Topoi (2009) 28:53–62
DOI 10.1007/s11245-008-9046-z

Enactivism and the Extended Mind

Mark Rowlands

Published online: 8 January 2009

Abstract According to the view that has become known as the *extended mind*, some token mental processes extend into the cognizing organism's environment in that they are composed (partly) of manipulative, exploitative, and transformative operations performed by that subject on suitable environmental structures. *Enactivist* models understand mental processes as (partly) constituted by sensorimotor knowledge and by the organism's ability to act, in appropriate ways, on environmental structures. Given the obvious similarities between the two views, it is both tempting and common to regard them as essentially variations on the same theme. In this paper, I shall argue that the similarities between enactivist and extended models of cognition are relatively superficial, and the divergences are deeper than commonly thought.

Keywords Ability · Enactivism · Extended mind · Sensorimotor knowledge

1 The Extended Mind

According to the view known variously as the *extended mind* (Clark and Chalmers 1998), *vehicle externalism* (Hurley 1998; Rowlands 2003, 2006a, b) *active externalism* (Clark and Chalmers 1998), *locational externalism* (Wilson 2004) and *environmentalism* (Rowlands 1999), at least some token mental processes extend into the cognizing organism's environment in that they are composed, partly (and, on most versions, contingently), of manipulative, exploitative, and transformative operations performed by that subject on suitable environmental structures. More precisely, what I shall refer to as the thesis of the *extended mind* (*EM*) is constituted by the following claims:

- The world is an external store of information relevant to processes such as perceiving, remembering, reasoning … (and possibly) experiencing.
- At least some mental processes are hybrid—they straddle both internal and external operations.
- The external operations take the form of *action*: manipulation, exploitation and transformation of environmental structures—ones that carry information relevant to the accomplishing of a given task.
- At least some of the internal processes are ones concerned with supplying a subject with the ability to appropriately use relevant structures in its environment.

As I shall understand it, therefore, the thesis of the extended mind is (1) an *ontic* thesis, of (2) *partial* and (3) *contingent* (4) *composition* of (5) *some* mental processes.[1]

1. It is *ontic* in the sense that it is a thesis about what (some) mental processes are, as opposed to an *epistemic* thesis about the best way of understanding mental processes. This ontic claim, of course, has an epistemic consequence: it is not possible to understand the nature of at least some of the mental processes without understanding the extent to which that organism is capable of manipulating, exploiting, and transforming relevant structures in its environment (Rowlands 1999). However, this epistemic consequence is not part of *EM*

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

[1] There are other possible ways of understanding *EM*, but this was the status of the thesis I developed and defended in Rowlands (1999).

itself. Indeed, the epistemic claim is compatible with the denial of *EM*.[2]

2. The claim is that (some) token mental processes are, *in part*, made up of the manipulation, exploitation or transformation of environmental structures. There is always an irreducible internal—neural and, sometimes, also wider bodily—contribution to the constitution of any mental process. No version of *EM* will claim that a mental process can be composed entirely of manipulative, exploitative or transformative operations performed on the environment.[3]
3. It is possible to understand *EM* as asserting a necessary truth about the composition of mental processes: that, *necessarily*, some mental processes are partly constituted by processes of environmental manipulation, etc.[4] It is possible, but inadvisable. The underlying rationale for *EM* is provided by a liberal form of functionalism.[5] And the entire thrust of liberal functionalism is to leave open the possibility of different ways of realizing the same (type of) mental process. By understanding *EM* as asserting a necessary truth, therefore, the proponent of *EM* is at risk of undermining his or her own primary motivation.
4. *EM* is a claim about the *composition* or *constitution* of (some) mental processes. Composition is a quite different relation than *dependence*. Thus, *EM* is a stronger and more distinctive claim than one of environmental embedding; and the thesis of the extended mind must be clearly distinguished from that of the *embedded* mind. According to the latter, some mental processes function, and indeed have been designed to function, only in tandem with certain environmental structures; so that in the absence of the latter the former cannot do what they are supposed to do or work in the way they are supposed to work. Thus, some mental processes are dependent, perhaps essentially dependent, for their operation on the wider environment. *EM*, on the other hand, does not simply claim that mental processes are, in this way, *situated* in a wider system of scaffolding, a system that facilitates, perhaps in crucial ways, the operation of these processes. That would be a claim of dependence. Rather, it claims that things we do to this wider system of scaffolding in part compose or constitute (some of) our mental processes.

