PSYCHE

Sensorimotor Activity

Mark Rowlands
Philosophy Group
School of Humanities
University of Hertfordshire

M.Rowlands@herts.ac.uk

PSYCHE 12 (1), March 2006

KEYWORDS: action, knowing-how, knowing-that, normativity, representation

ABSTRACT: This paper explores the concept of *sensorimotor activity* that is central to the enactive model of visual perception developed in Alva Noë's book, *Action in Perception*. The appeal to sensorimotor activity is, I shall argue, subject to a dilemma. On one interpretation, such activity presupposes representational states, and therefore is unable to aid us in the project of understanding how an organism is able to represent the world. On the other interpretation, sensorimotor activity fails to accommodate the essential normativity of representational states, and is therefore also unable to aid us in the project of understanding representation. The solution, I argue, lies in a new conception of sensorimotor activity, according to which such activity is normative, but where this normativity is not inherited from prior representational states.

1. Alva Noë's book, *Action in Perception* (MIT Press, 2004) has been eagerly awaited; and does not disappoint. In it we find the combination of empirical investigation, philosophical sophistication, and analytical penetration that we have come to expect of Noë's work—all of which is, as usual, set out in a lucid, cogent, manner. The result is a consistently excellent development of the *enactive* or *sensorimotor* account of visual perception.

Being a vehicle externalist in general, and almost entirely convinced by the enactive approach to visual perception in particular, this paper will rehearse the sort of disagreement it is possible to have only with someone with whom you agree on so very much. Indeed, even to characterize the content of this paper as a 'disagreement' with

Noë, is not really accurate. Instead, I want to identify what I think is a problematic aspect of the enactive approach defended by Noë. This is done not with the aim of arguing that the enactive approach is false; rather, to gesture how the enactive approach needs to be developed in order to overcome this problem. The problematic concept is that of *sensorimotor activity*. This can take two forms: *knowing* and *acting*.

2. In earlier work, in particular in his important collaboration with Kevin O'Regan (e.g. O'Regan and Noë 2001), Noë's work was often thought of as anti-representationalist—hostile to, and hence eschewing the need for, internal representations. In an important clarification, made early in the book, Noë emphasizes that this hostility is more apparent than real:

> No doubt perception depends on what takes place in the brain, and very likely there are internal representations in the brain (e.g. content-bearing internal states). What perception is, however, is not a process in the brain but a kind of skilful activity on the part of the animal as a whole (p.2).

Perception may *involve* the construction of content-bearing internal states, but it does *consist* in the construction of such states. Much of the role traditionally assigned to internal content-bearing states is, in fact, carried by 'skillful activity on the part of the animal as a whole'.

3. This skilful activity consists, roughly, in the ability of an animal to keep track of the systematic connections between what it does and what it experiences. The organism's sensory input is, in complex but systematic ways, dependent on its actions; and having visual experience is a matter of identifying these dependencies. More precisely, the role traditionally assigned to internal representations can, to an extent (a perhaps not exhaustive but certainly significant extent) be played by a combination of:

(1) The ability to act on the world—i.e. to probe and explore environmental structures by way of the visual modality.
(2) Knowledge of the *sensorimotor contingencies* that relate such activity to changes in visual input.

Both (1) and (2) refer to certain kinds of *activity*—acting and knowing—and the role traditionally assigned to representations is, at least in part, taken over by these activities. Thus, when I talk of *sensorimotor activity*, I shall be talking of both acting and knowing in the above senses.

The first condition must be handled with care. It is not that any specific instance of perception requires action on the part of the perceiving animal. What it does require, however, is that the animal possess the relevant sensorimotor knowledge concerning how

its sensory stimulation would alter if it were to perform a given action. And this knowledge cannot be possessed in the absence both of the *ability* to act and the *exercise* of that ability on at least some prior occasions.

4. To see how this works, consider Noë's explanation of *perceptual presence*. If you look at a tomato, for example, you experience it as three-dimensional and round, even though you only see its facing side. And you experience it *as* a tomato, and not as a pair of non-contiguous tomato parts, even though it may be occluded by the pepper pot standing in front of it.

The traditional view addresses the problem of perceptual presence by supposing that we build up an internal model of the tomato: the brain *fills in*—i.e. embellishes, embroiders, and elaborates upon the relatively sparse information presented to it in sensation—and the result is an internal representation of the tomato. Visual perception—as opposed to visual sensation—consists in the construction of this representation. As Noë points out, the phenomenon of *change blindness* provides an important empirical source of pressure for this traditional view of perception (pp.49-55).

