a very specific and peculiar way of enacting a lived-world, and one that lies outside the primary realm of "direct realism." This leads me, in conclusion, to reiterate my main point. I suggest that it may be helpful to recognize more thematically the *radical plurality* of the incommensurable realities that can be enacted by living organisms in coupling with their environment; to include science within this plurality as an item that certainly has its own particularities (but then so does each of the other "realities"); and to state upfront that these particularities do not confer on "science" any ontological privilege. In this way, we may be able to take a significant step away from the pitfall of objectivism.

John Stewart took his initial degrees in physics and then in genetics. He subsequently worked in various areas, including sociology, the radical critique of science, and setting up "Science Shops" in France. In 1990 he met Francisco Varela, and has since then worked on developing the paradigm of enaction.

RECEIVED: 13 JANUARY 2016

ACCEPTED: 16 FEBRUARY 2016

Who is “We”? Some Observations on Sensorimotor Direct Realism

John Pickering

Warwick University, UK

j.a.pickering/at/warwick.ac.uk

> Upshot • Sensorimotor direct realism may describe how animals engage with their surroundings. But human beings are not typical animals. Their engagement can be metaphorical as well as direct, in which case the theory has less plausibility.

« 1 » In advocating sensorimotor direct realism, Michael Beaton adds the subtitle “How we enact our world” to his target article. A lot hangs on “we,” that is, just who is doing the enacting?

« 2 » In discussing the evolution of animal behaviour, Ernst Mayr (1974) distinguished “open” from “closed” evolutionary strategies, although the distinction is more a continuum rather than separate categories with clear boundaries. Simplifying, closed strategies are ones where the action repertoire of an animal is predominantly innate and adjusted to particular environmental conditions. Open strategies are ones where innate skills, while necessary, are not sufficient for full development, and where the animal is less dependent on particular environmental conditions. However, conditions need to be such that it is possible for the animal to acquire the complement of learned skills needed to attain competent adulthood. Closed strategies are characteristic of animals that inhabit relatively unchanging and predictable niches. Open ones, by contrast, are characteristic of long-lived social animals, often with extended periods of development and where habitation within a niche is more flexible and adventitious, requiring long- and short-term learning.

« 3 » Human beings are such an extreme example of an open strategy that they can hardly be said to lie on the open-closed continuum at all. Crucially, human perception-action skills include one that may in fact be a human monopoly, or if it is not, is only found elsewhere in the animal kingdom in a vestigial state. This is the capacity for what

may be called metaphorical perception and mimetic action.

« 4 » Animals predominantly engage with the world by seeing literally what it is that they are able to do in a given situation or with given objects. That is, they perceive the world “as is,” namely, as a dynamic field that provides opportunities to exercise their particular action repertoire. In terms of James Gibson’s theory of direct perception, they perceive affordances and events. Humans can do likewise, and when they are engaged in some activities, playing fast-moving sports might be a good example, they are doing nothing more than what an animal might do.

« 5 » However, and crucially, humans can also perceive and act towards the world metaphorically and mimetically, that is, seeing and acting “as if” rather than “as is.” This ability may mark a major milestone in the evolution of the human mind (e.g., Merlin 1993).

« 6 » For perception, seeing “as if” could mean an object being envisaged as other than it is or as it might be once changed after being acted on. For example, a rock might be perceived as a potential cutting-scraping tool once parts of it had been removed. This illustrates that rather than merely altering superficial features of objects and situations, humans are able to change the behavioural meanings of objects and situations radically, through metaphorical thought.

« 7 » For action, mimesis, acting “as if” something or someone, including the actor, were other than they are, opens the way to a significant new arena of communication not available to animals. To be human is in part to have acquired a repertoire of vocalisations, gestures, body movements mostly geared to language-like communication with other human beings involving shared attention to objects and situations, some of which may be distant in time and space.

« 8 » Broadly, this is the foundation of creative, symbolic action that supports human cultural life. Humans can treat an object or situation as other than it actually is, and can easily find novel uses for things. Crucially, animals cannot; indeed, when Martin Heidegger noted that “the animal is incapable of ever properly attending to something as such,” he meant that affordances related to an instinctive action repertoire are all that

animals can perceive (Heidegger 2008: 249). Perceiving an object as something detached from its usual behavioural meanings is a human monopoly.

« 9 » In the human case, an object can be examined, manipulated and modified and used creatively in ways that go far beyond the various examples of play and tool-use seen in the animal kingdom. The importance of creative play in childhood is recognised by educators as vital to cognitive and emotional development. The nurturing of metaphorical cognition is crucial since it opens the way for the child to inhabit the adult world where understanding is an active process of social construction that is, hence, significantly metaphorical and that needs to be understood hermeneutically (e.g., Berger & Luckmann 1966).

« 10 » It is this cultural arena that raises issues for Beaton’s discussion of sensorimotor direct realism. Animals, not inhabiting a cultural system of any significance, predominantly engage with the world via perceivable affordances, in James Gibson’s terms, or via the signifying parts of their *umwelt*, in Jakob von Uexküll’s (see, e.g., Brentari 2015). In this case, the variety of direct realism advocated by Beaton is plausible. Indeed, so much so that it would be otiose to suggest anything else.

« 11 » The human case is profoundly different. Modern human beings are products of the cultural evolution of the past two million years, albeit that the vehicle for the human condition shares a far longer heritage of biological evolution with animals in general and our close relatives such as the apes in particular. This heritage is signalled by the remarkable amount of overlap between the human genetic makeup and that of apes such as the bonobos.

« 12 » But, as Michael Tomasello (2001) points out, that overlap notwithstanding, the cognitive skills of humans and apes are profoundly different, and the difference comes from the externalised products of human cognitive activity that constitute culture.

« 13 » Much of the human *umwelt* is human-made. It might be more accurate to say that most of it is, since the parts of the *umwelt* about which developing humans most need to learn in order to become fully human are those that have been put there by other human beings, often for just that

purpose. By this is meant human creations such as symbol systems, the built environment and the practices that go with them. It is through assimilation of these during development that human infants become human adults. In assimilating the cultural world, cognitive skills, of both perception and action, are developed and deployed. These skills are unique to humans and make possible a form of conditional, flexible and creative perception and action not available to animals.

«14» Seen from this perspective, sensorimotor direct realism is perhaps less straightforward than it might be when considering the experiential world of animals. Engaging with the world metaphorically, that is, perceiving and acting “as if” rather than “as is,” adds a qualitatively new element to the situation. Perceiving is no longer direct, but via an intrinsic process of interpretation and cyclic re-adjustment (e.g., Neisser 1976). Moreover, whether the affordances of the human world are directly perceivable is a controversial case, as witness Gibson’s remarks about postboxes. Kurt Koffka (1935) advanced the notion of “demand characteristics,” that is, what objects invite or require observer-actors to do. Door handles demand turning, cups invite holding and drinking from, postboxes offer the posting of letters in and so on. Gibson accepted the idea of demand characteristics, but preferred to call them affordances. He also challenged Koffka’s assertion that demand characteristics depend on the needs of the perceiver. Thus, when someone wants to post a letter:

“For Koffka it was the phenomenal postbox that invited letter-mailing, not the physical postbox. But this duality is pernicious. I prefer to say that the real postbox (the only one) affords letter-mailing to a letter-writing human in a community with a postal system. This fact is perceived when the postbox is identified as such, and it is apprehended whether the postbox is in sight or out of sight.” (Gibson 1979: 138f)

«15» It is not productive to try to resolve this issue here. We can, however, take from it the point that for human beings, perceiving what can be done with a postbox or any other culturally created artefact requires interpretation and involves culturally conditioned needs and intentions.

«16» So if “we” enact our world, perhaps some qualification may be needed to specify the “we” being referred to. Both animals and humans may enact their worlds in order to perceive and engage with them. Only culturally shaped humans, however, have the ability to enact in ways that are creative and metaphorical; and those ways may be far from direct.

Following degrees from Edinburgh and Sussex universities in the UK and postdoctoral fellowships in the US, at Rochester and Stanford, **John Pickering** has worked at Warwick University in the UK, where he lectures on psychology, philosophy and environmental issues. His principal research interests are consciousness, process thought, ecological psychology and biosemiotics.

RECEIVED: 12 FEBRUARY 2016

ACCEPTED: 18 FEBRUARY 2016

The Epistemological Dance: Difference, Experience and Representation

Hugh Gash

Dublin City University, Ireland

hugh.gash/at/dcu.ie

> **Upshot** • Accepting the biological origins and limits of what we know is a foundation stone of radical constructivist (RC) research. A corollary is that RC considers realism as allowing an impossible comparison between knowledge and reality. Recent works such as that presented in the target article have a more nuanced position in relation to “reality.” Points of similarity and difference between RC and direct realism are discussed in this commentary.

«1» Jean Piaget (1954) was introduced to the English-speaking psychological world at a time when Burrhus Frederic Skinner’s behaviourism was prominent in many psychology departments. At that time, psychology and philosophy were usually studied separately in universities and often the philosophical nuance evident in the target article was absent in psychological writings.

Skinner (1974: 12) maintained the future of psychology was in observable behaviour and that “mentalistic explanations allay curiosity and bring inquiry to a stop.” Alternatively, Piaget’s cognitive developmental psychology was presented as an interaction between children and their environments and focussed on developmental changes in the way mind organises experience (Kohlberg 1969). In Ernst von Glasersfeld’s (1984) analysis of the cognitive developmental “person-environment” interaction, knowledge was not about discovering “reality,” but about organising understanding. Of course, this meant “reality” has a radical status in radical constructivism (RC) that is dissonant with everyday use. This target article is welcome because it offers the opportunity to reflect on the relationship between knowledge and “reality” and to revisit the observable and the unobservable in psychological domains that contrast with the position taken in RC.

«2» Cognition and perception depend on noticing differences. As will be apparent in what follows, focussing on one part of a difference means that the other part is obscured or ignored. The author presents perception in sensorimotor experience as an active engagement with the world (§6). One of the remarkable things about watching children’s sensorimotor play is noticing their experience of novelty, their reflections and creativity. I would describe these experiences as noticing discrepancies between their experience of the world and their expectations. How does this work in Michael Beaton’s account? This seems to be a counter example to the position taken in the target article because noticing discrepancies is an internal event. The unproblematic apple is itself the result of hidden distinctions between oranges, pears and apricots (§15). So what Michael Beaton presents as a direct encounter with an apple requires a complex construction built up over time and dependent on sensori-motor learning that involves subtle differences, including specific ranges of shapes and colours. I am not in favour of DR for a variety of reasons I raise in this commentary, and I think Ernst von Glasersfeld would have completely rejected this idea. Unfortunately, as far as I am aware he never discussed McDowell.

«3» In RC, von Glasersfeld (1984) emphasised personal interpretation, as was the

case in John Dewey's epistemological work (Dewey 1960). Also, for both Glasersfeld and Dewey, knowledge and concepts were the results of mental operations. "Reality" had become mysterious and this is a major feature of RC, but ideas were to be checked through recursive experiment. These constructivist approaches offered a perspective on objectivity as "parenthesised," in Maturana's phrase (1988). This raises a central issue in the target article, what role has objectivity "without parenthesis" (§§1–4)?

« 4 » Representationalism as presented in §2 seems like RC in prioritising internal experiences and structures. However, RC would neither argue that external "reality" causes the representations nor deny that the representations are the result of a person's "objectivity in parenthesis" interface. A critical difference between RC and both representationalism and direct realism (DR) is that RC denies that it is possible to make claims about the relation between experience and "reality." In this denial, it falls into the category of an "externalist theory of perception" (BonJour 2013). So RC, having given up on objective reality but not "objectivity in parenthesis," seems to be less concerned about the "error" implied by hallucinations and illusions (§§34–39). Of course, the role of the constructing mind in making sense of illusions is a standard classroom demonstration of the importance of the constructing mind in RC. How does one check one's version of DR without an internal representation?

« 5 » DR as presented in §3 also has some shared features with RC. In particular, it is clear throughout the target article that "reality" plays a part in perception and cognition through the filter of the structures of one's understanding (§75). This is a change from the naïve realist position presented in the early part of the target article (§1) and a change that recurs in other philosophical works such as Luciano Floridi's essay on constructionism (2011). I have some comments. First, can we regard this as abandoning the idea of objective reality in DR as suggested in RC? Second, RC emphasises that we have contact with our best viable construction. This is different from denying that people have contact with the world. So I agree with Michael Beaton (that), "The point is that what I have access to is not something

fundamentally different from what you may gain access to, [...] even though I have access to it in a different way" (§9). However, I disagree that RC denies that external objects play a constitutive role in experience (§24). How does DR resolve different interpretations without recourse to representation and experience?

