

Understanding the ‘active’ in ‘enactive’

Mark Rowlands

Published online: 22 September 2007
© Springer Science + Business Media B.V. 2007

Abstract Much recent work on cognition is characterized by an augmentation of the role of *action* coupled with an attenuation of the role of *representation*. This coupling is no accident. The appeal to action is seen either as a way of explaining representation or explaining it away. This paper argues that the appeal to action as a way of explaining, supplementing, or even supplanting, representation can lead to a serious dilemma. On the one hand, the concept of action to which we appeal cannot, on pain of circularity, be a representational concept. Such an appeal would presuppose representation and therefore can neither explain it nor explain it away. On the other hand, I shall argue, if the concept of action to which we appeal is *not* a representational one, there is every reason for supposing that it will not be the sort of thing that can explain, or supplement, let alone supplant, representation. The resulting dilemma, I shall argue, is not fatal. But avoiding it requires us to embrace a certain thesis about the nature of action, a thesis whose broad outline this paper delineates. Anyone who wishes to employ action as a way of explaining or explaining away representation should, I shall argue, take this conception of action very seriously indeed. I am going to discuss these issues with respect to a influential recent contribution to this debate: the *sensorimotor* or *enactive* model of perception developed by Kevin O’Regan and Alva Noë.

Keywords Action · Representation · Perception · Sensorimotor · Myth of the giving

M. Rowlands (✉)
Department of Philosophy, University of Miami, Coral Gables, FL 33124, USA
e-mail: mrowlands@miami.edu

Action, representation, and the myths of the giving

Much recent work on cognition is characterized by an augmentation of the role of *action* coupled with an attenuation of the role of *representation*. A general theoretical framework for this is provided by what has become known as *vehicle externalism* or *the extended mind*.¹ The general contours of this position can be delineated by way of the following claims:

- (1) The world is an external store of information relevant to cognitive processes such as perceiving, remembering, reasoning, and so on.

It is such a store because, in all essentials, information is *ubiquitous*. One item will carry information about another in virtue of appropriate relations of conditional probability between them. But such relations can be instantiated in the environment just as much as in the relation between an internal representation and its external correlate. In virtue of this, information exists in the environment, and there are certain environmental structures that carry information relevant to cognition.

- (2) Cognitive processes are essentially hybrid—they straddle both internal and external forms of information processing.
- (3) The external processes involve manipulation, exploration, exploitation and transformation of environmental structures that carry information relevant to the accomplishing of the cognitive task at hand.
- (4) At least some of the internal processes involved are ones concerned with supplying the cognizing organism with the ability to appropriately use relevant structures in its environment.

The traditional construal of representations sees them as internal configurations of a subject, individuated by way of their higher-order physical or functional properties. Different versions of vehicle externalism have different conceptions of how much of the function of representations, thus construed, can be usurped by action. At one extreme are eliminativist treatments that see essentially no role left over for traditional representations.² Other versions are less sanguine, and see the role of action as supplementing, rather than supplanting, that of traditional representations. (See for example, Clark 1997; Clark and Toribio 1994; Rowlands 1999.) However, all forms of vehicle externalism agree that at least *some* of the role of representations, traditionally construed, can be taken over by action. Accordingly, an appeal to action is to be found at the core of any such externalism.

This appeal to action, I shall argue, requires serious scrutiny. In particular, I shall argue that the appeal to action as a way of explaining, supplementing or supplanting, representation can lead to a serious dilemma. On the one hand, the concept of action to which we appeal cannot, on pain of circularity, be a representational concept. Such an appeal would presuppose representation and therefore cannot explain, supplement or supplant it. On the other hand, I shall argue, if the concept of action to which we appeal is *not* a representational one, there is every reason for supposing

¹ This view of the mind has, in recent years, been defended by Donald (1991), Hutchins (1995), Clark (1997, 2001), Wilson (1997, 2004), Clark and Chalmers (1998), Hurley (1998), Rowlands (1999, 2003).

² Certain dynamicist accounts fit this profile. See, for example, van Gelder (1995).

that it will not be the sort of thing that can explain, or supplement, let alone supplant, representation.

Another way of putting this is that the appeal to action as a means of supplementing or supplanting representation is liable to fall victim to two distinct versions of a mistake initially identified by Susan Hurley: the *myth of the giving* (Hurley 1998). I shall argue that if the concept of action to which we appeal is representational, then our appeal falls foul of one version of the myth. If the concept is not representational, then it falls victim to the other. The resulting dilemma, I shall argue, is not an unavoidable one. But avoiding it requires us to embrace a certain thesis about the nature of action. In this paper, I can only sketch the general outlines of this thesis. The burden of this paper is to show that anyone who wishes to defend vehicle externalism should take this conception of action very seriously indeed.

I am going to discuss these issues with respect to a recent elaboration of the vehicle externalist position: the *sensorimotor* or *enactive* model of perception developed by Kevin O'Regan and Alva Noë.

The enactive or sensorimotor approach

Traditional approaches to perception have been driven by a supposed gap between *sensation* and *perception*. Suppose you are a blind person holding a bottle. Under normal circumstances you will haptically perceive the presence of a bottle in your hand. However, the incoming sensory information that forms the basis of perception is extremely impoverished. Tradition reconciles these two facts by way of the claim that brain supplements, extends and embellishes the impoverished information contained in sensory stimulation with what are, essentially, various forms of inferential process. The result is the construction of an *internal representation* of the bottle.