   Just as *EM* must be distinguished from the thesis of the *embedded* mind, so too must it be distinguished from that of the *embodied* mind. There are different ways of understanding this thesis; some ontic, some epistemic. According to ontic readings, at least some—not all by any means, but some—mental processes are constituted not just by brain processes but by a combination of these and wider bodily structures and processes (Shapiro 2004; Damasio 1994).[6] While this is an interesting and controversial thesis in its own right, it is not equivalent to *EM*. The latter is a claim of partial environmental constitution, not partial bodily constitution.
5. Finally, *EM* does not make a blanket claim about all mental processes. *EM* can view with equanimity the strong likelihood that the composition of some, even many, mental processes is exclusively neural. *EM* claims only that exclusive neural composition is not true of all mental processes. Indeed, the focus, until recently, has largely been on a sub-category of mental processes: cognitive processes. *EM* claims that some, but not all, cognitive processes are ones partly and contingently composed of processes of environmental manipulation, exploitation, and transformation. Thus, contrary to popular belief, *EM* is compatible with the possibility of a brain in a vat. It is just that, if *EM* is true, the mental life exhibited by the brain would be somewhat truncated.

2 The Mind Enacted

Suppose you are a blind person holding a bottle. You have the feeling of holding a bottle. But what tactile sensations do you actually have? Without slight rubbing of the skin, tactile information is considerably reduced, and information about temperature will soon disappear through adaptation of receptors, etc. Nonetheless, despite the poverty of sensory stimulation, you have the feeling of having a bottle in your hand. According to the traditional approach, the brain supplements, augments, and embellishes the impoverished information contained in sensory stimulation with what are, in effect, various inferences or 'guesses' about the sort of thing most likely to be responsible for this stimulation. The result is the construction of an internal *haptic* (i.e., tactile) representation of the bottle.

[2] This is because this epistemic claim is also a corollary of a weaker claim to be discussed shortly: the thesis of the *embedded mind*.

[3] It is truly incredible how often one finds it necessary to repeat this obvious point.

[4] Someone, with an enthusiasm bordering on the rabid, might even find themselves tempted to claim: some mental processes are *necessarily* constituted by processes of environmental manipulation. This *de re* version of the necessity claim would be even less plausible than the modalized *de dicto* claim.

[5] See Clark (2008) for the connection between *EM* and functionalism.

[6] Since *EM* is primarily an ontic thesis, ontic readings of the thesis of the embodied mind, rather than epistemic interpretations, are logically most proximal to *EM*. Hence, for purposes of distinguishing it from *EM*, I shall ignore epistemic readings of the embodied mind thesis.

However, according to Mackay (1967), there is an alternative explanation: 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. More precisely, 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 along the bottle surface, 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 slide your hand upwards far enough, the size of what you are encompassing with your hand would diminish (because you are moving to the bottle's neck). Your sense of holding a bottle is made up of these *anticipations* of how your experience would change if you were to perform certain types of action. In this, Mackay was drawing (explicitly) on an account of phenomenological presence developed by Edmund Husserl (1913).

According to Mackay, again following Husserl, *seeing* a bottle is, at least in one respect, analogous to touching it. 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 bottle's label is fixated by foveal vision, and so on.

Mackay's discussion provides an important early illustration of what has become known as the *enactive* approach to perception, an approach that has received significant recent theoretical development by O'Regan and Noë (2001), Noë (2004) and Thompson (2007). I shall refer to this as the thesis of the *mind enacted* (*ME*). In discussing the relation between this thesis and *EM*, I shall focus on Noë's (2004) account, although I believe similar conclusions can be obtained, *mutatis mutandis*, via the analysis of any major enactivist account.