« 6 » Coming from an RC perspective, I would argue that the realist position in (§25) needs clarification. Yes, I agree we are in contact with the world, but what we know is internal and an interpretation. For example, where Maurice Merleau-Ponty writes about features and things (§§26f), an RC position views these features and things as constructions. Yes, I agree that we do not experience raw sensations (for the most part) because our existence and sanity require building ways to organise our sensations. So it seems to me that for the author to deny "reality" plays a role in the RC account is overstatement. Is DR's position on "reality" close to RC's if it is clear RC does acknowledge an interface with "objectivity in parenthesis"?

« 7 » Is it the case that the author and I are in a kind of dance with our ways of trying to get to grips with the intricacies of how to describe and explain objects and more complex social phenomena? Philip Boxer and Vincent Kenny (1990) argued for the need for a third order cybernetics to overcome the ways essential features of systems are obscured by particular perspectives. In the present article, at one moment DR emphasises direct interaction with "reality" and suggests "experience" may best be written out of the script (§23). In another phase, RC is shown to emphasise private experience (§64), with the implication that "reality" has no role in the construction of this experience. The dance flows because there is a constantly moving focus of attention in our respective momentarily dominant world views.

« 8 » This dance depends on changes in emphasis in the conversation as to what is centrally important in describing and explaining thought. It shows how the emphasis on the need for direct contact with "reality" in DR leads to its own concerns in terms of how to deal with imaginary objects, illusions (§§34–39), locked-in syndrome (§40–42) and introspection (§§47–51). These emphases are highlighted for this reviewer in the

theory-laden nature of statement "I certainly accept that there is such a thing as introspection (although, as discussed, I do not accept that it involves inner perception)" (§51). However, I find congenial the description of introspection "as a self-reflexive transition within an agent's understanding" (§71), what is this if not a form of inner perception and experience?

« 9 » Wearing my RC hat, the de-emphasis of inner perception in DR seems a function of the emphasis on direct contact with (structure interpreted) reality. Similarly, maybe the following quote "that meaning always is, and only can be, about the construction of arrangements of private experience" (§64) ought to be considered a function of RC's insistence on "reality" being mysterious. In RC, whether or not a person has interpreted the raw material viably "in private experience" depends on finding others who agree; or as Floridi (2011) put it, we know we have misinterpreted when we miss the train.

« 10 » I wonder if we can find a solution to this dance. In an RC position, I have no difficulty agreeing that some features of "reality" play a constitutive role in cognitive constructions, and what I know depends on my human biology. I have no difficulty accepting that some objects in "reality" are public. Viability in RC depends on reflective iterative encounters with experience. How does DR account for varying interpretations of "reality"?

« 11 » The issue is surely to move beyond the different foci and to see where we really are in agreement and where is the difference between DR and RC? On reading the target article, I find I agree with a representationalist position accepting that the content of experience is caused by the interface between the active cognitive agent and her experience. I agree with a DR view that accepts that "objectivity" is parenthesised. I am at a loss to understand why DR dismisses internal perception. DR as presented here recognises the mind-dependence of "reality" and also the need for a sensorimotor account of the origins of concepts (§4) and accepts the RC insight that knowledge is a coordination (§62). However, this coordination must be initially an individual act, and while we can hope for interpersonal agreement, agreed access to a publicly shared world (§66) is

something for which we can strive but about which we cannot be certain. Is the major impasse between RC and other epistemological approaches to perception that RC has abandoned questions of the truth of perceptions in favour of questions of viability and so can be classified as having an “externalist theory of perception” (BonJour 2013)? Or is the impasse that RC has turned away from classical epistemology and accepted the need to incorporate a psychological account of knowing?

Hugh Gash worked at St. Patrick's College Dublin, now incorporated into Dublin City University, until 2010.

Gash is a member of the International Institute for Advanced Studies in Systems Research and Cybernetics.

He has published extensively on educational applications of constructivism, details of which may be found on his website, <http://staff.spd.dcu.ie/gashh>

RECEIVED: 17 JANUARY 2016

ACCEPTED: 12 FEBRUARY 2016

Representationalism and the Sensorimotor Theory

David Silverman

Paris Descartes University, France
davsil/at/gmail.com

> Upshot • In light of the construal of sensorimotor theory offered by the target article, this commentary examines the role the theory should admit for internal representation.

« 1 » One aspect of the sensorimotor theory (SMT) that has often been a subject of controversy is the precise theoretical role it should be read as admitting for representation. One of the target article's virtues is that it makes useful progress toward resolving this controversy, by highlighting what I believe is the most crucial sense in which SMT is properly construed as anti-representationalist. I agree with everything said in the target article, and the purpose of this commentary is to build on it by clarifying the role SMT properly admits for internal representation, given the construal recommended by Michael Beaton.

Two kinds of representationalism

« 2 » There are a number of distinct positions referred to as “representationalism.” I will address two of them.

« 3 » One, which I label “constitutive representationalism,” is the claim that perception is essentially a kind of internal representation. Constitutive representationalism is sometimes featured in philosophy of perception, not always of a naturalistic kind. It can be contrasted, as it is in the target article, with direct realism about the epistemology of perception.

« 4 » Constitutive representationalism can also be found, albeit in a slightly different form, in cognitive science and its philosophy. Cognitive scientists betray a commitment to constitutive representationalism when they suggest that cognition, including perception, consists, as a matter of conceptual necessity, or even definition, of the deployment of neurally-encoded representations. This commitment is identified and rejected by, for example, William Ramsey (2015) and Alva Noë (2004). One reason SMT serves as a *scientific* complement to direct realism is that its arguments are directed against constitutive representationalism as espoused by cognitive scientific theories of perception, not only philosophical ones.

« 5 » “Enabling representationalism,” as I will call it, claims that perceptual experience is realised by the subpersonal deployment of internal representations, but does not claim that perception is constituted by (i.e., identical to) an activity of internal representation. John McDowell (1994) argues that perception is plausibly enabled, subpersonally, by a process of internal representation, even though it is not constituted by one. This shows that you can endorse enabling representationalism even while giving a skill-based account of perception that is actively incompatible with constitutive representationalism.

« 6 » McDowell uses this point in an attempt to reconcile James Gibson's (1966) anti-representationalism (which rejects constitutive representationalism) with David Marr's (1982) representationalism (which McDowell supposes only endorses enabling representationalism). Noë (2004) rejects the proposed reconciliation on the ground that Marr actually endorses constitutive representationalism, and not mere enabling rep-

resentationalism. This is betrayed by Marr's claim that “vision is the [representational] process of discovering from [retinal] images what is present in the world, and where it is” (Marr 1982: 2, emphasis added). One ill-effect of this claim, Noë observes, is that it prejudices Marr's account of vision's enabling features.

« 7 » Nonetheless, McDowell's broader point is sound. You can deny that perception is a kind of representation without ruling out the possibility that it is enabled subpersonally by representations. By this token, it would not compromise SMT's constitutive account of perception, as Beaton characterises it, to allow that perception might be enabled by subpersonal representations. Although SMT rejects constitutive representationalism, it is not committed to rejecting enabling representationalism.

Enabling representationalism

« 8 » While McDowell merely intended to show that representationalism about perception's enabling features is compatible with anti-representationalism about its constitutive features, there is a respect in which McDowell's constitutive/enabling distinction actually makes enabling representationalism easier to defend.

« 9 » Consider one prominent enactivist argument against enabling representationalism. Daniel Hutto and Erik Myin (2013; henceforth H&M) state that cognitive scientific accounts of perception should not make any explanatory appeal to representation whatsoever. They argue that the most promising accounts of content all depend on the idea that co-varying with something is identical to representing it, at least when certain further conditions are met. H&M claim that there is nothing compatible with naturalism that could adequately motivate the claim, for example, that a tree's rings bear truth conditions (and so content) about the tree's age, as opposed to merely co-varying with its age. They conclude that enabling representationalism should therefore be rejected.

« 10 » One straightforward response the representationalist can make it a brute stipulation that by representational content they mean covariance and nothing more. But this move is troubling, because we cannot tell if it is deflating the notion of representation (which would be fine) or inflating the notion

of covariance (which violates naturalism). In view of this, a stalemate beckons. But the argument is potentially resolved in Hutto and Myin's favour by considerations of conceptual hygiene: if "representation"-talk is apt to being interpreted in a harmful way, it is best avoided.

« 11 » The best way to show that "representation"-talk is being used in a harmless, deflationary way, and hence to overcome H&M's criticism, is to show that it is not being used to naturalise anything usually considered subject to an explanatory gap, for example the phenomenal quality or intentionality of perceptual consciousness. The best way to show this is to reject constitutive representationalism explicitly, and claim instead that perception is constituted by the exercise of sensorimotor skills.

« 12 » The moral is that SMT, even when construed as a scientific complement to direct realism, may nonetheless be further developed scientifically with reference to internal representation, and moreover that SMT, construed in the way Beaton recommends, in one respect makes the case for "representation"-talk at the subpersonal level more secure.

Avoiding constitutive representationalism

« 13 » It is important that SMT does not lapse into constitutive representationalism, however. As Beaton underlines (§6), sensorimotor knowledge must be construed as practical knowledge, i.e., know-how. This know-how must not be parasitic upon knowing-that, since this would suggest that perception is essentially a process of internal representation.

« 14 » A puzzle here presents itself. Sensorimotor knowledge includes knowledge about the consequences of movements that need not actually occur, as Beaton (§7) notes. Indeed, SMT must appeal to counterfactual knowledge to do justice to the phenomenology espoused by Noë, which claims that you can visually experience the presence ("in absence"; Noë 2004: 128) of the back of a tomato without making the movements required to come into sensory contact with the back of the tomato. The problem is that knowing how your sensory inputs *would* change in line with movement is, on the face of it, a kind of knowing-that, not purely a knowing-how.

« 15 » I propose that sensorimotor knowledge consists of the ability to carry out bodily actions that betray a sensitivity to the changes in sensory input that would occur as a result of possible movements. The knowledge can be ascribed in a similar manner to the way Daniel Dennett's (1987) "intentional stance" ascribes beliefs and desires. We look at an agent's behaviour, ascribe to her a goal-state, and on this basis ascribe to her knowledge of a particular set of sensorimotor contingencies.

« 16 » To revisit an old example, consider a guided missile following a plane (O'Regan & Noë 2001). We can ascribe to the missile the goal of keeping the plane aligned in the centre of its sensor, and explain its success by ascribing to it the knowledge that turning its nose to the right or left would cause the image of the plane to shift a corresponding degree further to the left or right in its sensor. When the target appears in the centre of the sensor, implicit knowledge of those contingencies is manifested by the missile's *not* changing course. Similarly, to experience the tomato's hidden side, the relevant sensorimotor contingencies do not have to be actualised, i.e., you do not have to come into sensory contact with the back of the tomato. It suffices that you act in a way that manifests a sensitivity to the sensory consequences of possible movements, even in cases where some of those movements do not actually occur.

« 17 » Notice that although sensorimotor knowledge, so understood, is logically dependent on capacities to perform particular goal-directed actions, it is not identical to those capacities. In this sense, SMT is not an action-oriented theory. Sensorimotor knowledge is the capacity to respond, regardless of the particular goal, with sensitivity to the ways sense inputs are prone to change if particular movements occur. All the same, sensorimotor knowledge is grounded in your ability to act, and can therefore be construed as purely practical knowledge.

Avoiding constitutive representationalism: Part two

« 18 » SMT can this way account for the perceptual presence of absent features without identifying perception, at the personal level, with knowledge-that (and so

representation). This personal level view is compatible with perception being enabled, subpersonally, by representations.

« 19 » Tom Roberts (2010) endorses a variant of SMT's skill-based view of perception at the personal level. But noting that the content of perceptual experience includes environmental features with which the perceiver is not presently engaged in bodily interaction – a truism that we have just seen is endorsed by SMT's own peculiar phenomenology – he hints that perceptual experience therefore actually *requires* subpersonal representation. In other words, perception appears to be subject to what Andy Clark and Josefa Toribio (1994) call "representation hunger."