Here, however, is an alternative model, advanced by a forerunner of the enactive approach, D.M. Mackay (1962).³ Information is present in the environment over and above that contained in sensory stimulation, and this information is sufficient to specify that you are holding a bottle. In what does this information consist? According to Mackay, in this: your brain is *tuned* to certain *potentialities*. For example, it is tuned to the fact that if you were to slide your hand very slightly, a change would come about in the incoming sensory signals that is typical of the change associated with the smooth, cool surface of glass. Furthermore, your brain is tuned to the fact that if you were to move your hand upwards, the size of what you are encompassing with your hand would diminish (because you are moving to the bottle's neck). And so on.

What does this talk of 'tuning' mean? Basically, your brain has extracted various laws of what O'Regan and Noë (2001, 2002) call *sensorimotor contingency*. Very roughly, your brain has extracted, and has now activated, certain laws pertaining to the way motor action will be accompanied by changes in sensory input; it has, that is, extracted a certain mapping function from motor activity to sensory input. This

³ This example of Mackay's is cited and endorsed by O'Regan and Noë (2001: 945).

provides the additional information lacking in sensory stimulation, information that specifies that you are holding a bottle.

According to the enactive approach, seeing a bottle is an analogous state of affairs. You have the impression of seeing a bottle if your brain has extracted knowledge concerning a certain web of contingencies. For example, you have knowledge of the fact that that if you move your eyes upwards towards the neck of the bottle, the sensory stimulation will change in a way typical of what happens when a narrower region of the bottle comes into foveal vision. You have knowledge of the fact that if you move your eyes downwards, the sensory stimulation will change in a way typical of what happens when the green label is fixated by foveal vision.

O'Regan and Noë have argued, convincingly, that visual perception, just like haptic perception, obeys its own rules of sensorimotor contingency. These contingencies can be divided into two sorts: (i) apparatus-related contingencies, and (ii) object-related contingencies. Contingencies of the former sort are related to the structure of the visual apparatus. Here is a trivial example: in the contingency that the eyes close, the visual stimulation becomes uniform (i.e. blank). Here's a less trivial one (O'Regan and Noë 2001: 941). As the eyes rotate, the sensory stimulation on the retina shifts and distorts in a very particular way, determined by the size of the eye movement. In particular, as the eye moves, contours shift, and the curvature of lines change. For example, if you are looking at the midpoint of a horizontal line, the line will trace out an arc on the inside of your eyeball. If you now fixate upwards, the curvature of the line will change—represented on a flattened out retina, the line would now be curved. In general straight lines on the retina distort dramatically as the eyes move, somewhat like an image in a distorting mirror. Since contingencies deriving from the structure of the visual apparatus characterize all vision, they constitute, O'Regan and Noë, claim, the defining characteristics of visual sensation, and they are what distinguish visual sensation from sensation in other modalities.

Object-related contingencies, on the other hand, derive from the structure of the various objects of perception. We have already encountered some of these. As I slide my eyes up the bottle, from the label to the neck, the change in sensory stimulation typical of what happens when a narrower region of the bottle comes into foveal vision, is a sensorimotor contingency that derives from the structure of the bottle.

Each form of perception has its own contingency rules, and, according to O'Regan and Noë, what differentiates visual perception from other forms is the structure of the rules governing the sensory changes produced by various motor actions. The sensorimotor contingencies within each sensory modality are subject to different invariance properties, and so the structure of the rules that govern the perception in these modalities will be, in each case, different. To learn to perceive visually is to learn the rules of sensorimotor contingency governing the relation between changes in the orientation of the visual apparatus and the resulting changes in the character of the perceived world.

If the enactive approach is correct, there is little need to explain the perception of the bottle in terms of the production or activation of an internal representation. The work of such a representation can be performed by the bottle itself. This is true of both haptic and visual perception. The bottle is an external structure that carries information over and above that present in any sensory stimulation the bottle is

currently inducing in the hand. How does it carry such information? By providing a stable structure that can be probed or explored at will by the haptic or visual modalities.

Thus we arrive back at the general framework for vehicle externalism. Visual perception is essentially hybrid, made up of internal processes (extraction and activation of the laws of sensorimotor contingency) plus external processes (the probing or exploration of information bearing structures in the environment). Visually perceiving is a process whereby the world – understood as an external store of information – is probed or explored by acts of perception, and the results of this exploration are mediated through the laws of sensorimotor contingency.

The myths of the giving

The concept of the *myth of the giving* was introduced by Susan Hurley. In this section, I shall identify two distinct forms this myth can take.⁴ The difference between these forms is, for our purposes, important. Each form of the myth has ramifications for the enactive model of perception, and for vehicle externalism in general, but these ramifications are, in each case, quite distinct.

The myth of the giving (1): Action as an unproblematic given One form of the myth involves taking the content of an action as an unproblematic given, and then using this to *explain* or *explain away* representation. That is, action is used either to explain what representation is, or used to supplement, or even supplant, an appeal to representation in an account of cognition. This is a mistake for the simple reason that action, in its traditional sense, presupposes representation and therefore cannot explain, supplement or supplant it. This form of the myth is perhaps most familiar in connection with a certain response to Wittgenstein's ruminations over the possibility of meaning. Wittgenstein, famously, developed a paradox concerning the possibility of rule-following. Many commentators have thought that the key to solving this paradox lies in Wittgenstein's appeal to *practice*: 'And hence also "obeying a rule" is a practice. And to think one is obeying a rule is not to obey a rule. Hence it is not possible to obey a rule "privately": otherwise thinking one was obeying a rule would be the same thing as obeying.' (Wittgenstein 1953: section 202).