Noë's enactive model provides a very different account of perceptual presence:

> Our perceptual sense of the tomato's wholeness—of its volume and backside, and so forth—consists in our implicit understanding (our expectation) that the movements of our body to the left or right, say, will bring further bits of the tomato into view. Our relation to the unseen bits of the tomato is mediated by patterns of sensorimotor contingency. Similar points can be made across the board for occlusion phenomena (63).

Sensorimotor contingencies take two forms: *movement-dependent* and *object-dependent*. Movements of your body can modulate sensory stimulation, and when this occurs, the resulting contingencies are movement-dependent ones. However, movements of the object can also produce sensory changes, and in such cases, the resulting contingencies are object-dependent ones.

5. Noë claims that sensorimotor knowledge—knowledge of the relevant sensoriotor contingencies—is, in either form, a type of 'implicit practical understanding' (p.66); it is 'practical not propositional' (p.117); and 'consists in the possession of practical abilities' (p.117). Noë identifies two reasons for insisting on the practical, non-propositional, nature of sensorimotor knowledge.

Firstly, 'it is unlikely that perceivers (human and otherwise) actually have that knowledge.' (p.118) It is unclear, to say the least, that any perceiver would be able to say exactly what it knows in virtue of which it is able to have experiences of the world. Secondly, 'propositional grasp of counterfactuals … could not be the basis of the grasp of spatial content, because the counterfactuals themselves presuppose a prior grasp of such content.' (p.118). While I agree that the attempt to cast sensorimotor knowledge as

propositional should be resisted, it is not clear to me that either of these reasons is compelling.

6. The claim 'it is unlikely that perceivers (human and otherwise) actually have that knowledge' (p.118) is ambiguous. It could mean that the perceiver must be regarded as in possession of *all* the propositional knowledge necessary to have experience of the world in general. On this reading, the sentence is almost certainly true, but, I think, irrelevant: the requirement that a perceiver be in possession of such knowledge is almost certainly too strong. On the other reading, the perceiver must, *in any particular instance* of perception, be in possession of propositional knowledge that specifies how sensory input will be modulated contingent upon movements of either the perceiver or of the object. On this second interpretation, it is far from clear that the perceiver is not in possession of the relevant propositional knowledge.

To see this, consider the following scenario. You are facing our imagined tomato which is occluded by the pepper pot. An interlocutor then asks you questions such as: 'How would your experience of the tomato change if you were to move your head to *this* point?' 'How would your experience change if I were to move the tomato to *that* point?' And so on. In such cases, you should be able to provide answers that specify, at least vaguely, how your experience will change. That is, you should be able to say things like: 'If I were to move to *this* point, then I will see more of *this* region of the tomato, while, correspondingly, some of *that* region will now be occluded behind the pepper pot.' If so, then you do, in fact, possess propositional knowledge concerning movement-dependent and object-dependent sensorimotor contingencies. This knowledge will be there in dispositional, rather than occurrent form. Moreover, as in the above examples, it is likely to be specified by way of propositions that essentially involve indexicals such as 'here', 'there', 'this', 'that', and these propositions will, accordingly, be *incomplete* (in a roughly Fregean sense). But neither their dispositional nor their incomplete status is sufficient to undermine their status as providers of the content of propositional knowledge.

7. The claim 'propositional grasp of counterfactuals … could not be the basis of the grasp of spatial content, because the counterfactuals themselves presuppose a prior grasp of such content' (p.118) is also ambiguous. And depending on the reading will lead to distinct problems of distinct levels of severity. Noë certainly identifies one legitimate worry: 'The idea here is that it is *because* we experience the tomato as three-dimensional and voluminous that we are committed to the relevant counterfactual conditionals' (p.117). And again, 'This seems to put the cart before the horse; knowledge of the propositions (tacit or otherwise) is (as Peacocke has put it, in conversation) *consequent* on the experience' (pp.119-20).

While this worry is certainly a legitimate one, it is not clear that it is compelling. My intuitions are, on this matter, not as robust as Noë's (or for that matter Peacocke's). In particular, it is not clear which—the experience or the sensorimotor conditionals—come first. Why could it not be that the experience of the tomato, for example, is constituted by the grasp of the conditionals concerning how sensory stimulation will change in the event of certain movements on the part of perceiver or object of perception? Indeed, it seems to

be that is precisely what the enactive model *should* claim. Experience is *literally* constituted by a grasp of these dependencies. After all, it is not as if the enactive approach can allow that the experience is logically prior to a grasp of these dependencies. Then we would have to give some other account of the experience and this, it seems, would be to abandon the enactive model.