« 20 » McDowell's (1994) approach would reject constitutive representationalism even if representation hunger made representational explanation at the subpersonal level indispensable. This is because, in McDowell's outlook, perception is a personal-level activity, and subpersonal representations can at most enable perception, even if they play a necessary role.

« 21 » However, it is not clear that sensorimotor theorists should accept that perception is necessarily a personal-level (or agent-level) phenomenon, and Noë (2004) indeed suggests that there is no clear 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. If perception as construed by SMT is subject to representation hunger, it by this light appears to entail constitutive representationalism.

« 22 » We could concede representation hunger while resisting constitutive representationalism by claiming that perception is constituted not just of representation, but of skilful bodily interaction that draws on internal representation. The appeal to bodily skill, here, would make SMT compatible with direct realism. But SMT, so understood, would not lend any extra support to direct realism, since it could as easily be

interpreted, instead, as a scientific complement to the view that the epistemic access that perception gives us to the world is mediated by a representation. This would be a pity.

« 23 » A better response is to reject the notion of representation hunger (see Degeenaar & Myin 2015). We should not do this by dispensing with the characteristics that make SMT *appear* subject to representation hunger, such as its claim that perceptual experience presents absent features such as the back of the tomato. Instead, we may conceive of these characteristics as entailing *prima facie* representation hunger. *Prima facie* representation hunger, I propose, does not entail that representation is indispensable, merely that representation *could* do the necessary enabling work. We should maintain that there are in principle non-representational ways of implementing sensorimotor knowledge (explored, for instance, by Thomas Buhrmann and Ezequiel Di Paolo 2014, and Martin Fultot 2016).

« 24 » It does not matter to SMT whether perception does or does not happen to be enabled subpersonally by representations. Denying that perception *necessarily* draws on internal representation, and this way rejecting constitutive representationalism, is sufficient to ensure that the support given by SMT to direct realism, as highlighted by Beaton, is secure.

Acknowledgement

This work was financed by ERC Advanced Grant 323674 “FEEL.” I am further developing the points made in this commentary in a full paper, currently in progress.

David Silverman is a postdoctoral researcher working as part of J. Kevin O'Regan's ERC funded project, FEEL. His work focuses on defending and developing the philosophical component of the sensorimotor approach to perception and consciousness.

RECEIVED: 11 FEBRUARY 2016

ACCEPTED: 14 FEBRUARY 2016

Phenomenal Promiscuity

John Mark Bishop

Goldsmiths, University of London,

UK • m.bishop/at/gold.ac.uk

> **Upshot** • Sensorimotor direct realism is too promiscuous in its account of sensation.

« 1 » In arguing against what he takes to be the near consensus view in cognitive science of “representationalism,” Michael Beaton presents a serious attempt to rehabilitate direct realism as a viable, scientifically testable, theory of mind by making more explicit the links to Kevin O'Regan and Alva Noë's sensorimotor theory of perception (2001).

« 2 » The fundamental tenets of direct realism (DR), as outlined by Beaton, can be summarised as stating that: “we are directly in contact with the world” (§25); “we can and do directly perceive reality” (§1); and that “perceiving is the same thing as engaging in (or being poised to engage in) meaning-filled, physical action in the world,” (§14).

« 3 » The central tenet of sensorimotor theory (ST), as conceived by O'Regan and Noë (2001), is a reconceptualization of visual perception, away from analysis of the raw visual patterns of stimulation, to focus on the law-like changes in visual stimulation brought about as a result of an agent's actions in the (light-filled) world; in this way ST offers a radical enactive approach (Varela, Thompson & Rosch 1991) to (visual) perception that emphasises the role of motor actions and their effect on sensory stimuli.

« 4 » A key consequence of this change is an alternative way of interpreting objects by the unique set of “sensorimotor correspondences” that define the characteristic changes in objective appearance brought about by the agent-object interactions [in the world]. These characteristic correspondences – relating the movement of any object relative to the agent – define its sensorimotor dependencies [qua world]; an agent's practical knowledge of these sensorimotor dependencies constitutes its visual experience.

« 5 » Thus in O'Regan and Noë's sensorimotor theory we have a rich, testable, psychological theory that accounts for why our conscious experience of the world appears as it does, a theory that Beaton suggests fits per-

fectly with DR (§3); this combination forming the foundation of his composite account of phenomenal perception, sensorimotor direct realism (SDR).

« 6 » Although I am broadly sympathetic to the SDR approach Beaton outlines (as I am to ST), it seems to me that at least one of the challenges that has been levelled at ST also appears unresolved in SDR: the challenge of, what I term, “phenomenological determinism,” whereby our phenomenal experience of the world is uniquely determined by our sensorimotor coupling to it: “perceiving is the same thing as engaging in (or being poised to engage in) meaning-filled, physical action in the world” (§14); and with respect to colour, “to perceive a colour is to perceive (to pick out, to master the existence of) the constancy in all this change (change in actual and available interactions)” (§13).

« 7 » Phenomenological determinism is problematic for both ST and SDR as, if phenomenal experience is merely contingent upon exercising the appropriate sensorimotor profile (in interaction with the world), it implies a broad degree of promiscuity regarding the set of systems that are able to have perceive sensation. Put baldly, any system (biological or say, robotic) that exercises the right profile will undergo the same perceptual experience. As Andy Clark and Josefa Toribio wryly observed in their response to O'Regan and Noë's magnum opus (O'Regan 2001):

“A good ping-pong playing robot, which uses visual input, learns about its own sensorimotor contingencies, and puts this knowledge to use in the service of simple goals (e.g., to win, but not by too many points) would meet all the constraints laid out. Yet it seems implausible to depict such a robot (and they do exist – see, e.g., Andersson 1988) as enjoying even some kind of modest visual experience. Surely someone could accept all that O&N offer, but treat it simply as an account of how certain visual experiences get their contents, rather than as a dissolution of the so-called hard problem of visual qualia.” (Clark & Toribio 2001: 980)

« 8 » However, in later writings Noë appears to retreat from this position; for example, in *Action and Perception*, he highlights that:

“Nothing in our view committed us to saying that the robot would be perceptually conscious.

All we committed ourselves to is the possibility that the robot could be perceptually conscious if it acquired the relevant practical skills.” (Noë 2004: chapter 7, footnote 12)

« 9 » This point is further finessed in private communication in which he further clarified,

“I am inclined to think that if you start out with a spark of consciousness, the sensorimotor enactive account can explain how you end up with all the varieties of consciousness. But sms [Sensorimotor Theory, M.B.] alone won't tell you how we get started. In my view – and here I do agree with Thompson – you need life to get the ball rolling.” (Noë, 8 October 2013)

« 10 » In contrast, in O'Regan's later writing, for an agent to have the conscious experience of a quality, say redness, the agent must – via its interaction with the environment – both instantiate appropriate:

- *sensorimotor dependencies* appropriate to pertaining to redness; and
- *cognitive access* to the actions being performed, such that the agent may claim, “I am doing this.”

Such “cognitive access” introduces additional requirements for the agent to have a notion of a self and thus, *having knowledge about its own body, mind and social context*, access to the “experienced” quality.

« 11 » Thus, in an attempt to close the “absolute gap” via sensorimotor theory it seems one is obliged to follow either Noë and reify a link between mind and life, or tread in O'Regan's footsteps and insist on the need for additional explicit cognitive access to phenomenal consciousness; in the process over-complicating, and hence losing the elegance of the foundational sensorimotor theory. And it appears that Beaton endorses O'Regan's view:

“A closely related claim, which is perhaps less explicit in the work of Noë, perhaps more so in the work of Kevin O'Regan, but that I would wish to make, is that this action structure, of which I have mastery, is the structure of my experience of the shape of the object.” (§7)

« 12 » Thus a corollary of Beaton's endorsement of O'Regan's later ST is that it implicitly endorses a formal (computation-

al) account of sensation (phenomenal consciousness) and hence is vulnerable to the various critiques of machine consciousness (Bishop 2009a).

« 13 » One argument that I have developed that questions the possibility of machine consciousness, is the Dancing with Pixies (DwP) reductio (Bishop 2002, 2005, 2009a, 2009b). Its underlying thread derives from positions originally espoused by Hilary Putnam (1988), Tim Maudlin (1989) and John Searle (1990), with subsequent criticism from David Chalmers (1996) amongst others (for early consideration of these themes see the special issue “What is Computation?” of *Minds and Machines*, Harnad 1994).

« 14 » In the DwP reductio, instead of seeking to secure Putnam's claim that “every open system implements every finite-state automaton” (FSA) and hence that “psychological states of the brain cannot be functional states of a computer,” I establish the weaker result that, over a finite time window, every open physical system implements the execution trace of an FSA Q on a given input vector (I).

« 15 » That this result leads to panpsychism should be clear as, equating FSA Q(I) to a finite computational system that is claimed to instantiate phenomenal states as it executes, and employing Putnam's state-mapping procedure to map a series of computational states to any arbitrary non-cyclic sequence of states, we discover identical computational (and ex hypothesis phenomenal) states lurking in any open physical system (e.g., a rock); little pixies (raw conscious experiences) “dancing” everywhere.

« 16 » Baldly speaking, DwP is a simple reductio ad absurdum argument to demonstrate that: *if* (the assumed claim is true: that an appropriately programmed computer really does instantiate genuine phenomenal states) *then* (a vicious form of panpsychism is true).

« 17 » However, if against the backdrop of our immense scientific knowledge of the closed physical world and the corresponding widespread desire to explain everything ultimately in physical terms, we are led to reject panpsychism, then the DwP reductio suggests computational processes cannot instantiate phenomenal consciousness.

« 18 » If I am correct, the DwP reductio highlights the continued phenomenological promiscuity of O'Regan's (contra Noë's) conception of sensorimotor theory (Bishop 2014), a version that Beaton's SDR endorses; in my view, both accounts would be stronger if they engaged the “strong-embodiment of brain and body” (and concomitant environmental and social context) more seriously (cf. Varela 1991; Bickhard 1995; Thompson 2007; Deacon 2012; etc.).

Mark Bishop is Director of the Tungsten Centre for Intelligent Data Analytics and Professor of Cognitive Computing at Goldsmiths, University of London. He has edited two collections of essays and published over 170 articles in the field of cognitive computing. Mark is particularly interested in strongly embodied, enactive, embedded and ecological approaches to cognition.

RECEIVED: 13 FEBRUARY 2016

ACCEPTED: 18 FEBRUARY 2016

The Role of External Objects in Perceptual Experience

Bryony Pierce

University of Bristol, UK

bryonypierce/at/btinternet.com

> Upshot • This commentary is broadly sympathetic to the claims made in the target article. I start by questioning whether we can have direct access to an external reality in such a way that our experience is not intrinsically private. I then suggest that the argument for direct realism presented here is inconclusive with regard to whether external objects play a causal or a constitutive role.

« 1 » Michael Beaton presents a well-argued case for combining direct realism and sensorimotor contingency theory, showing ways in which the versions of these views that he defends in the target article lend support to each other. In this commentary, I discuss the extent to which we can have access to an external reality in such a way that the objects of our perceptual experience can be studied from a third-person perspective, disputing the notion that experience is

devoid of raw sensations. The controversial question of whether objects in the external world are constitutive parts of experience is another primary concern. I aim to defend the view that whilst sensorimotor activity, or at least the capacity for sensorimotor activity, is an essential and ineliminable constituent of perceptual experience, we have reason to withhold judgement on whether external objects, even if directly involved in veridical perception in some way, play a causal or a constitutive role in perceptual experience more generally.

« 2 » I take as my starting point a position close to that outlined by Kevin O'Regan (2011), though I remain receptive to further arguments in support of direct realism. O'Regan (2011: 28) says that information processed during visual perception "constitutes a kind of representation," and that seeing is not caused by this representation, but is caused, or rather constituted, by the activity in which we engage when exploring a scene. Whereas O'Regan stops short of claiming that external objects are constituents of experience, as do I, Beaton (2014: 155) describes perceptual states as "ways of interacting with the world" in which we have direct access to external objects that play an essential, constitutive role in experience. Beaton says that "the fluent access which I have to the detail in the world involves that worldly detail itself. I cannot have access to the detail, if the detail is not there" (ibid: 164).