As many have pointed out, however, understood as a *constructive* solution to the rule-following paradox, this appeal is deeply puzzling. The worry is that the appeal to practice presupposes content, and therefore cannot explain it. A practice is *what we do*. However, doing, it seems, is a form of acting and, as such, is essentially

⁴ Hurley does not explicitly distinguish the two forms. Therefore, it is not clear if she would endorse the distinction I am going to draw. However, in my view, that there are at least two distinct forms of the myth is implicit in the distinct uses to which she puts the myth in *Consciousness in Action*. One use, for example, pertains to a certain mistaken interpretation of Wittgenstein's appeal to the concept of practice, and this corresponds to what I am going to call the myth of the giving (1). Another use concerns the parallels between perception and action, in particular the ways these parallels emerge in the context of a neo-Kantian treatment of the unity of consciousness. This use of the myth corresponds to what I am going to call the myth of the giving (2).

connected to intentional states. But intentional states are individuated by their content—which is precisely what the rule-following paradox calls into question.

More precisely, both the *status* of an event as an action, and its *identity* as the particular action it is, depends on its connection to prior intentional states. The precise nature of this connection will depend on which theory of action you endorse, but that there is *some* appropriate connection is asserted by all theories. For example, suppose you are patting your head while rubbing your stomach. Consider, first, what makes this an action. On any traditional philosophical account of action, its status as an action depends on its standing in some appropriate connection to intentional, hence representational, states. The term ‘appropriate’ is defined only *within* a theory. On a causal theory of action, for example, ‘appropriate’ is explained in terms of certain sorts of causal relations – the movement constitutes an action because it is caused by some prior intentional state – an intention, volition, belief–desire complex, etc. Other theories give very different accounts of what an ‘appropriate’ relation is but all assert that bearing *some* relation to other intentional states is essentially involved in being an action.

Secondly, let’s shift focus from the question of the status of an action to that of its identity. How many actions do we have here? Is patting your head while rubbing your stomach one action, or two, or many? Again, on traditional accounts, the individuation of actions is essentially bound up with the individuation of other intentional states. Returning to the causal theory, for example, the idea would be that if my intention or volition is a single state, then the action counts as one rather than many. Thus, if my intention is to pat my head while rubbing my stomach – a single intention – then the action counts as one action. If, on the other hand, I am the subject of two distinct intentions – to pat my head, and rub my stomach – which just happen to be contemporaneously activated, then the action counts as two, rather than one.

In this way, traditional accounts of action make both the *status* of an event as an action, and its *identity* as the particular action it is, essentially dependent on its relation to other intentional states. But intentional states are individuated by their content. And content arises through representation. So, the appeal to action presupposes representation and therefore cannot explain it. And, *a fortiori*, it cannot supplement or supplant it.

The same sort of point applies not just to actions in the usual sense, but *activity* in the broader sense. Any appeal to, for example, mental activity – for example, *knowing* – will fall foul of this first form of the myth of the giving if the activity in question is individuated by way of its content.

The myth of the giving (2): Action as merely causally related to representation In its second form, the myth concerns a certain conception of the boundary between representational states and actions. To see this conception, consider first its brother myth, that of the *given*. In its most familiar form, the myth of the given is the mistake of supposing perceptual experience to be a matter of ‘pure input from the world to mind with no active contribution from the receiver’ (Hurley 1998: 14). If this conception of experience were correct, then the world could play no *normative* role in the formation of one’s experiences. Perceptual experiences are normative in

this sense: if a perceptual experience with the content *that p* occurs, then the world *should be p*. The myth of the given robs experience of this normative dimension: the experience is the result of purely causal impingements on the senses, impingements that do not fall within the *space of reasons*. A perceptual experience is caused by whatever does, *in fact*, cause it not what *should* cause it. Therefore, we cannot explain the normative character of perceptual experience in causal terms alone. The myth of the given, in effect, establishes, or presupposes, a particular type of *boundary* between experience and the world, a boundary that can be straddled only by causal impingements. As a consequence, we may, as McDowell, memorably, puts it, have *exculpations* for our experiences, but we do not have *justifications* (McDowell 1994: 8).

In parallel, the myth of the giving is the mistake of supposing action to be a matter of 'pure output from mind to world' (Hurley 1998: 15). The myth involves an analogous boundary between action and intention, volition, or trying. A naive causal theory of action provides the most obvious example of the myth of the giving. Such a theory identifies action with bodily movement alone. This bodily movement counts as an action in virtue of its causal antecedents such as intentions, volitions, or tryings. There is nothing intrinsic to an action that makes it an action. It is its extrinsic relations that constitute it as such. Causal theories of action, of this simple sort, are not currently popular, and for good reason. The standard objection is that they make actions contingently actions. Since the same bodily movement can occur with or without the prior trying, and since actions are identical with bodily movements, the same action can exist as a non-action. The problem, however, is that actions are necessarily things you do, but mere movements are not things you do at all. So an action cannot be identical with a movement.