Noë is worried that to claim this is 'make the behaviourist error of supposing the effects are logical constructions of their causes' (p.118). But this, I think, is not, in fact, the case: what counts as cause and effect is precisely what is at issue in the enactive approach. The cause of the sensorimotor contingencies associated with the visual state-of-affairs of the tomato being occluded by the pepper pot is the state-of-affairs itself, even though this is irreducibly relative to the perceiver. This state-of-affairs causally gives rise to a set of contingencies that are then available for the perceiver to grasp. And the grasping of these contingencies is in what the experiencing of the presence of the tomato consists. This is not a *confusion* of causes and effects. Rather, it is the *assertion*—and I take it that this is the point of the enactive approach—that experience is thus constituted.

8. Therefore, I do not endorse this form of the worry. Nevertheless, I think there is a genuine concern about putting the cart before the horse. And it is here that the requirement that knowledge of sensorimotor contingencies be non-propositional has real bite.

As we have seen, the basis of the enactive approach is that the need to appeal to visual representations, as a way of explaining visual experience, can be avoided, or at least mitigated, by (1) the ability to act on the world—i.e. to probe and explore environmental structures by way of the visual modality, and (2) knowledge of the *sensorimotor contingencies* that relate such activity to changes in visual input. However, in its propositional form, knowledge is a representational state. Hence any victory over the traditional representational approach would appear to be a Pyrrhic one. Even if we have eliminated the need to appeal to visual representations in a traditional sense, this is only because we have appealed to representations somewhere else in our explanation. We avoid visual representations, but only at the cost of introducing knowledge representations.

9. If this is correct, then it means that the language Noë uses to describe our sensorimotor knowledge is troubling. For it often seems to betray a commitment to a representational conception of this knowledge. Consider, for example, the following:

> Our perceptual sense of the tomato's wholeness … consists in our implicit understanding (our expectation) *that* movements of our body to the left or right, say, will bring further bits of the tomato into view (p.63, emphasis mine)

> My sense of the presence of the whole cat behind the fence consists precisely in my knowledge, my implicit understanding, *that* by a movement of the eye or head or the body, I can bring bits of the cat into view that are now hidden. This is one of the

> central claims of the enactive or sensorimotor approach to perception (pp.63-4, emphasis mine).
>
> In general, our sense of the perceptual presence of the detailed world does not consist in our representation of all the detail in consciousness now. Rather, it consists in our access now to all of the detail, and to our knowledge *that* we have this access (p.63, emphasis mine).

These passages are, I think, entirely representative. The problem, however, is that these passages appeal to knowing, understanding, and expecting *that* by moving my body thus-and-so, I can bring about certain changes in the character of my experience. And knowing, understanding, and expecting *that* are all propositional states—they are individuated by the content of the sentence that follows the *that*-clause. And so, while Noë is rightfully wary of understanding sensorimotor knowledge in propositional terms, he, nonetheless, seems to tacitly employ this understanding in his explanations of sensorimotor knowledge in particular cases.

10. The same problem arises, in a slightly different form, in connection with the ability to *act* on the world—also central to the enactive approach. While different accounts of action can vary significantly, common to all models of action is the idea that both the *status* of an action as an action and the *identity* of an action as the particular action it is depend on it standing in some or other appropriate relation to distinct representational states. Different accounts of action will have very different explanations of what counts as the 'appropriate relation'. But that there is some appropriate relation is asserted by all accounts.

Suppose, for example, you are patting your head while rubbing your stomach. What makes this an action, as opposed to a bodily movement? On a straightforward causal account, for example, what you do counts as an action because it is caused by an appropriate intention, volition or trying. And it counts as one action, rather than two, or many, if this appropriate intention, volition or trying is a single one, rather than two or many (for example, if you are trying to pat you head and rub your stomach at the same time, rather than having two distinct intentions—to pat your head and rub your stomach—which just happen to be contemporaneously activated). Other models of action differ with regard to the connection in which the movement must stand to other representational states for it to count as an action. However, all accounts assert that there must be *some* connection to intentional, hence representational, states.[1]

Therefore, the appeal to action, as a way of avoiding appeal to orthodox visual representations, or reducing the role of such representations, is, again, in danger of amounting to nothing more than a Pyrrhic victory. We may have obviated the need for visual representations, but only by bringing in representational states somewhere else in our explanatory schema—in this case, as individuators of actions.