« 3 » There are two main concerns I want to raise about Beaton's defence of direct realism: firstly, whether access is indeed direct, and secondly, whether a causal explanation would suffice in lieu of the constitutive claim presented, in which experience is fundamentally world-involving. My first point is that access may not be direct in the sense required for veridical experience of external reality. Beaton acknowledges that the structure of the world as we perceive it depends, in turn, on the way our understanding is structured (§4). I would argue that the range of colours and sounds we can detect, and how they seem to us, is determined not only in terms of their functional properties, which would normally be tracked effectively by sensorimotor skills, but also in terms of qualitative features that vary inter-subjectively. Beaton denies

the existence of "raw feels": "there are no raw sensations of colour" (§13), but sensorimotor activity arguably does not account adequately for the qualitative character of colour perception all the way down to the most basic of sensations; it only explains differences in experiences of colour relative to one another across a range of contexts relating to functional roles. Whether we can get full-blown qualitative character from sensorimotor skill acquisition alone is debatable.

« 4 » What O'Regan calls "raw feels," Peter Fazekas calls "monadic markers" (Fazekas 2011: 24). A monadic marker is defined by Fazekas as an "unstructured representation," one of the fundamental building blocks of perceptual experience, a representation with no constituent structure of its own, which starts out as an arbitrary symbol and acquires meaningfulness when it represents that which acquires significance for the perceiver. I am not offering an alternative explanation of raw feels, other than to suggest tentatively that something like Fazekas's monadic markers could underlie the range of qualitatively differing sensations experienced as sensorimotor skills are acquired. I raise this point as a potential problem, if we accept that there are raw sensations, for Beaton's claim that all aspects of experience are accessible for study from the third-person perspective (§9).

« 5 » Let us pursue this question further by considering the example of intra- and interpersonal variations in colour perception (cataracts can affect factors such as whether colours previously appearing to differ significantly can be distinguished from each other), or of adaptation to inversion goggles (Kohler 1964), when our experiences of external objects alter and diverge from those of others. A case can perhaps be made for the accessibility of objects of perception for study by taking features affecting the nature of vision or other sensory modalities into consideration (after all, we are able to establish that cataracts have an adverse effect on the ability to distinguish colours, or to study the effects of wearing inversion goggles experimentally). However, there seems intuitively to be a residual element of experience that is accessible only from the first-person perspective. This becomes clearer, perhaps, if as well as cases of visual impairment affecting colour perception, we consider cases such as hallucinations in patients unable to give reliable

verbal reports, where a third-person perspective seems unlikely to gain access to aspects of the experience that have no discernible effect on behaviour, as well as bearing no direct relation to objects in the immediate environment. Even the qualitative character of quasi-visual experiences produced by tactile sensory substitution systems (Bach-y-Rita 2004) introduces types of experience whose nature is arguably inaccessible to those who have no first-person knowledge of the extent to which the sensations produced resemble vision as opposed to other sensory modalities. Beaton's account might need to be clarified a little more before we can judge how well it deals with objections such as these.

« 6 » A second concern is that (direct or indirect) access to worldly detail could arguably be provided by an ongoing causal relation between the external world and the perceiver. Although Beaton presents fairly extensive and persuasive arguments in support of the constitutive nature of external objects, they seem to be inconclusive, unless we share his intuitions. I am in agreement that perceptual experience depends on sensorimotor activity for its qualitative character, and find the arguments presented by O'Regan, Erik Myin and Alva Noë in their 2004 paper particularly compelling, but resist the further claim made in the target article that external objects are constitutive of perceptual experience, favouring a causal account.

« 7 » I believe a causal account is best placed to explain cases of hallucination. In these cases we can view causation as indirect, temporally extended and involving a recombination of [elements of] the original stimuli: what is perceived in a hallucinatory experience is not entirely novel and unrelated to all prior experience; it contains elements of previously encountered external objects, either roughly as originally perceived, as when a person with Charles Bonnet syndrome believes they are seeing their deceased partner in the room (Chaudhury 2010), or reconstituted/distorted in some way. So whether a theory relying on causal explanation would be disjunctivist, in the sense of treating object-involving experience as fundamentally distinct from non-object-involving experience, depends on whether the introduction of temporally extended causation of perception, in which recombination plays a part,

during hallucination or dreaming, amounts to a *fundamental* difference in the nature of experience. Thus, in cases of hallucination, dreams, and so on, the capacity for sensorimotor activity may be sufficient for experience “as of” an apple, in the absence of a concurrently present apple, although some apple or apples previously encountered play a necessary, but indirect, causal role in producing the illusory experience.

« 8 » Lastly, I am mindful of Susan Hurley’s comments on the causal-constitution debate: “If extended multidimensional dynamic patterns provide the best explanation of quality type in some cases, why assume that external dimensions are merely causal while internal are constitutive?” (Hurley 2010: 126). She also advocates the adoption of a bottom-up approach, guided by the explanatory value of psychological accounts. (I note that Beaton’s account does just this, in drawing on empirical research.) At the same time, I share Fred Dretske’s concerns about the false dichotomy, widespread in the literature on philosophy of mind and action, between internal and external causes of action. Dretske (1988: 22) says that there is “no hard and fast line separating internal from external causes,” and that “by tracing the causal sequence far enough back in time, one can, sooner or later, find external causes for *every* change or bodily movement.” The boundaries thus become blurred in causal explanation, depending on how causes are individuated, and distinctions break down. This leads me to question, by analogy, whether what is or is not a constitutive part of perceptual experience might sometimes be relative to individual, context-dependent conceptualisations and/or indeterminate, or whether it might be contingent upon coherence with some further (psychological or other) account of perceptual experience.

Bryony Pierce recently completed her PhD at the University of Bristol. Her doctoral thesis was on The Role of Consciousness in Action and she has published papers on consciousness, philosophy of action and experimental philosophy. She is a Founder Member of Experimental Philosophy Group UK, and a former member of the European Science Foundation CNCC “CONTACT” research group.

RECEIVED: 13 FEBRUARY 2016

ACCEPTED: 17 FEBRUARY 2016

How Far Can Sensorimotor Direct Realism Go?

Matteo Mossio

IHPST (CNRS/Paris1/ENS), France
matteo.mossio/at/univ-paris1.fr

> Upshot • The target article convincingly argues in favor of the idea that the sensorimotor account of perception provides a positive scientific context for direct realism. In some cases, however, perception and experience do not seem to fit easily with sensorimotor direct realism. This raises a question of scope that requires further elaboration.

« 1 » In his target article, Michael Beaton convincingly argues in favor of the idea that the sensorimotor account of perception provides a positive scientific context for direct realism. Beaton’s central claim is that, according to the sensorimotor account, perceiving consists of engaging in sense-making interactions with the world, through which sensorimotor couplings are established. The practical understanding (the “mastery,” which does not require explicit knowledge) of these couplings is the same thing as perceiving. The perception of an object or an event is an action, is the very fact that an agent is realizing a sensorimotor coupling with that object or event.

« 2 » Beaton argues that, in this conception, there is no need to postulate that perception is mediated by an (internal) representation of the (external) world. Agents *directly* perceive the apple by interacting with it, they do not perceive a representation of the apple. Moreover, the sensorimotor account vindicates a form of perceptual realism insofar as what is perceived – the couplings – are *real* structures in the sense of being open to third-person objective scientific examination. As a result, the sensorimotor account of perception carries on a form of direct realism.

« 3 » As Beaton appropriately emphasizes, direct realism should be distinguished from naïve realism, i.e., the idea according to which the object of perception is mind-independent. By construction, indeed, the sensorimotor account claims that what is perceived is co-constituted by the

agent, which does not have access to an independent reality: therefore, the account is realist, direct, but not naïve. In this respect, it integrates insights coming from both realism and constructivism, without contradiction.

« 4 » I find the core theoretical framework philosophically compelling and scientifically fecund. Yet, it seems to me that sensorimotor direct realism still has to deal with several challenges that all have to do with its *scope*, i.e., the set of capacities and phenomena to which it is supposed to apply. Indeed, there are cases of perceptions and experiences that do not seem to fit easily with the sensorimotor account and specifically with its version in terms of direct realism. In the target article, Beaton does discuss most of these cases, but I think that these issues require further elaboration in order to assess adequately the prospects of the theory.

Non-veridical experiences

« 5 » The first set of uncooperative cases includes what Beaton labels “non-veridical experiences,” i.e., all those experiences that do not correspond to perceptions of present objects, such as imaginations, hallucinations, illusions, or dreams. Non-veridical experiences are supposed to be a challenge for direct realism insofar as they occur in the absence of the “real” object.

« 6 » One straightforward way of handling these cases would be to restrict the application of sensorimotor direct realism to veridical experiences, i.e., to perception. Sensorimotor direct realism, one could argue, is a theory of perception, not a general theory of cognition. As a consequence, it might be the case that while perception does not involve representations, other cognitive capacities do.

« 7 » Beaton seems to adopt a different strategy. He does emphasize the differences between veridical and non-veridical experiences, the former being, among other things, “richer” than the latter in the sense that they continuously transcend us, outrun us, and surprise us. This seems to me well taken, although Beaton wants to maintain that sensorimotor direct realism would *also* be able to account for non-veridical experiences. According to him, the sensorimotor structures underlying such experiences are

different, and yet “relevantly similar” to those involved in perception.

« 8 » In my view, this line of thought is problematic insofar as it seems to lead to the idea that non-veridical experiences also rely on sensorimotor couplings. But how should these couplings be understood in these cases? In veridical perception, the meaning of the two terms to which the “sensorimotor” term makes reference, at least to a first approximation, is clear: “sensory” refers to the sensations impinging on our organism, while “motor” is about the movements made by the same organism. What kind of sensorimotor couplings would be established in the case of non-veridical experiences? Either there are none, or the meaning of the concept is stretched and altered so as to become applicable: in both cases, I think that sensorimotor direct realism is in serious trouble with regards to these capacities. A clarification and probably a theoretical development seem mandatory.

« 9 » The trouble goes even deeper, for the following reason. Even if it were able to provide a convincing explanation of non-veridical experiences, sensorimotor direct realism should still explain how an agent could be able to establish the “relevant similarities” mentioned above without invoking some kind of representation. If an agent knows that she is currently imagining an apple, that is because she is comparing the (supposed) ongoing sensorimotor coupling with a different one that she would have experienced in perceiving an apple. But this second coupling is not realized, it is merely potential. So with what kind of structure is the agent comparing the ongoing mental image of the apple?

« 10 » As for the cases discussed below, non-veridical experiences raise, therefore, the question of the scope of sensorimotor direct realism. Should it be understood as a theory of perception, while a broader version of the sensorimotor account, which would presumably include some form of representation, would apply to non-perceptual experiences? I would like to underscore that, although it would include representations, such a broader sensorimotor account would still be original and innovative with respect to classical representationalist approaches, insofar as it would shift the focus

onto *sensorimotor* representations, in conformity with its inherent interactive and constructivist dimension.

Perceptions relying on potential sensorimotor couplings

« 11 » A second set of recalcitrant cases are those in which perception (i.e., veridical experience) does not seem to rely on the mastery of ongoing sensorimotor couplings. A relevant example is discussed by Beaton in §8 of the target article. There, he discusses the situation in which an agent perceives that a distant tree has the same size as a closer tree. As Beaton explains:

“In seeing a distant tree to be the size that it is, one is understanding (practically) that if one moved closer to it, then the reaching and looking movements necessary to delimit its shape would be exactly the same as those required to delimit the shape of a closer tree of its size.” (§8)

« 12 » This explanation sounds convincing from a general sensorimotor perspective, but it seems at odds with its interpretation in terms of direct realism. The reason is that an actual, ongoing perception (the distant tree being as big as the closer tree) would consist of the mastery of a potential, non-ongoing sensorimotor coupling (the way in which the shape of the distant tree *would* change with the reaching and looking movements): the agent perceives even though she is not engaged in an interaction with the environment. In a first approximation, this is in conflict with the definition of perception given at the beginning.

« 13 » As with the non-veridical experiences discussed above, it seems that sensorimotor direct realism has difficulties in accounting for these perceptual capacities. In §§41f, Beaton does advocate the idea that sensorimotor direct realism could pertinently apply to perceptions relying on potential sensorimotor couplings and, more generally, to perceptions relying on counterfactual situations. In my opinion, however, his argument is not compelling precisely because it does not deal with the main issue, which is the *hiatus* between an ongoing perception and a potential coupling.