It is not my purpose here to evaluate this type of attack. I am concerned with what the attack is generally taken to show—that an adequate account of actions needs to bring the inner trying (or other intentional state) into more intimate relation with the action. Central to the causal theory is a boundary between actions and intentional states that is bridged by causal impingements alone. But we need more than a merely causal relation between trying and action. We need the boundary between trying and action to be more *permeable*. That is, we need to allow the *content* of the trying to flow *into* the action—to become an intrinsic part of the action, to *inform* the action (in a, roughly, Aristotelian sense). And to do this is to reject the myth of the giving.

We need to avoid this myth because actions make normative claims just as much as do experiences. The latter make normative claims in this sense: if an experience with the content *that p* occurs, then the world *should be p*. Similarly, actions make normative claims on tryings: 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* . An action is caused by whatever does, *in fact* cause it, and not necessarily by what *should* cause it. The causal theory of action cannot accommodate this normative dimension to action, and its failure lies in its falling victim to the myth of the giving. It has presupposed a boundary between trying an action that is bridged by causal impingements alone.

The enactive account, I shall now argue, is in danger of falling victim to both of these myths. Worse, the only way it can avoid each myth is by falling victim to the

other. This problem is not fatal, but avoiding it, I shall argue, requires that we embrace a certain, distinctive, view of the nature of action and its relation to representation.

Action, enaction and the first myth of the giving

The underlying idea behind enactive approaches is that seeing fundamentally involves the ability of a perceiving organism to keep track of the systematic connections between what it does and what it experiences. The organism's sensory input is systematically dependent on its actions, and having visual experience is a matter of identifying these dependencies. The role traditionally assigned to internal representations can, to a considerable 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 a certain kind of *activity* – doing and knowing – and the myth of the giving can, accordingly, arise in connection with both claims.

Consider, first, the idea of *knowing* sensorimotor contingencies. O'Regan and Noë's official position is that knowledge of sensorimotor contingencies is a form of 'practical knowledge':

Visual experience is a mode of activity involving practical knowledge about currently possible behaviors and associated sensory consequences. Visual experience rests on know-how, the possession of skills. (O'Regan and Noë 2001: 946)

On this view, then, knowledge of sensorimotor contingencies is a form of *knowing-how*. And, indeed, O'Regan and Noë reinforce the practical character of this knowledge by frequent use of concepts such as *mastery*:

Visual perception can now be understood as the activity of exploring the environment in ways mediated by knowledge of the relevant sensorimotor contingencies. And to be a visual perceiver is, thus, to be capable of exercising mastery of vision-related rules of sensorimotor contingency. (O'Regan and Noë 2001: 943)

To see is to explore one's environment in a way that is mediated by one's mastery of sensorimotor contingencies, *and* to be making use of this mastery in one's planning, reasoning, and speech behavior. (O'Regan and Noë 2001: 944)

The suggestion seems to be that knowing the laws of sensorimotor contingency is like the exercise of a skill one has mastered—as one might have mastered the art of riding a bicycle, or driving a car. Indeed, driving a car is one of the examples they employ to explain their general position.

Consider, as an example, the feeling of driving a Porsche, and its distinctive qualitative character. What does this feeling consist of... There are characteristic ways in which the vehicle accelerates in response to pressure on the gas pedal. There are definite features of the way the car handles turns, how smoothly one can change gears, and so on. What it is like to drive a Porsche is constituted by all these sensorimotor contingencies, and by one's skilful mastery of them. (O'Regan and Noë 2001: 961)

Driving a Porsche, one might suppose takes, as we might say, a *little getting used to*. But when you are used to it, you have mastered the art of driving a Porsche.

So far, then, the idea seems to be that knowledge of sensorimotor contingencies amounts to a type of practical knowledge importantly akin to knowing how to drive a car. However, this interpretation does not seem to square with other claims they make. Many of their explanations of the character of visual experience appeal to a form of knowing *that* rather than knowing *how*. Thus, in explaining the character of our experience of red, they write:

In what does your focusing on the red hue of the wall consist? It consists in the (implicit) knowledge associated with seeing redness: the knowledge *that* if you were to move your eyes, there would be changes in the incoming information that are typical of sampling with the eye; typical of the nonhomogeneous way the retina samples color; knowledge *that* if you were to move your eyes around, there might be changes in the incoming information typical of what happens when the illumination is uneven, and so on. (O'Regan and Noë 2001: 961)

This exercise [of our mastery of sensorimotor contingencies] consists in our *practical understanding that* if we were to move our eyes or bodies or blink, the resulting changes would be those that are typical of red, and not of green patches of light. (O'Regan and Noë 2002: 84. Emphasis is mine.)

In the first passage, there is a shift to talk of knowing *that*, rather than knowing *how*. In the second passage there is a very curious running together of *practical knowledge* with *knowing that*, and this strongly suggests that whatever they mean by practical knowledge, it cannot be knowing *how* in the traditional sense. Moreover, in explaining the character of the haptic perception of a knife, they write:

You know *that* if you move your fingers upwards, you will encounter the ring attached to one end of the knife, and if you move it the other way, you will encounter the smoothness of the plastic surface, and the roughness of the corkscrew. It is this knowledge that constitutes the haptic perception of the object. (O'Regan and Noë 2002: 88. Emphasis is mine.)