11. The denial that sensorimotor knowledge is propositional in character, then, is well-motivated. We need to preclude the propositional status of sensorimotor knowledge on pain of merely replacing visual representation with some other form of representation. However, the denial is, at the same time, problematic. This is evidenced in the fact that the clearest illustration of sensorimotor knowledge all seem to be forms of propositional knowledge—knowing, understanding, expecting, and anticipating *that* sensory input will change contingent upon certain actions on the part of the subject of experience or movements on the part of the object of experience. Equally well-motivated, and equally problematic, is the denial that the action involved in perception—the probing, exploring and exploiting of visual structures in the environment—is representational, in the sense of deriving its status as action, and its identity as the particular action it is, by way of its connection to prior representational states. If this conception of action is not precluded, then it seems we will again be merely replacing one form of representation—visual representation—with another.

Thus the desire for a non-propositional conception of sensorimotor knowledge and the desire for a non-representational conception of sensorimotor action have a common root: the need to avoid a merely Pyrrhic victory over orthodox representational conceptions of the mind.

12. However, this is only the beginning of the problem. The danger of a merely Pyrrhic victory over orthodox representational accounts is, in effect, one horn of a dilemma. The other horn is best glimpsed, initially, by way of the appeal to action. Suppose, influenced by the need to avoid appeal to a representational conception of action, we attempt to identify a conception of sensorimotor activity purged of all representational conceptions. Two problems are immediately evident. First of all, it is far from clear that such conception of activity is available. Secondly, it is even more unclear whether such a conception could serve the role required of it by the enactive model. To the extent that such a conception of activity is available, it seems that action will amount to nothing more than bodily movement of some sort. And while bodily movement can, certainly, help with the identification of some sensorimotor contingencies, it is far from clear that it can underwrite the identification of all such contingencies. Passive movement and willed movement do not, in general, facilitate the discovery of the same sensorimotor contingencies (Hurley 1998: pp.384-5). But suppose these problems could be overcome: it is possible, let us suppose, to identify a non-representational conception of action that can play the required role in the identification of all action-dependent sensorimotor contingencies. Then a further, more serious, problem remains.

13. Representation is a *normative* phenomenon. This is one of its essential, and most recalcitrant, features. It a representation with the content *that p* is instantiated, then the world *should*, in an appropriate sense, be *p*.[2] The traditional approach involves trying to capture this normativity by way of a set of relations obtaining between internal configuration and external state-of-affairs which explain why if the former occurs then the latter *should* occur also. A common strategy, for example, is to see the normativity of representation as a function of the normativity inherent in *teleolosemantic* relations. The

enactive approach, as we have seen, does not eschew internal representations, but does see them, in effect, as being designed to function only in combination with acts of environmental probing or exploration, and mediation of this activity through knowledge of the relevant sensorimotor contingencies.

14. Suppose, now, we think of this probing, exploratory activity in non-representational terms. The activity approximates to bodily movement. Then, the danger is that such activity would provide us merely with new ways of causally impinging on the world. But no number of causal impingements can ever add up to normativity. The activity of probing and exploring the world could, in John McDowell's memorable phrase, only ever give us *exculpations* for our visual experience; it could never give us *justifications* for that experience. However, the enactive approach is committed to the idea that the normatively constrained content of an experience exceeds that which can be provided by a perceiver's internal configurations alone. Therefore, if we conceive of sensorimotor activity in purely causal, non-normative, terms such activity could never play any role in explaining the content of visual experience: this content is normative, and sensorimotor activity, conceived of in this way, is not. This worry—that sensorimotor activity will fail to capture the normativity of visual perception—is the second horn of the dilemma.

15. Therefore, with regard to sensorimotor activity, the trick is to identify a conception of activity that satisfies two conditions. Firstly, any token instance of sensorimotor activity must acquire neither its *status* as activity nor its *identity* as the particular activity it is from its connection to prior intentional, hence representational, states. This condition must be satisfied if we are to avoid the suggestion that the enactive approach's victory over traditional representational models, is a merely Pyrrhic one. This distinguishes sensorimotor activity from action in the traditional philosophical sense: for, in the latter sense, actions do acquire both their status and identity from their connection to representational states. Secondly, sensorimotor activity must provide us with more than merely the ability to causally impinge on the world in additional ways: the activity must be more than *merely* causal, it must be normative. This requirement must be satisfied if sensorimotor activity is to play any role in explaining an essentially normatively constrained phenomenon such as the content of visual experience. Satisfying this requirement is sufficient to distinguish sensorimotor activity from mere bodily movement.