« 14 » If the agent is not actually (here and now) engaged in a sensorimotor inter-

action, then sensorimotor direct realism implies that she does not perceive, by definition. Perception *is* engaging in sensorimotor interactions. Vice-versa, if one wants to maintain that there is indeed perception, then it remains unclear how sensorimotor direct realism can consistently account for it. Again, one solution would be to redefine the scope of the sensorimotor direct realism even *within* perceptual capacities. If their advocates do not want to follow this path, then additional justifications seem to be required.

Non-sensorimotor perceptions

« 15 » The third and last set of noncompliant cases are those perceptual situations that do not seem to rely on sensorimotor couplings at all. These cases constitute a general and classical objection to the sensorimotor theory as a whole (and particularly to its constructivist dimension) and not just to his interpretation in terms of direct realism. As a matter of fact, there seem to be perceptual experiences that do not require being engaged in a sensorimotor coupling. Consider the case of a person who is listening to music with headphones. In this situation, she is actually perceiving the music even though there is no coupling between the sensations and her movements: the sensations do not vary with the movements, the two being completely decoupled. More generally, hearing does not seem to require sensorimotor couplings as a necessary condition, at least in humans.

« 16 » Again, it might be possible to apply the sensorimotor account to these cases by broadening the meaning of the term “sensorimotor,” for instance by claiming that when one listens to music with headphones, there is a sensorimotor coupling between sensations and, say, brain dynamics. Yet that could be not just an extension but probably a useless dilution of the approach, which would lose its specificities and explanatory power. Beaton does not explicitly deal with these cases in the target article. Yet it seems to me that any sensorimotor theory of perception and experience should make explicit whether and, at least tentatively, how it intends to include them within the framework. That would greatly enrich the theory and avoid useless debates.

Conclusion

« 17 » Beaton's article convincingly argues in favor of a sensorimotor-centered direct realism in perception. There are a number of situations in which perceptual capacities could be adequately explained in terms of the engagement of an agent in sensorimotor couplings with the world, without appealing to representations. Yet my comment was aimed at raising one main question: How far can sensorimotor direct realism go? As a matter of fact, the sensorimotor theory seems unable to account for some kinds of perceptions and experiences, while embracing direct realism. That raises

an issue of scope that calls for further elaboration.

« 18 » I can envision two strategies for taking up the challenge. Either the scope of sensorimotor direct realism is restricted to (some kind of) perceptual experience, which raises the question of whether the sensorimotor theory should embrace some (original) form of representationalism to account for the cases in which direct realism does not apply. Or an adequate justification is provided to show how sensorimotor direct realism can apply to difficult situations and, thereby, that it constitutes a general theory of experience. Both strategies could

be pursued; let us see how the philosophical and theoretical debate, hopefully nourished by future experimental results, will deal with them in the future.

Matteo Mossio is a research fellow (tenured) at the *Centre National de la Recherche Scientifique*, and full member of the *Institut d'Histoire et de Philosophie des Sciences et des Techniques (IHPST)* in Paris, France. His research interests mainly focus on the principles of biological organization and autonomy, and their relations with cognitive capacities.

RECEIVED: 16 FEBRUARY 2016

ACCEPTED: 19 FEBRUARY 2016

Author's Response The Personal Level in Sensorimotor Theory

Michael Beaton

> **Upshot** • I offer responses to the commentaries on my target article in five short sections. The first section, about the plurality of lived worlds, concerns issues of quite general interest to readers of this journal. The second section presents some reasons for rejecting "enabling" as well as "constitutive" representational approaches to understanding the mind. In the remaining three sections, I clarify aspects of sensorimotor direct realism relating to the self, qualia, counterfactuals, and the notion of "mastery."

An introductory comment

« 1 » I wish to thank the authors of the commentaries for their thoughtful and helpful responses to my target article. It is pleasing to note that the commentators were overall rather sympathetic towards my proposals, even though I presented a philosophical and scientific approach to perception that is very different from that taken in most mainstream cognitive science today.

The plurality of worlds

« 2 » I would like to thank **John Stewart** for a particularly careful commentary. His points are well-made and well-taken; I

would not have been able to make them myself in the same way, and they complement my target article well. Nevertheless, I wish to try to defend myself on the points at which **Stewart** quite rightly pushes me.

« 3 » **Stewart** is correct that the *umwelt* of a tick and the *umwelt* of an oak tree are quite different from each other (§2), and that neither will ever see the world as the other sees it. Nevertheless, I would suggest that this is a cognitive limitation of oak trees and ticks that is not sufficient to show that these two forms of life do not, in fact, share a world. Two issues are raised at this point: whether agents share a world, and whether or not they are aware that they do. I will address the latter issue first.

« 4 » A seagull has a quite different *umwelt* from mine, yet it sees me as an agent, as I do it: it understands at least some of my motivations (though it misunderstands others), as I do its. Does a tick, or an oak tree, view me as an agent? I suppose not. Do at least some insects view me (at least implicitly, at least some of the time) as an agent, given the way in which some of their responses to me are structured (albeit that these are evolved, not learnt, action structures)? Yes, I suspect so. Do many higher animals view me as an agent, as I do them? Yes, certainly. When agents can manifestly see each other as agents (which certainly seems to be the case as between us and many higher animals), I think there must be less objection to the claim that they and we share a world, in some important sense.

« 5 » However, even in the case of the tick and the oak tree, where they certainly cannot see each other as agents, I do not think that their worlds are completely incommensurable with one other. They each have a world structured around basic positive and negative valence, at the least, as does any agent. My target article concentrated mainly on human experience. **Stewart** has said more, and better than I could, about the experience of much simpler beings.¹ Nevertheless, I persist in the claim that, in the end, the tick, the oak tree, and I all live in the same, shared, world;² despite that fact that we experience very different parts of it, very differently; and despite the fact that not all of us can recognise that we do share a world. I think that this claim is compatible with (indeed, follows from) the otherwise somewhat relativist and idealist tone of my approach, precisely because I think there is *some* overlap between the mental lives³ of

1| However, I suspect that Jakob von Uexküll's description of the tick, which **Stewart** endorses (§2), probably radically underestimates the behavioural range of the tick (I suspect that **Stewart** might agree, however).

2| Indeed, it seems that von Uexküll might agree, given that he uses the metaphor of partially overlapping soap bubbles for his *umwelten* (von Uexküll 1957: 29).

3| Of course, overlap between mental lives only entails overlap between worlds on a view in which the shape of mental lives determines the shape of worlds, but that is exactly the view that I, **Stewart** and von Uexküll endorse.

all agents; albeit that certainly not all agents can recognise this overlap.

« 6 » **Stewart** additionally suggests that my privileging of physics is another sign of not-so-latent objectivism on my part (§4 and *passim*). Once again, his points are well-made and well-taken, but I do not wish to retract what I said. I might put it this way. I think that physics examines aspects of the structure of the life of any creature; aspects that are very implicit and very deep, but nevertheless omnipresent. General relativity and quantum mechanics, for instance, are the two most well-tested, quantitatively successful scientific theories we have ever had. Neither quantum mechanics nor general relativity has ever been shown to be even slightly wrong, in any empirical measurement, up to many, many decimal places of accuracy. (I certainly agree that these theories may, nevertheless, eventually be overthrown; and, indeed, that it is almost universally thought that they must eventually be overthrown, or at least in some fundamental way revised, due to apparent incompatibilities between them.)

« 7 » What does this have to do with the life of the oak tree or the tick, or indeed (**Stewart's** additional example, §3) the peasant farmer? Nothing, in a sense; for, of course, none of these agents are concerned *with* the facts revealed by relativity and quantum mechanics. Yet everything, in another sense; for the actual structure of these agents' lives in the world does (and must, to the very best of our knowledge) accord with these regularities that we have discovered. At the very, very fine level of detail, the way in which the tick, the oak tree, and I move *is* affected by these theories;⁴ and affected in ways that need not, but sometimes can, have full-blown macroscopic effects. Science, and physics in particular, is all about discovering such very abstract regularities in the world. "Abstract" in precisely the sense that most of life, including most of human life, has nothing to do with being concerned *with* these regularities. But, if I am right, these regularities are nevertheless present deep in the structure of *how* we all live. This, I would argue, is a more sophis-

ticated way of clarifying why physicists are quite right to claim that what they study is, in a sense, privileged. Just as **Stewart** is quite right to claim that, in another sense, it is not.

« 8 » **John Pickering** also addresses the plurality of worlds. He offers the strong endorsement that:

“In [the] case [of animals], the variety of direct realism advocated by Beaton is plausible. Indeed, so much so that it would be otiose to suggest anything else.” (§10)

« 9 » Nevertheless, he then goes on to make his central claim that, for human observers, our “creative,” “metaphorical,” “culturally shaped” ways of interacting with the world are “far from direct” (§16). Perhaps we are talking at cross purposes here, but I dispute **Pickering's** claim, in the sense in which I mean “direct” in my target article. I fully agree that the world that a human inhabits is fundamentally shaped by culture, symbol use and metaphor. However, I would reject any claim that we layer such interpretation onto some simpler layer of perception (that we perhaps share with animals). On the contrary, I would agree with exactly what James Gibson says, in a quote that **Pickering** himself offers (§14):

“the real postbox (the only one) affords letter-mailing to a letter-writing human in a community with a postal system. This fact is perceived when the postbox is identified as such” (Gibson 1979: 130)

« 10 » Gibson, I think, means what he says. The postbox, as such, is *perceived*. It is our perception itself that is deeply culturally modulated, not just some further layer of interpretation that occurs after perception. We (directly) perceive postboxes, as such, by engaging in a richly culturally modulated, enactive dance with them. Indeed, echoing Heidegger (quoted in my own target article, §27), the “postbox-ness” of the postbox is much closer to us than any details of its three-dimensional shape. Far from “interpreting” something simpler (that we might be supposed to perceive more directly) we actually have to do work to recover the allegedly “simple” properties of what we perceive (as any artist knows well). Nevertheless, sensorimotor theory as I have laid it out makes

explicit certain non-obvious regularities of action that are necessarily involved in perceiving things – including postboxes – as having certain “simple” properties (such as shape or colour), at all.

« 11 » The above points are relevant to a question that **Hugh Gash** poses in his commentary:

“Is DR's position on ‘reality’ close to RC's if it is clear RC does acknowledge an interface with ‘objectivity in parenthesis?’” (§6)

« 12 » I understand **Gash** to be asking whether or not the “reality” of my position is actually the same thing as what he terms “objectivity in parenthesis.” It is quite correct to say that my position's “reality” is fundamentally and irrevocably cognitively structured. Radical constructivism equally emphasises that an agent's world is fundamentally and irrevocably cognitively structured (Glaserfeld 1991). For all that, “reality”⁵ on my position goes beyond us, surprising us, confirming or denying our expectations, and so on.

« 13 » Thus, my answer to **Gash** is that direct realism (DR) would be very close indeed to radical constructivism (RC), *if* it was accepted that radical constructivism acknowledges an interface with “reality” as I have tried to describe reality. But it is far from clear to me whether Ernst von Glasersfeld's radical constructivism (Glaserfeld 1991) *can* be consistently read as acknowledging an interface with anything like the intersubjectively shared “reality” of which I talk. Radical constructivism shares with representationalism the idea that whatever cognitive structures we have are related (if they are related at all) to an external world that we can never directly know. However, as **Gash** says:

“A critical difference between RC and both representationalism and direct realism (DR) is that RC denies that it is possible to make claims about the relation between experience and ‘reality.’” (§4)

« 14 » In the target article, I say that “we have no way of accessing the world, except via our cognitive structures” (§68). For the

4 | The recent detection of gravity waves confirms that general relativity does indeed correctly predict minute, but empirically observable, microscopic consequences, as well as macroscopic ones.

5 | Or reality without quotes, as I would prefer to say at certain places, given that I have tried to defend the validity of the notion if used carefully enough.

radical constructivist, the latter part of this quote would simply express why we cannot access the world. This is because, for the radical constructivist, it is clear that our cognitive structures do not contain parts of the external world. Therefore, it is equally clear that we cannot directly access the external world, if our cognitive structures are our only potential means of access to it. For me, however, the latter part of my quote expresses *how* we access the world. A central thesis of my target article is that our cognitive structures literally do contain parts of the external world (direct realism is a radical, but serious, position), and hence that we can and do access the world. I would emphasise that this is not meant to belittle the claim that our world is cognitively structured, but rather to fit with it. John McDowell believes that this is Kant's view. So do I. For all these reasons, whilst I do not know how to fully answer Gash's question, I do believe there is certainly more than enough room for continued fruitful dialogue here.