And, indeed, knowing *that* fits far more closely with every explanation they give of the exercise of one's mastery of sensorimotor contingencies. Thus, for example, they cite with approval Mackay's account of the haptic and visual perception of a bottle, which seems to clearly involve knowledge *that* rather than knowledge *how*. That is, as you slide your eyes up the bottle, you anticipate *that* the sensory stimulation reaching you will change in a manner consistent with the narrowing of the bottle neck. Indeed, you don't even need to slide your eyes up the bottle, you just

need to know *that if* you were to slide your eyes up the bottle, *then* your incoming sensory stimulation would change in a manner consistent with the narrowing of the bottle's neck, and so on.

The vacillation between knowing how and knowing that is, I think, partially camouflaged by O'Regan and Noë's tendency to shift from talk of mastery of sensorimotor contingencies to knowledge of the *laws* of sensorimotor contingency. For example:

We shall say that the missile guidance system has mastery of the sensorimotor contingencies of airplane tracking if it "knows" the laws that govern what happens when it does all the things it can do when it is tracking the airplanes. (O'Regan and Noë 2002: 81.)

So, knowledge of laws, in this sense, is knowledge of what happens to your visual input when you do certain things. At other points, the laws involved seem quite different. Thus:

As a result of such differences, lawful changes in the neural influx occur as a function of the eyes' position. The laws underlying these changes, that is, the sensorimotor contingencies, are indicative of the fact that the patch is being sampled by the visual apparatus, and not via, say the olfactory or tactile modalities. (O'Regan and Noë 2002: 83.)

Here, the suggestion seems to be that the relevant laws are ones governing neural influx—which seem, on the face of it, to be very different from the laws governing the way one's experience changes contingent upon one's actions. In any event, whatever the form of the laws appealed to, the result is an unfortunate masking of the appeal to knowledge that. One speaks of knowledge *of* laws, and this may give one the impression that knowledge of laws is not a form of knowing that. But, of course; it really is. Knowledge of laws is knowledge of certain facts—very roughly, facts that are expressed in the form of universally quantified modal conditionals or biconditionals that support subjunctive and counterfactual statements. And, in another of their variations, O'Regan and Noë do employ the idea of knowledge of facts when explaining the idea of mastery of sensorimotor contingencies:

Yet so long as the missile guidance system is, for example, tuned to the *fact that* it can turn to bring the airplane back into the camera's sights, we would still say that the missile guidance system is currently visually tracking the airplane. (O'Regan and Noë 2002: 82. Emphasis is mine.)

Knowledge of facts is knowledge *that*. So, while O'Regan and Noë's official position is that knowledge of sensorimotor contingencies is a form of practical knowledge, what this knowledge in fact turns out to be is, I think, a form of knowing *that*. And this matters—deeply.

Knowing *that* is, of course, a representational activity. So, if our theory makes essential use of the concept of knowing that, we can hardly claim that the theory has explained the concept of representation. Nor can we claim that this knowing has supplemented or even supplanted the need for representation. Our appeal to knowledge – if we are indeed dealing with knowledge *that* – presupposes

representation, therefore cannot explain, supplement or supplant it. This is one way, then, in which the enactive model is in danger of falling victim to the first version of the myth of the giving.

Precisely the same danger accompanies the appeal to *action* in the sense of the probing and exploring of environmental structures by way of the visual modality. For, as we have seen, on its usual construal, action, itself, presupposes representation. Both the *status* of an event as an action, and its *identity* as the particular action it is, depends on its connection to prior intentional states. The precise nature of this connection varies from one theory of action to another, but that there is *some* appropriate connection is asserted by all theories. But intentional states are individuated by their content. And content arises through representation. So, the appeal to action presupposes representation and therefore cannot explain, supplement or supplant it.

None of this should, of course, be taken as fatal to the enactive model. The moral we should draw, however, is that if we want to appeal to the idea of activity – broadly construed to incorporate both knowing and acting – in an attempt to explain or explain away representation, then we must be careful not to presuppose an intentional, hence representational, conception of activity.

Action, enaction, and the second myth of the giving

It might be thought that, with respect to the problems described above, the enactive approach has an obvious escape route. The appeal to a representational concept of activity is, if made as an attempt to explain, supplement or supplant representation, obviously circular. Therefore, what the enactive approach requires is a conception of activity that is *not* representational. The problem with this strategy, however, is that it renders the enactive approach susceptible to the *second* form of the myth of the giving.

As we have seen, the second form of the myth concerns a certain conception of the boundary between representational states and actions—a boundary that can be straddled only by causal impingements. Its brother myth – that of the *given* – understands perceptual experience as ‘pure input’, the result of purely causal impingements on the senses. As such, the normative dimension of perceptual experience is lost. Similarly, the myth of the giving – in its second form – understands action as ‘pure output’ from mind to world and, as such, cannot accommodate the normative dimension of action. The root of the myths of the given and giving is the same: both myths situate perception and action behind (or in front of, depending on your perspective) a boundary that only causal impingements can traverse. In the case of perception, this means that the contents of the world can exert no normative pressure on perceptual experience. In the case of action, this means that the contents of tryings, intentions or other representational states can exert no normative pressure on actions.

There is a parallel picture into which vehicle externalism in general, and the enactive approach in particular, is in danger of falling. To appeal, in the attempt to explain, supplement or supplant representation, to a non-representational conception of action is to advance what is, in effect, a *dual component theory* about visual

perception, where the label is intended to draw attention to the corresponding way of understanding good old fashioned *content* – rather than vehicle – externalism.

