

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.