Representation, representation, representation

«15» I agree with David Silverman that McDowell (1994, 1996: 55) is correctly read as endorsing what Silverman calls “enabling representationalism”: the position that positing internal representations may be useful to explain the inner workings of the brain. Silverman also correctly says that McDowell rules out what Silverman calls “constitutive representationalism”; that is, McDowell strongly rejects the claim that the contents of our personal level experience are in any way to be identified with the contents of any such “enabling,” sub-personal representations (McDowell 1994). Now if someone reads my work, or McDowell's, or Silverman's, and thereby comes to understand why it is a mistake to equate having an experience with having an internal representation (with the right content, playing right functional role) then I am already happy.

«16» But actually, I would wish to take what is arguably a stronger line than McDowell's, here,⁶ and certainly a stronger line

than Silverman's, who argues strongly for “enabling representational” explanations, even whilst agreeing with me that “constitutively representational” explanations cannot work. I accept that internal representational explanations can do some work, as far as they go. But I strongly suspect that they do not go far enough: that an explanation of cognition or perception in terms of internal representations will always, necessarily, miss the possibility of perfectly good (and, in important cases, correct) alternative explanations as to how a given task is performed. To insist on a representational explanation of a given cognitive or perceptual task is effectively to rule out silently, at the outset, the possibility that the world itself is a constitutive part of how the task is performed. However, as I argued in my target article (§§19f), enactively inspired cognitive science has already given us many examples in which interesting, non-trivial, cognitive and/or perceptual tasks are performed in ways that are fundamentally world-involving (and thus, at the very least, not fully representational). Nothing that we know rules out the possibility (indeed, I would say, the likelihood) that our own perception is like this, in various fundamental ways.

«17» Agreeing with Silverman (§4), I would once again emphasise that sensorimotor theory is a scientific theory, not just a philosophical one, precisely because it strongly suggests that these other types of explanation of perceptual experience will be fruitful in understanding human perception (far from being ruled out almost *a priori*, as some representationalists seem to feel). The scientific work on perception being carried out in Kevin O'Regan's lab (for an overview, see O'Regan 2011) also strongly bears out the claim that this is a fruitful scientific framework in which to work.

«18» I would like to make one further point about representation. It is confusing, but bear with me. The point is that McDowell does not balk at using the term “representation” at the personal level (e.g., McDowell 1996: 162). However, I must clarify that, in doing so, McDowell is absolutely not falling

into the trap of supposing that the contents of our mental states are carried by internal representations. Instead, when McDowell uses the term “representation” in this way, he is using it as an *entirely* personal level concept. Thus, when we say that someone's experience “represents” a tree, in this sense, then all we are saying is that their experience is either “of” or “as of” a tree. They are either veridically seeing the tree,⁷ or else they are having an experience that is, for them, experientially as if a tree were present, even though it is not.⁸ That is all. There is no further implication, whatsoever, that internal states with matching content are required to explain this personal level phenomenon. Thus, McDowell is actually using the term “representation,” at the personal level, in exactly the same sense in which Kant (1996) uses the term “representation,” or *Vorstellung* in the original, which is sometimes translated into English as “presentation,” precisely in order to avoid any misleading impression that it has anything to do with sub-personal states. This way of using the term “representation” might very well be misleading since, in the day-to-day English usage of the word “representation,” one thing refers to something else that it is not. Here, I just want to point out that a considerably different usage, which is perhaps misleading, but (for the reasons just stated) is genuinely not perniciously representationalist, exists in certain parts of the relevant literature.

The self in sensorimotor theory

«19» Mark Bishop (§12) takes it to be the case that I support O'Regan's proposal (O'Regan 2011) that sensorimotor theory should be supplemented with Thomas Metzinger's self-model theory (Metzinger 2003), and that I support O'Regan's related claim that conscious perception involves contemporaneous self-knowledge of what one is doing. This is incorrect. I mentioned my

7| This means correctly seeing the tree, when and because the tree is there. When veridically seeing, the tree itself is a constitutive part of the experience, according to McDowell's and my position; however, this latter claim is neither affirmed nor denied merely by using the notion of “representation” or “presentation” correctly, in this Kantian way.

8| Or is not in the right way (Noë 2003).

6| Though I do not think that what I say is necessarily incompatible with McDowell's fairly guarded position on all this; McDowell is, I think, correctly read as simply stating that nothing he

says actively rules out “enabling representational” explanations, without ever positively endorsing the claim that such explanations are necessarily good, helpful or even useful, in any given case.

position on O'Regan's proposals in footnote 3 of my target article: "I would have reservations about some of the philosophical additions to the theory that O'Regan (2011) has recently proposed, in particular around the correct treatment of the self." Here I wish to restate clearly that I reject O'Regan's recent additions to sensorimotor theory for much the same reasons that Bishop does: they are overly cognitivist and internalist.

« 20 » Indeed, in my target article, I attempted to outline a quite different treatment of the self in sensorimotor theory, based around Sydney Shoemaker's (1996) analysis of self-knowledge. In Shoemaker's theory, there are no self-models,⁹ and there are no detectors of internal states. Instead, we are only concerned with personal-level, rational (i.e., reason-respecting) connections between mental states. If I see food and I am hungry, then, all other things being equal, I will want the food and try to get the food. A creature that is hungry does not have to "detect" its inner state of hunger in order to want food; being hungry simply is being motivated to act in the appropriate way(s) in response to food.¹⁰ According to Shoemaker's account, one similarly does come to know that one has some mental state (of hunger, pain, perception, experience, knowledge, etc.) by "detecting" it. Instead, to learn the meaning of a concept¹¹ such as pain or hunger, in self application, is to learn to say (or think) that one is in the relevant state, as and when one is.¹² No inner detection is needed

for such acts of introspection, any more than inner detection of the feeling of hunger is needed in order to act hungrily.

« 21 » Shoemaker's is an account of self-knowledge, not an account of conscious feeling. Creatures that are far too simple to have the concept of pain can certainly still feel pain, on Shoemaker's account as on mine, and even creatures that have the concept of pain do not need to apply it to feel it. O'Regan, on the contrary, suggests that a creature must have and apply the (proto) concept of pain (for example) in order consciously to feel pain. I disagree with this. So I disagree with O'Regan on two fundamental counts. I disagree on the correct model of self-knowledge (Shoemaker's vs. Metzinger's); and I disagree on whether active self-knowledge needs to be occurring, right now, in order for conscious feeling to be occurring, right now. For all that, I share O'Regan's instincts in this area to a significant extent. I agree that the correct treatment of the self is an important part of the full elaboration of the sensorimotor theory of perception; and I agree that it is important that conscious states be, at least, the right type of states to be introspectible – the kind of thing that a sufficiently advanced creature could *learn* to introspect. If they were not, I have argued (Beaton 2009b), they could hardly be the conscious states that we spend so much time discussing! For all that, I feel that O'Regan is, unfortunately, currently endorsing a model of self-knowledge (i.e., Metzinger's) that is too cognitivist and internalist ever to be a good match for sensorimotor theory. To be clear, I think that this philosophical point about self-knowledge sits somewhat at the edge of the sensorimotor framework, at least in as much as it might guide scientific work in perceptual psychophysics. For that reason, I definitely do not think that this mistake (as I see it) invalidates O'Regan's scientific work based on the fundamental principles of the sensorimotor framework that, of course, he himself helped to develop.

« 22 » The paragraph that Bishop (§11) read as my endorsement of O'Regan was meant simply to say that I find O'Regan to say more, and more explicitly, than Alva Noë does about the claim that the action-structure of our sensorimotor engagement with the world should be identified with the phenomenal structure of our experience. What I

said in that paragraph was misleading, for the claim in question is absolutely central to sensorimotor theory as I have presented it. Thus, the reader might well be confused as to how I could possibly think that Noë does not make that very same claim. Actually, I *do* think that Noë makes that very same claim; extensively, but arguably largely implicitly. Nevertheless, I should more rightly have said simply that O'Regan gives more and different examples of this point than Noë does, for instance in O'Regan's mathematical work on colour perception (Philipona and O'Regan 2006) and perception of the dimensionality of space (Philipona et al. 2003).

« 23 » The above account of introspection relates to one of the reservations about my view that Bryony Pierce expresses. She worries that my direct realism is not fully compelling, in that the points that I make do not rule out an alternative account on which "access to worldly detail [is] provided by an ongoing causal relation between the external world and the perceiver" (§6). There is a misunderstanding here – though certainly not a trivial one. For whilst my account does indeed state that objects in the world are constitutive parts of experiences, it does not thereby deny that experience involves an ongoing causal relation between the world and the perceiver. On the contrary, according to my account, experience *is* the ongoing causal relation between the world and the perceiver. What Pierce is actually proposing is that there may still be room for an account on which experiences (construed as occurring inside the perceiver) might be only causally related to objects outside the perceiver. That is as may be, and that is not my account. The mistake that I believe Pierce makes is to suppose that my account is straight-forwardly opposed to an account on which there is "an ongoing causal relation between the external world and the perceiver" (§6). It is not.

« 24 » On the direct realist account that I have set out, experiences extend beyond perceivers: the tree that I am looking at is not part of me, but it is part of my experience. An upshot of this, which I explicitly noted in my target article (§§47–50), is that I can introspect things that are not part of me, though they are part of my experiences. This sounds ridiculous on an inner perception account of introspection: of course I cannot perceive, inside me, something that is outside me.

9| At least, not in a subpersonal, computationalist sense of "self-model." I accept the psychological observation that, at least in some cases, we can find that we know ourselves no better than we know a stranger. Thus, certainly, we do have presuppositions about ourselves, some of which are wrong. However, this is a personal-level phenomenon, which should not be conflated with any subpersonal (and quite possibly extended) explanation of the phenomenon.

10| We are not motivated by the hunger. The feeling of hunger is the feeling of being thus motivated. (With the addition, perhaps, of the literal feeling of an empty stomach, etc.)

11| Or proto-concept; see footnote 18.

12| This does not imply that introspection is infallible. It does imply that a failure of introspection is a failure of rationality; but such failures are perfectly possible.

However, I have rejected that account of introspection in favour of an account in terms of reason-respecting transitions in thought.

« 25 » This links to the following rhetorical question in Gashi's commentary:

“However, I find congenial the description of introspection ‘as a self-reflexive transition within an agent’s understanding’ [...] what is this if not a form of inner perception and experience?” (§8)

« 26 » So far from being “a form of inner perception,” introspection can be (and I believe is) exactly what I have set out, here. There are no internal representations in this account. Perception does not involve them; introspection does not involve “looking at” them.

Phenomenological determinism

« 27 » Bishop (§§6ff) is worried that direct realism has the consequence that our qualitative feels (our “qualia”) must be determined entirely by what we perceive since, it seems, there is nothing else in direct realism that could determine them. Bishop feels, as do I, that this cannot be right: that there is more to how the colour red (for example) feels to me than can be determined simply by determining that I am seeing red. Bishop calls this apparent problem with direct realism “phenomenological determinism.” I myself initially felt that direct realism suffered from this problem when I first encountered the theory. I now think that there is, in fact, no such entailment. This is not because I think that direct realism should be patched up by putting back in representations (or any other sub-personal states) painted with qualia. Rather, it is because I think there is, objectively, more richness to our agent-level responses that are specific to a given type of stimulus (the colour red, say) than is determined simply by what we are responding to.

« 28 » Colour is actually a particularly complex case, because colours are not entirely “out there,” but also not entirely “in here” (Thompson 1994). The colours I can see – the distinctions I can make, and the conditions under which I can make them – are determined not just by “what colours things have” (which frequencies of light they reflect and emit), but also by what visual system I have (particularly pertinently,

by the frequency response profiles of the cones in my retina).

« 29 » Nevertheless, given two agents with exactly the same visual system (in this respect: capable of making exactly the same distinctions, and discriminating only exactly the same colours), one can still ask whether these two agents must, necessarily, perceive colour in qualitatively the same way. Despite first appearances, I do not think that direct realism entails any such consequence, for the following reasons.