Content externalism is, minimally, the claim that the content of at least some mental states does not supervene on what is occurring inside the skin of a mental subject. The dual component theory provides one way of interpreting this minimal characterization. The dual component interpretation of content externalism asserts that propositional mental states comprise two components. On the one hand, there is a component that supervenes on neurophysiology, consisting of the qualitative/action-guiding/causal-explanatory properties of the state. On the other, there is a component that does not thus supervene, comprising the state's representational properties. Each component is logically distinct from the other.

This interpretation can legitimately be thought of as *reactionary* because it tries to hold on to the general Cartesian division between mind and world, albeit in a somewhat different form. Specifically, the Cartesian idea of the mind as a self-contained interiority is preserved: it's simply that its bounds are redrawn. The boundaries between inner and outer now correspond to the distinction between, on the one hand, the qualitative/action-guiding/causal-explanatory component of a mental state and, on the other, its representational component. And traditional Cartesian theses about the mind – both ontic and epistemic – could now be asserted about the inner component rather than the entire mental state. The Cartesian view of the mind, thus, persists in recognizable form in the dual component theory.

This type of interpretation can be reiterated at the level of vehicles. Any case of representation can be factored into two components: a genuinely representational one, and an action-based component whose function is to facilitate this genuinely representational component. Applied to the enactive model the idea would be that visual experience is composed of two components. Internal representations of the world exist, but they are not the rich, complex and detailed maps that orthodoxy has taken them to be. Instead, they are rough, partial, and incomplete—providing their possessors with only the general gist of the visually presented world. These proto-representations have been designed to function only in conjunction with environmental probing and exploration by way of the visual mode and knowledge of the relevant sensorimotor contingencies.

The connection between internal and external components may be extremely tight. Typically, the internal component may have been *designed* to function only in tandem with the external component, so that the former cannot fulfill its function in the absence of the latter. However, what makes this a dual component interpretation is the idea that the external component – consisting in appropriate action on the environment – is not itself a genuinely representational component. It may *facilitate* the representational component in performing its representational function, and it may be essential to its performing this function. But the external component is not itself representational.

This dual component interpretation is reactionary in that it acknowledges the force of the arguments in favour of the enactive approach, but tries to limit their significance.⁵ The distinction between action and representation is preserved, but the

⁵ For a wealth of empirical evidence in favour of the enactive approach, see O'Regan and Noë (2001).

scope of the latter is severely attenuated. Representation is logically independent of action, but it does not, at least typically, add up to visual experience in the absence of action.

Most importantly for our purposes, this dual component interpretation reiterates precisely the sort of boundary implicated in the second form of the myth of the giving. How do internal representational vehicles reach out to the world and thus represent it? They are helped by appropriate action—probing, exploring, and exploiting. But this help is, in fact, no help at all—not if our goal is to provide an explanation of representation. If the facilitating action is not representational, then it simply supplies us with the ability to causally impinge on the world in ways in which we would otherwise not have been able. But content, hence representation, is normative and causal impingements are not. No amount of causal impingements can add up to representation. The second form of the myth of the giving stems from our conceiving of the relation between representations and the actions in terms of a boundary that was bridged only by causal impingements. By appealing to action, but understanding this in non-representational terms, we have done nothing to change this boundary. Therefore, our appeal falls foul of the second form of the myth of the giving.

Representation in action

We are now in a position to appreciate the force of the myth of the giving. If we are to avoid the first form of the myth, then, in our attempt to understand representation, we cannot appeal to a concept of action that presupposes representation. If we do, then our explanation will be *circular*. The concept of action we employ must, it seems, be a non-representational one. However, if we are to avoid the second form of the myth, then, in our attempt to understand representation, we cannot appeal to a concept of action that does *not* presuppose representation. If we do, we simply reiterate a conception of the boundary between mind and world as one that is bridged only by causal impingements, and these, as the second form of the myth tells us, can never add up to representation. Therefore, our explanation will be *inadequate*.

This dilemma should not, I think, be regarded as a fatal one for the enactive approach. There is, in fact, a way of avoiding this dilemma and, as far as I'm aware, *only* one way of doing so. This involves looking more closely at the first horn. The appeal to a representational concept of activity will be circular only if that activity acquires its representational status from something else—for example, from a prior intentional state. The case is at its clearest in the case of *action*, so let us focus on this. If, in an attempt to explain or explain away representation, the concept of action to which we appeal is a representational one because it renders the status and identity of any token action dependent on its relation to *prior* representational states, then we shall, of course, be assuming representation to explain representation. Suppose, however, that the representational status of the action is not derived from anything else. Suppose the action is representational because of what it is in itself, and not because of its relation to a prior, or distinct, representational state. Then, the appeal to this type of action would avoid the problem of circularity. Moreover,

understanding how this action acquires its representational status might help us in understanding how representation is possible. And in employing this form of action in understanding representation, we would not be reiterating the sort of boundary between mind and world that is implicated in the second form of the myth of the giving. We would not be committed to thinking of this boundary as something bridged only by causal impingements.