16. Action is commonly regarded as a normative phenomenon. However, this normativity is regarded as inherited from that of the intentional states to which the action is constitutively connected: if an action of the type ϕ-ing occurs, then this *should* be preceded by an intentional state of the sort we would characterize as *trying to ϕ*. This model, however, is one we cannot endorse for sensorimotor activity—on pain of merely replacing traditional visual representations with other forms of, equally traditional, representation. The key is to smuggle enough normativity into one's conception of sensorimotor activity to enable such activity to play a role in explaining the normative character of representation, but without regarding this normativity as inherited from prior

representational states. The normativity of sensorimotor activity must be *sui generis*. The question is: how can this be?

17. The answer, I want to suggest, is that sensorimotor activity acquires its normative status from the same source, and for the same reasons, as did internal configurations on the traditional model of representation. On most models, the normative status of internal representations is a matter of them satisfying a *teleological* constraint, or some weighted combination of a teleological and *informational*, constraint. Roughly:

> *Teleological Constraint*. Any representational item, R must have the *proper function* either of tracking the environmental feature that *produces* it, or of enabling an organism or other representational *consumer* to achieve some task in virtue of tracking such a feature.
>
> *Informational Constraint*. Any representational item, R, must carry information about the environment; i.e. it must track some environmental feature.

Sensorimotor activity, I would argue, satisfies both these conditions.[3] Sensorimotor activity both carries information about environmental states-of-affairs and, indeed, has the function of tracking environmental states-of-affairs, or of enabling the perceiving animal to achieve some goal in virtue of tracking such states-of-affairs. Defending this claim would take us too far beyond the brief of this paper. However, the consequences of these claims are important, and can be at least gestured towards here.

Firstly, the normative status of sensorimotor activity, required for it to play a role in explaining the nature of representation, would be safeguarded. This status would not derive from prior, and distinct, normative states. Rather, the normative status of sensorimotor activity would be *sui generis*. This would protect the enactive approach from the charge of scoring a merely Pyrrhic victory over traditional representational approaches.

Secondly, in satisfying teleological and informational constraints, sensorimotor activity has gone a long way towards satisfying the major constraints on representation. This suggests an intriguing possibility: that sensorimotor activity has as much warrant to the status of representational as internal configurations traditionally construed. There are certainly other constraints on representation, besides teleological and informational, but these two constraints are at the core of most accounts of representation. Other constraints, such as misrepresentation and decouplability, are typically thought of as derivative upon these core constraints. So, I submit, as a possibility meriting serious consideration, that sensorimotor activity is itself representational activity—and not because it derives this status from prior representational states. This would be a truly radical interpretation of the enactive approach. Instead of seeing the enactive approach as antithetical to representation, instead representation permeates it to its very core. But the concept of representation, here, has been reinterpreted. Representation is not a matter of an internal configuration of an animal somehow reaching out to its object. The means by which an

animal represents the world—the vehicles of representation—do not stop at the animal's skin. Rather, representation extends out into the animal's behaviour. This behaviour is not, as on traditional accounts, an outer expression of some internal representational core. Rather, the representational status of the behaviour is *sui generis*. The representational activity of the perceiving animal is representational *all the way out*!

References

Hurley, Susan (1998) *Consciousness in Action*, Cambridge, MA., Harvard University Press.

Noë, Alva (2004) *Action in Perception*, Cambridge, MA: MIT Press.

O'Regan, Kevin, and Noë, Alva (2001) 'A sensorimotor approach to vision and visual consciousness', *Behavioral and Brain Sciences* 24, 5, 939-73.

Rowlands, Mark (2006) *Body Language: Representation in Action*, Cambridge, MA: MIT Press. (forthcoming).

Notes

1. Broadly speaking, there are four possibilities extant in the literature: (1) actions are bodily movements caused by intentional states, (2) actions are intentional states that cause bodily movements, (3) actions are a combination of intentional states and bodily movements, (4) actions are intentional states individuated by way of their effects—i.e. successful tryings. The precise nature of each account is, for our purposes, unimportant. What is important is that each asserts that if anything is to count as an action, it must stand in some appropriate relation to an intentional state (on more than one account that relation is identity).

2. What is 'appropriate', of course depends on the nature of the representational state in question. What is appropriate for a visual representation will not be for a memory representation, etc.

3. Indeed, I do argue this elsewhere. See my *Representation in Action* (MIT Press, forthcoming 2006).

4. This paper was supported by the AHRC