« 30 » Red reminds me of blood, *inter alia*. It reminded Kant of “heavy cinnabar” (Kant 1996: A101). It also reminds me of traffic lights; and I have been trained to think of it as a symbol for warning and danger (whether or not this is a consequence of red being the colour of blood is a further question). Blue reminds me of the sea, and ice, and the sky. Green reminds me of trees and leaves. Thus, it seems, exactly what I associate with these different stimuli is not fixed, simply by fixing which stimuli I can perceive.

« 31 » Furthermore, not only is learnt association¹³ not fixed by fixing what I can perceive. Neither is affect (i.e., emotional valence). Thus, for example, I find sharp spiked objects naturally somewhat off-putting, and soft fluffy objects naturally somewhat comforting. This seems to me to be a consequence of my kind of embodiment: sharp, spiky objects are naturally likely to be damaging, soft fluffy objects are, typically, naturally less so. Nevertheless, it seems coherent to imagine an alien that finds soft fluffiness quite repellent, and sharp spikiness comforting and attractive.

« 32 » Thus it seems that both affect and association can vary independently of what is perceived, even for two agents with the same type of visual system who are perceiving exactly the same object or property. I have tried to argue (Beaton 2009a) that such variations of affect and association (when considered alongside the details of *what* is perceived on which I concentrated in the target article) are exactly the right kinds of potentially introspectible differences to count as qualia; and even to account (though not completely) for philosophers’ intuitions about inverted spectra. Such properties could indeed dif-

13| Nor innate association, if such a thing exists.

fer, even between two agents with exactly the same type of visual system (in respect of discriminatory abilities) who are perceiving the very same coloured object. This is very reminiscent of inverted qualia, and I would argue that it is certainly sufficient to respond to Bishop’s problem of supposed phenomenological determinism. That said, these are behaviourally detectable differences; there is no space in my theory for behaviourally undetectable, completely private differences in phenomenal states. I reject the claim that such differences are possible.

Counterfactuals and mastery

« 33 » Matteo Mossio (§§5ff) worries that the theory that I have presented may be a good account of perception in cases where the subject actually is interacting with an object, but a bad account of non-veridical experience, such as illusion and hallucination. This is not what I want to achieve. I do not want, for instance, to provide a non-representational account of veridical experience but to have to resort to inner representations in order to account for our non-veridical experiences. The same type of objection has been applied to McDowell, who rather notoriously stated that the world itself is involved in our experience “when we are not misled” (McDowell 1996: 9, 143). But what about when we are misled? What is going on then? What are we misled by?

« 34 » I believe that this is a point where sensorimotor theory can bolster direct realism. I have already tried to explain why, both in my target article (§§38f) and in Beaton (2013), but I will reiterate the point. Normal, everyday science deals in what I will call “counterfactuals” all the time. In what *would* happen if we did some experiment. For example, we do not think that a proton is a proton only if it is (*per impossibile*) acting, at once, in all the ways characteristic of protons. We think that something is a proton as long as we have good reason to believe that it would interact in each of the ways characteristic of protons, if tested. I believe that sensorimotor theory, with this one theoretical extension in terms of counterfactuals (which I have tried to emphasise more, I think, than other authors) can indeed account for non-veridical experience, and can also provide detail, where McDowell could not, about what is going on when we *are* misled.

«35» The issue of counterfactuals is actually closely related to issues about “mastery,” a term that has been something of a bugbear to those who are relatively sympathetic to enactive approaches, but who do not quite “get” what it is that sensorimotor theory brings to the table. Noë and O’Regan have always (again, somewhat notoriously) said that you have to have mastery of¹⁴ the relevant sensorimotor contingencies in order to perceive. It should be noted that the requirement that we have mastered the relevant sensorimotor counterfactuals applies even to veridical perception: it applies when we are sitting stock still and staring at something that is also stock still. Many more primitive animals cannot perceive anything in such a situation. We can. Actually, it turns out that our eyes have to move, slightly, in order to continue to perceive in such cases, but I do not want to have to rely on this kind of fact; I think that would make a weak theory, prone to counterexamples. Instead, what I think is going on here – what O’Regan and Noë have always said is going on here – is that we have *mastered* the relevant sensorimotor contingencies; that we know, quite correctly, what *would* happen *if* we moved our head to the side, or *if* the tomato (Noë’s favourite example) started to rotate.

«36» I will return to the issue of non-veridical experiences in a moment, but firstly I would like to clarify two remaining points about “mastery” in sensorimotor theory. Firstly, mastery is fundamentally norm-involving. It is not enough that some system – a computer with a camera, say – is set up such that these counterfactual sensorimotor contingencies would, in fact, apply. What is required is that the agent *understands* that these counterfactual contingencies apply,¹⁵ and acts (or would act, if appropriately tested) in ways that demonstrate that it understands this. In other terms, this understanding must be fully integrated with the agent’s norms, such that the agent can, and typically will, use this understanding to go about getting what

it desires, and otherwise achieving its goals. Secondly, for reasons that I discuss in more detail in my target article (§10), and also in Beaton (2013), I believe that this sensorimotor understanding is necessarily not explicit, nor symbolic, nor verbalisable (not even in creatures that can verbalise). It is implicit and deep; but it is also incredibly rich and complex; it is a form of genuine understanding. It is integrated with, and crucial to, our more abstract forms of understanding. One might say that it is the base layer of our understanding. But it is still flexible and responsive. It *can* change and adapt,¹⁶ especially when the agent puts its mind to it – as examples such as the work of Ivo Kohler (which I and Noë have emphasised, for this reason) show.

«37» What, then, of non-veridical perception? What are we misled by? What, as Mossio asks (§9), are we comparing our non-veridical experience to, when we only think that we are seeing an apple? I am not trying to pull a sleight of hand when I say that I think these latter are the wrong questions. I think that sensorimotor direct realism can show us what is going on in these cases, even though it does not allow us to answer these questions quite as posed (with all their presuppositions). When we imagine that we are seeing an apple, our sensorimotor action profile is as if we are seeing an apple: if we are appropriately tested, or asked, we will move our hands or eyes as if we were seeing an apple.¹⁷ We can only do this – could only do this – because we know what apples are like; but knowing what apples are like does not mean having a stored sensory image of an apple. It means knowing how to act in these “apple shaped” ways.¹⁸ All very

16| I would accept that we have to allow that some of our sensorimotor understanding was developed over evolution, and that not all of it is plastic in the life of a given agent; but (a) I do not think there will be clear boundaries here, and (b) I think that this is equally true of all intelligence and understanding.

17| Though, as I clarify in my target article, not *exactly* as if, for the world is not there to let all the movements and counterfactual movements be exactly as they would be when actually interacting with an apple.

18| That is, it relies on having the sensorimotor (proto)concept of “apple,” which is something that non-verbal higher animals can perfectly well

well, the reader might say; that is a third-person story about action; but what are we, the subject, misled by when we have such an illusory experience? Well, the one thing we are not misled by is a sensory image of an apple; rather, on the account given here, to hallucinate or imagine an apple simply is to be prone to behave (if tested) in these ways. One’s base-level sensorimotor understanding has become misled (in ways that visual scientists, or psychologists, or pharmacists, might study) such that these apple-shaped ways of behaving seem appropriate. Being thus misled depends on one’s having apple-shaped ways of behaving in the first place. But it does not rely on *comparing* one’s apple-shaped ways of behaving to the ways in which one is behaving now. Rather, to have an illusory apple before one simply is to be prone to act, at the base sensorimotor level, as if there were an apple before one when there is not. This may all sound very behaviourist, but I urge once again that it is “actionist,” it is fundamentally about norms and understanding, not about mere meaningless behaviour. Indeed, this is a very radically “actionist” view in which perceiving *is* acting (and having mastery of counterfactual, or possible, acts). This would make no sense on a traditional view of things. It makes much more sense on a direct realist view, in which some of the actions that we can take fundamentally involve things in the world.

«38» We need to go down this route to avoid representationalism, but it is also the case that we can and should go down this route in order to do justice to experience: there are sufficient materials in sensorimotor direct realism to account for the qualitative, first-person, phenomenally rich experiences of the world (or, sometimes, only “as of” the world) that we all know that we have.

Acknowledgements

I gratefully acknowledge support by the Spanish Government MINECO Project, Reference FFI2014-52173-P.

RECEIVED: 1 MARCH 2016

ACCEPTED: 9 MARCH 2016

have: a practical understanding of what apples are, and what it is possible to do with them.

14| Not the alternative English reading of “mastery,” which would mean “being in the process of mastering,” so let us put that possible misunderstanding aside.

15| Tom Froese (2014) makes this same point while commenting on Anil Seth’s version of sensorimotor counterfactualism (Seth 2014).