Does the requisite concept of action exist? Are there such things as actions that possess representational status but that do not acquire this status from any other representational state? If there were, they would have to be the sorts of things that can satisfy the generally accepted criteria of representation, including:

- (1) Informational constraint. Any representational item, R, must carry information about the environment; i.e. it must track some environmental feature.
- (2) Teleological constraint. 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.
- (3) Misrepresentation constraint. R must be capable of misrepresenting some feature of the environment. The possibility of misrepresentation goes hand in hand with the possibility of representation: if R is not capable of misrepresenting a given environmental feature then it is not capable of representing that feature either.
- (4) Decouplability constraint. R must be decouplable from the environment in the sense of being able to occur (and guide behaviour) in the absence of the feature whose function it is to track.
- (5) Combinatorial constraint. R must occur not in isolation but only as part of a more general representational framework.

Not all of these constraints are uncontroversial. Worse, on some of their interpretations, they are not even mutually compatible. However, there is general consensus – at least among naturalistically minded approaches to representation – that if an item were to satisfy all five conditions, then it would count as representational if anything does.⁶

I think that there are, in fact, types of action that satisfy these constraints. Moreover, they do so independently of any connection they may bear to other representational states. In virtue of this independence, these types of action are very different from actions understood along traditional philosophical lines, where the connection to prior intentional states is regarded as essential. We might call these events *pre-intentional acts*, or *deeds*.

One example of a type of deed employed in visual perception would be saccadic eye movements. Land and MacLeod (2000) have conducted an important study of

⁶ This is not to say that it would count as *a* representation. To do that, it would also have to play an appropriate role in an agent's psychology; a role typically captured by way of a causal or explanatory constraint on representation. My concern here, however, is with what makes something the sort of thing that could be *about* something else, or take something else as its content. That is, my concern is with what makes something representational and not with what makes it a representation. The former is all that is required to avoid the first horn of the dilemma.

the saccadic eye movements involved in one particular visual task: those employed by a cricket batsman in hitting the ball.

The visuo-motor problem faced by the batsman is to make contact with a ball traveling at anything up to 100 miles per hour, traveling in an arc that sees the ball hit the ground somewhere in front of him, and subsequently bouncing in a manner dependent on the ball's speed, trajectory, orientation, and hardness of the ground. Due to the weight of the bat, and the resulting time it takes to make an adjustment with it, the batsman's judgment of where the ball is going to go must be based on information present in the first 300–400 ms of the ball's flight.

Land and MacLeod measured, by way of a head mounted camera, the eye movements of three batsmen as they faced balls delivered from a bowling machine at a velocity of 25 m/s. The three batsmen were of varying levels of ability. Mark was a professional cricketer, Charlie an accomplished amateur, and Richard an enthusiastic but distinctly unaccomplished amateur. The following results emerged:

The batsman views the ball at crucial moments during its flight.⁷ They fixate on it at the point of delivery—the moment it is released from the bowler's hand. The gaze is stationary for a period after delivery as the ball drops from the field of view. Second, they then saccade to, and fixate on, the anticipated point where the ball will bounce, and gaze is focused on this point for a period of about 200 ms after the bounce. This profile was common to all three batsmen. However, there were also clear individual differences.

Mark, the most accomplished of the three, exhibited more pursuit tracking – smooth, unbroken, tracking of the ball rather than a saccadic jump – than Charlie or Richard. Thus, typically, the saccade accounted for only 48+/-11% of his total pre-bounce gaze change, compared to 69+/-8% for Charlie and 77+/-12% for Richard. Richard, the least skilled batsman, was slower to respond to the appearance of the ball, taking at least 200 ms to initiate the pre-bounce saccade. Thus, the times to the midpoints of his saccades were consistently greater than those of the other players. It seems that Richard was not anticipating the movement of the ball, and was waiting until it completed a large portion of its flight to the bounce point before starting the saccade. This 'catch-up' saccadic behaviour has an unfortunate consequence. Once the ball reaches a certain velocity, Richard's saccade would have been too late to enable him to see the ball bounce—either because he would not have started his saccade at this point, or because the bounce would have occurred at mid-saccade, during which saccadic suppression would briefly suspend vision. By contrast, even with very 'short' balls, (those that bounce soonest after delivery) Mark and Charlie reached the bounce point 100 ms before the ball.

The sorts of saccadic movements employed by the batsman in solving the problem of where and when the ball is going to reach him are things the batsman does rather than things that happen to him. However, they fall short of actions in the traditional sense since they are individuated by way of neither an antecedent intention nor an intention-in-action. In general, we have no idea of the fine-grained saccadic movements we employ in accomplishing visuo-motor tasks. And one cannot have an intention to ϕ if one has no idea what ϕ -ing is or when one is doing

⁷ Received wisdom tells you to 'keep your eye on the ball'. This received wisdom is, in fact, physically impossible to follow.

it. In general, saccadic eye movements occur beneath the level of intention. The identity of saccades as the particular events they are cannot be explained by way of their connection to distinct intentional states, and so saccadic eye movements cannot be individuated by way of such states. This has one crucial consequence: if saccadic eye movements were to possess representational status, this is not a status inherited from other intentional states.

Saccadic eye movements can, I have argued, satisfy the usual constraints required for an item to qualify as representational.⁸ The most important of these is, I think, the teleological constraint. The ball is moving with a velocity and trajectory that will inevitably take it outside the batsman's field of vision. In such circumstances, visual factors such as image expansion and rate of binocular disparity cannot be applied. Therefore, a distinct strategy has to be employed, and this is what results in the general framework of fixation-plus-saccade adopted by all three batsmen. This strategy is something that is acquired through trial-and-error learning. This learning provides the strategy with a *history*. Those who fail to identify the strategy will fail to track the ball. Grossly oversimplifying, you employ the strategy on a given occasion t_3 *because* on earlier occasions t_1 when you did not employ the strategy you failed to track the ball and on other occasions t_2 when you did employ the strategy you did succeed in tracking the ball. The framework – fixate on delivery followed by saccade to bounce point – is employed by an individual on a given occasion because it has worked in the past for that person. And it has worked in the past for that person because it is the only strategy that provides a viable solution to the problem of tracking a moving object traveling in a given type of arc with a given type of velocity. Thus, the employment of the strategy on a given occasion is the result of certain *reproduced* features of the involved framework. And for this reason, the strategy has a *history*. This is sufficient for giving the strategy a proper function in the etiological sense: it has the proper function of tracking the trajectory of the ball (and enabling the batsman to hit the ball in virtue of tracking it). The eye movements employed by batsmen in solving the problem of how to hit a cricket ball, therefore, satisfy the teleological constraint.

Satisfaction of the teleological constraint is, in general, sufficient for satisfaction of the misrepresentation constraint: the teleological account is introduced precisely as a means of securing this constraint. Satisfying the teleological condition gives you, in essence, a way of typing the saccadic scan path in a way that is independent of the actual trajectory of the ball. What is crucial is the trajectory the scan path is *supposed* to track, and not the actual trajectory of the ball. Therefore, when the etiologicaly supposed trajectory and the actual trajectory do not coincide, the scan path misrepresents.

The lack of a necessary individuating relationship between scan path and ball trajectory also entails that the scan path can satisfy the decouplability constraint. The scan path can vary independently of the ball and vice versa. And the scan path – like all complex behaviour – can also be functionally decomposed into constituents or lexical elements in a manner that allows it to satisfy the combinatorial constraint (when this is properly understood). Finally, the scan path carries information about

⁸ For defence of this, see Rowlands (2006).

the trajectory of the ball to no greater extent but also no lesser extent than inner representations traditionally construed, and so satisfies the informational constraint—again, when this properly understood.

I have defended these claims in much more detail elsewhere (Rowlands 2006), and space does not permit me to rehearse the arguments here. However, if the arguments are correct, then two consequences emerge fairly smoothly from it. The first is the possibility of a more radical interpretation of the enactive approach, and of vehicle externalism in general. This interpretation denies the sort of separation of representation and action that was, in effect, common to both eliminativist and dual component interpretations. There is no possibility of separating off the genuinely representational from the action-based components of visual perception, or cognition in general. Actions and representations do not make even notionally separable contributions to the overall task of representing the world. The second, and more general, consequence is a new picture of the nature and role of action – or, more accurately, deeds – in representation. The deeds we employ in representing the world do not merely facilitate our representation; they are themselves parts – *representational parts* – of our representation of the world. Representation – non-derived representation – is not the sole preserve of internal configurations of a subject. Representation extends all the way out into a subject's deeds. Our representation of the world is representational all the way out: it does not stop short of the world.

References

- Clark, A. (1997). *Being-there: Putting brain, body and world back together again*. Cambridge: MIT.
- Clark, A. (2001). *Mindware*. Oxford: Oxford University Press.
- Clark, A., & Chalmers, D. (1998). The extended mind. *Analysis*, 58, 7–19.
- Clark, A., & Toribio, J. (1994). Doing without representing. *Synthese*, 101, 401–431.
- Donald, M. (1991). *Origins of the modern mind*. Cambridge: Harvard University Press.
- Hurley, S. (1998). *Consciousness in action*. Cambridge: Harvard University Press.
- Hutchins, E. (1995). *Cognition in the wild*. Cambridge: MIT.
- Land, M., & MacLeod, P. (2000). From eye movements to actions: How batsman hit the ball. *Nature Neuroscience*, 3, 1340–1345.
- Mackay, D. M. (1962). Theoretical models of space perception. In C. A. Muses (Ed.), *Aspects of the theory of artificial intelligence*. New York: Plenum.
- McDowell, J. (1994). *Mind and world*. Cambridge: Harvard University Press.
- O'Regan, K., & Noë, A. (2001). A sensorimotor account of vision and visual consciousness. *Behavioral and Brain Sciences*, 24, 939–1031.
- O'Regan, K., & Noë, A. (2002). What is it like to see: A sensorimotor theory of perceptual experience. *Synthese*, 129(1), 79–103.
- Rowlands, M. (1999). *The body in mind: Understanding cognitive processes*. Cambridge: Cambridge University Press.
- Rowlands, M. (2003). *Externalism*. London: Acumen.
- Rowlands, M. (2006). *Body language: Representation in action*. Cambridge: MIT.
- van Gelder, T. (1995). What might cognition be if not computation? *Journal of Philosophy*, XCII, 345–381.
- Wilson, R. (1997). *Cartesian psychology and physical minds*. Cambridge: Cambridge University Press.
- Wilson, R. (2004). *Boundaries of the mind*. Cambridge: Cambridge University Press.
- Wittgenstein, L. (1953). *Philosophical investigations (trans. G. E. M. Anscombe)*. Oxford: Basil Blackwell.