Combined References

- Andersson R. L. (1988) Robot ping-pong player. MIT Press, Cambridge MA.
- Bach-y-Rita P. (2004) Tactile sensory substitution studies. *Annals of the New York Academy of Sciences* 1013: 83–91.
- Beaton M. (2009a) An analysis of qualitative feel as the introspectible subjective aspect of a space of reasons. Doctoral Thesis. University of Sussex UK.
- Beaton M. (2009b) Qualia and introspection. *Journal of Consciousness Studies* 16: 88–110.
- Beaton M. (2013) Phenomenology and embodied action. *Constructivist Foundations* 8(3): 298–313.
► <http://constructivist.info/8/3/298>
- Beaton M. (2014) Learning to perceive what we do not yet understand: Letting the world guide us. In: Cappuccio M., Froese T. (eds.) *Enactive cognition at the edge of sense-making: Making sense of non-sense*. Palgrave Macmillan, Basingstoke UK: 153–180.
- Beer R. D. (2003) The dynamics of active categorical perception in an evolved model agent. *Adaptive Behavior* 11: 209–243.
- Berger P. & Luckmann T. (1966) *The social construction of reality*. Anchor Books, New York.
- Bickhard M. H. & Terveen L. (1995) *Foundational issues in artificial intelligence and cognitive science*. North Holland, Amsterdam.
- Bishop J. M. (2002) Dancing with pixies: Strong artificial intelligence and panpsychism. In: Preston J. & Bishop J. M. (eds.) *Views into the Chinese room: New essays on Searle and artificial intelligence*. Oxford University Press, Oxford: 360–378.
- Bishop J. M. (2005) Can computers feel? *The AISB Quarterly* 199: 6.
- Bishop J. M. (2009a) A cognitive computation fallacy? *Cognition, computations and panpsychism*. *Cognitive Computation* 1(3): 221–233.
- Bishop J. M. (2009b) Why computers can't feel pain. *Minds and Machines* 19(4): 507–516.
- Bishop J. M. & Martin A. O. (2014) Contemporary sensorimotor theory: A brief introduction. In: Bishop J. M. & Martin A. O. (eds.) *Contemporary sensorimotor theory*. Springer, Heidelberg: 1–22.
► <http://cepa.info/2525>
- BonJour L. (2013) Epistemological problems of perception. In: Zalta E. N. (ed.) *The Stanford encyclopedia of philosophy* (Spring edition).
<http://plato.stanford.edu/archives/spr2013/entries/perception-episprob/>
- Boxer P. & Kenny V. (1990) The economy of discourses: A third-order cybernetics? *Human Systems Management* 9(4): 205–224.
► <http://cepa.info/2383>
- Braitenberg V. (1984) *Vehicles: Experiments in synthetic psychology*. MIT Press, Cambridge MA.
- Brentari C. (2015) *Jakob von Uexküll: The discovery of the Umwelt between biosemiotics and theoretical biology*. Springer, New York.
- Buhrmann T. & Di Paolo E. (2014) Non-representational sensorimotor knowledge. In: Del Pobil A. P., Chinellato E., Ester Martinez-Martin E., Hallam J., Cervera E. & Morales A. (eds.) *From animals to animats 13*. Springer, New York: 21–31.
► <http://cepa.info/2521>
- Byrne A., Logue H. (2009) *Disjunctivism: Contemporary Readings*. MIT Press, Cambridge MA.
- Carroll L. (1895) What the Tortoise said to Achilles. *Mind* 4: 278–280.
- Chalmers D. (1996) *The conscious mind: In search of a fundamental theory*. Oxford University Press, New York.
- Chaudhury S. (2010) Hallucinations: Clinical aspects and management. *Industrial Psychiatry Journal* 19(1): 5–12.
- Clark A. (2001) *Mindware*. Oxford University Press, New York.
- Clark A. & Toribio J. (1994) Doing without representing? *Synthese* 101(3): 401–431.
- Crane T. (2006) Is there a perceptual relation? In: Gendler T. S. & Hawthorne J. (eds.) *Perceptual experience*. Clarendon Press, Oxford: 126–146.
- Davidson D. (1974) On the very idea of a conceptual scheme. *Proceedings and Addresses of the American Philosophical Association* 47: 5–20.
- Deacon T. (2012) *Incomplete nature: How mind emerged from matter*. W. W. Norton, New York.
- Degenaar J. & Myin E. (2015) Representation-hunger reconsidered. *Synthese* 191(15): 3639–3648.
- Dennett D. C. (1987) *The intentional stance*. MIT Press, Cambridge MA.
- Dewey J. (1960) *The quest for certainty*. Capricorn, New York. Originally published 1929.
► <http://cepa.info/1889>
- Di Paolo E. (2009) *Extended Life*. *Topoi* 28: 9–21.
► <http://cepa.info/322>
- Di Paolo E., Barandiaran X., Beaton M. & Buhrmann T. (2014) Learning to perceive in the sensorimotor approach: Piaget's theory of equilibration interpreted dynamically. *Frontiers in Human Neuroscience* 8: 00551.
<http://journal.frontiersin.org/article/10.3389/fnhum.2014.00551/full>
- Dretske F. (1988) *Explaining behavior: Reasons in a world of Causes*. MIT Press, Cambridge MA.
- Fazekas P. (2011) Cognitive architecture and the epistemic gap: Defending physicalism without phenomenal concepts. *Philosophia* 39(1): 21–29.
- Fernández-Ruiz J. & Díaz R. (1999) Prism adaptation and aftereffect: Specifying the properties of a procedural memory system. *Learning and Memory* 6(1): 47–53.
- Floridi L. (2011) A defence of constructionism: Philosophy as conceptual engineering. *Metaphilosophy* 42(3): 282–304.
► <http://cepa.info/2384>
- Froese T. (2014) Steps toward an enactive account of synesthesia. *Cognitive Neuroscience* 5: 126–127.
► <http://cepa.info/2526>
- Fultot M. F. (2016) Counterfactuals versus constraints: Towards an implementation theory of sensorimotor mastery. *Journal of Consciousness Studies*. In press.
- Gibson J. J. (1966) *The senses considered as perceptual systems*. Houghton Mifflin, Boston.
- Gibson J. J. (1979) *The ecological approach to visual perception*. Routledge, London.
- Glaserfeld E. von (1984) An introduction to radical constructivism. In: Watzlawick P. (ed.) *The invented reality*. Norton, New York: 17–40. Originally published in German as: Glaserfeld E. (1981) *Einführung in den Radikalen Konstruktivismus*. In: Watzlawick P. (ed.) *Die Erfundene Wirklichkeit*. Piper, Munich: 16–38.
► <http://cepa.info/1279>
- Glaserfeld E. von (1991) Knowing without metaphysics: Aspects of the radical constructivist position. In: Steier F. (ed.) *Research and reflexivity*. Sage, London: 12–29.
► <http://cepa.info/1420>
- Glaserfeld E. von (2005) Thirty years radical constructivism. *Constructivist Foundations* 1(1): 9–12.
► <http://constructivist.info/1/1/009>
- Haddock A. & Macpherson F. (2008) Disjunctivism: Perception, action, knowledge. Oxford University Press, Oxford.
- Harnad S. (ed.) (1994) *What is computation?* Special issue. *Minds and Machines* 4(4).

- Harvey I., Di Paolo E., Wood R., Quinn M., Tuci E. & Iridia E. T. (2005) Evolutionary robotics: A new scientific tool for studying cognition. *Artificial Life* 11: 79–98.
- Heidegger M. (1977) The origin of the work of art. In: Krell D. F. (ed.) *Martin Heidegger: Basic Writings*. Harper and Row, New York: 139–212. German original published in 1935.
- Heidegger M. (2008) The fundamental concepts of metaphysics: World, finitude, solitude. Indiana University Press, Bloomington. Originally a lecture course given in 1929/30.
- Hinton J. M. (1973) *Experiences: An inquiry into some ambiguities*. Oxford University Press, Oxford.
- Hume D. (1748) An enquiry concerning human understanding. London. <http://www.gutenberg.org/ebooks/9662>
- Hurley S. (2010) The varieties of externalism. In: Menary R. (ed.) *The extended mind*. MIT Press, Cambridge MA: 101–153.
- Hutto D. D. & Myin E. (2013) Radicalizing enactivism: Basic minds without content. MIT Press, Cambridge MA.
- Izquierdo E. & Di Paolo E. (2005) Is an embodied system ever purely reactive? In: Capcarrere M. S., Freitas A. A., Bentley P. J., Johnson C. G. & Timmis J. (eds.) *Proceedings of the 8th European Conference on Artificial Life*. Springer-Verlag, Berlin: 252–261.
- Kant I. (1996) Critique of pure reason. Translated by Werner S. Pluhar. Hackett Publishing Company, Indianapolis. German original published in 1781/1787.
- Koffka K. (1935) *Principles of gestalt psychology*. Lund Humphries, London.
- Kohlberg L. (1969) Stage and sequence: The cognitive-developmental approach to socialization. In: Goslin D. (ed.) *Handbook of socialization theory and research*. Rand McNally, Chicago: 347–480.
- Kohler I. (1964) The formation and transformation of the perceptual world. *Psychological Issues* 3: 19–134. Originally published in German as: Kohler I. (1951) *Über Aufbau und Wandlungen der Wahrnehmungswelt*. Österreichische Akademie der Wissenschaften. Sitzungsberichte, philosophisch-historische Klasse 227: 1–118.
- Latour B. & Woolgar S. (1979) *Laboratory life: The social construction of scientific facts*. Sage, Beverly Hills.
- Laureys S. (ed.) (2005) *The boundaries of consciousness: Neurobiology and neuropathology*. Elsevier, Amsterdam.
- Levine J. (1983) Materialism and qualia: The explanatory gap. *Pacific Philosophical Quarterly* 64: 354–361.
- Locke J. (1689) *An essay concerning human understanding*. London. <http://www.gutenberg.org/ebooks/10615>
- Marr D. (1982) *Vision*. MIT Press, Cambridge MA.
- Martin M. G. F. (2002) The transparency of experience. *Mind and Language* 4: 376–425.
► <http://cepa.info/2401>
- Martin M. G. F. (2004) The limits of self-awareness. *Philosophical Studies* 120: 37–89.
► <http://cepa.info/2402>
- Maudlin T. (1989) Computation and consciousness. *Journal of Philosophy* (86): 407–432.
- Mayr E. (1974) Behaviour programs and evolutionary strategies. *American Scientist* 62(6): 650–659.
- McDowell J. (1982) Criteria, defeasibility and knowledge. *Proceedings of the British Academy* 68: 455–479.
- McDowell J. (1994) The content of perceptual experience. *The Philosophical Quarterly* 44: 190–205.
- McDowell J. (1996) *Mind and world: With a new introduction by the author*. Harvard University Press, Cambridge MA.
- McDowell J. (2009) *Having the world in view: Essays on Kant, Hegel and Sellars*. Harvard University Press, Cambridge MA.
- Merleau-Ponty M. (1962) *The phenomenology of perception*. Routledge & Kegan Paul, London. French original published in 1945.
- Merlin D. (1993) *Origins of the modern mind: Three stages in the evolution of culture and cognition*. Harvard University Press, Cambridge MA.
- Metzinger T. (2003) *Being no one: The self-model theory of subjectivity*. MIT Press, Cambridge MA.
- Neisser U. (1976) *Cognition and reality*. Freeman, New York.
- Noë A. (1999) Thought and experience. *American Philosophical Quarterly* 36: 257–265.
- Noë A. (2003) Causation and perception: The puzzle unravelled. *Analysis* 63: 93–100.
- Noë A. (2004) *Action in perception*. MIT Press, Cambridge MA.
- Noë A. (2009) *Out of our heads*. Farrar, Straus and Giroux, New York.
- Omnès R. (1999) *Quantum philosophy: Understanding and interpreting contemporary science*. Princeton University Press, Princeton NJ. French original published in 1994.
- O'Regan J. K. (2011) *Why red does not sound like a bell: Understanding the feel of consciousness*. Oxford University Press, Oxford.
- O'Regan J. K. (2012) How to build a robot that is conscious and feels. *Minds and Machines* 22(2): 117–136.
- O'Regan K. Myin E. & Noë A. (2004) Towards an analytic phenomenology: The concepts of “bodiliness” and “grabbiness.” In: Carsetti A. (ed.) *Seeing, thinking and knowing*. Kluwer, Dordrecht: 103–114.
- O'Regan J. K. & Noë A. (2001) A sensorimotor account of vision and visual consciousness. *Behavioral and Brain Sciences* 24(5): 883–917.
► <http://cepa.info/2285>
- Peacocke C. (2001) Does perception have a non-conceptual content? *Journal of Philosophy* 98: 239–264.
- Philipona D. L. & O'Regan J. K. (2006) Color naming, unique hues, and hue cancellation predicted from singularities in reflection properties. *Visual Neuroscience* 23: 331–339.
- Philipona D. L. & O'Regan J. K. (2008) Reply to Johnson and Wright. *Visual Neuroscience* 25: 225–226.
- Philipona D., O'Regan J. K. & Nadal J. P. (2003) Is there something out there? Inferring space from sensorimotor dependencies. *Neural Computation* 15: 2029–2049.
► <http://cepa.info/2527>
- Piaget J. (1954) *The construction of reality in the child*. Ballantine, New York. Originally published in French as: Piaget J. (1937) *La construction du réel chez l'enfant*. Delachaux & Niestlé, Neuchâtel.
- Putnam H. (1988) *Representation and reality*. Bradford Books, Cambridge MA.
- Ramsey W. (2015) *Must cognition be representational? Synthese*. Online First.
- Roberts T. (2010) Understanding “sensorimotor understanding.” *Phenomenology and the Cognitive Sciences* 9(1): 101–111.
► <http://cepa.info/2528>
- Searle J. (1990) Is the brain a digital computer? *Proceedings of the American Philosophical Association* 64: 21–37.
- Sedivy S. (2008) Starting afresh disjunctively: Perceptual engagement with the world. In: Haddock A. & Macpherson F. (eds.) *Disjunctivism: Perception, action, knowledge*. Oxford University Press, Oxford: 348–375.

- Sellars W. (1963) *Science, perception and reality*. Routledge & Kegan Paul, London.
- 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: 97–118. ► <http://cepa.info/2529>
- Shoemaker S. (1996) *The first-person perspective and other essays*. Cambridge University Press, New York NY.
- Skinner B. F. (1974) *About behaviorism*. Knopf: New York.
- Speaks J. (2009) Transparency, intentionalism, and the nature of perceptual content. *Philosophy and Phenomenological Research* 79: 539–573.
- Thompson E. (1994) *Colour vision: A study in cognitive science and philosophy of science*. Routledge, London.
- Thompson E. (2001) *Between ourselves: Second-person issues in the study of consciousness*. Imprint Academic, Exeter.
- Thompson E. (2007) *Mind in life: Biology, phenomenology and the sciences of mind*. Harvard University Press, Cambridge MA.
- Tomasello M. (2001) *The cultural origins of human cognition*. Harvard University Press, Cambridge MA.
- Uexküll J. von (1957) A stroll through the worlds of animals and men. In: Schiller C. (ed.) *Instinctive behavior*. International Universities Press, New York: 5–80. German original published in 1934. ► <http://cepa.info/2403>
- Uexküll J. von (1992) A stroll through the worlds of animals and men: A picture book of invisible worlds. *Semiotica* 89(4): 319–391.
- Varela F. J., Thompson E. & Rosch E. (1991) *The embodied mind: Cognitive science and human experience*. MIT Press, Cambridge MA.
- Wittgenstein L. (1953) *Philosophical investigations*. Translated by G. E. M. Anscombe. Blackwell, Oxford.
- Wright R. D. & Ward L. M. (2008) *Orienting of attention*. Oxford University Press, Oxford.