Suppose you are looking at a cube. You can not of course, see the whole of the cube at any given moment; you see only some of its surfaces. Nonetheless, it appears to you that you are looking at a cube. Noë captures the basic idea of his enactive account in passages such as this:

> As you move with respect to the cube, you learn how its aspect changes as you move—that is, you encounter its visual potential. To encounter its visual potential is thus to encounter its actual shape. When you experience an object as cubical merely on the basis of its aspect, you do so because you bring to bear, in this experience, your sensorimotor knowledge of the relation between changes in cube aspects and movement. To experience the figure as a cube, on the basis of how it looks, is to understand *how* its look changes as you move. (2004, p. 77)

Alternatively, consider your visual experience of a tomato. If you look at a tomato you experience it as three-dimensional and round, even though you only see its facing side. Suppose, further, that your view of the tomato is blocked by the pepper pot that stands in front of it. Nevertheless, you experience it *as* a tomato, and not as a pair of non-contiguous tomato parts. The tomato is *phenomenologically present* to you, despite the apparent limitations of the visual scene. Traditional accounts would explain this in terms of the construction of a visual representation of the tomato—your brain's *guess* concerning what is causing your visual impressions. Noë, however, demurs:

> 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. (2004, p. 63)

Abstracting from the details, the general idea seems clear. Visually perceiving the world is made up of two things:

(1) Expectations about how our experience of an object will change in the event of our moving, or the object of our vision moving, relative to us (or some other object moving with respect to that object—for example, in front of it). Noë calls this *sensorimotor knowledge* or *knowledge of sensorimotor contingencies*. When our expectations are correct, this is because we have *mastered* the relevant sensorimotor contingencies.
(2) The ability to act on the world—i.e., to probe and explore environmental structures by way of the visual modality.

3 *EM* and *ME*: Doppelgangers?

Prima facie, of course, *EM* and *ME* seem to have much in common. To see how much, recall the earlier characterization of *EM*; in particular, the first three conditions. First:

- The world is an external store of information relevant to processes such as perceiving, remembering, reasoning … experiencing.

ME seems to make use of this claim in much the same way as *EM*. The role traditionally assigned to visual representations can be taken over, at least in part, by the fact that the visual world is a stable store of information that can be explored at will by the visual modality. The sense of phenomenological presence implicated in our visual experience of a tomato—our sense that in addition to the aspect it presents to us it has other systematically related aspects—is underwritten by the fact that the tomato is a continuous, structured and stable store of information, one to whose parts or aspects the visual subject is able to direct its attention at will. (In the same way, of course, the bottle is a stable store of haptically obtainable information that can be explored at will by the subject). Secondly:

- At least some mental processes are hybrid—they straddle both internal and external operations.

Again, this also *seems* to be a claim endorsed by *ME*. A representationalist account will explain seeing in terms of the production within the subject of an internal visual representation. Visually perceiving, therefore, begins where sensation—the distribution of light intensity over the retina—ends; and it consists in the internal processes responsible for the production of the visual representation. The enactivist approach, on the other hand, thinks that at least some of the role traditionally assigned to visual representations can be taken over by the probing and exploration of visually accessible structures by way of the visual modality. Clearly what is going on in the brain is going to be crucially important in this process. But, if the enactivist account is correct, it would be a mistake to suppose that it exhausts the process of visually perceiving the world. To the extent that visual representations are involved, they provide us, at most, with the gist of the visual situation; and the details have to be filled in by suitable probing and exploratory action.[7] If so, then *ME* seems to be committed to the hybrid conception of visual perception. Finally:

- The external operations take the form of *action*: manipulation, exploitation and transformation of environmental structures—ones that carry information relevant to the accomplishing of a given task.

Probing and exploration of visual structures in the environment do, of course, seem to be forms of action in this sense. If the visual task in question is, for example, producing (i.e., enacting) experiences that reflect the structure, richness and complexity of the visual environment surrounding the subject, then the enactive account denies that these features need to be reproduced internally—i.e., it denies that they need to be reproduced as features of the visual representation. Rather, in its probing, exploratory activities, the perceiving subject *exploits* the structure, richness and complexity contained in stable external stores of information, and then uses this to enact experiences that reflect this structure richness and complexity. Thus, *ME* also seems to conform to condition three of our characterization of *EM*.

On the surface at least, *ME* seems to follow closely the characterization I have given of *EM*. It is, therefore, initially tempting to think of *ME* as simply a version of *EM*. Indeed, I was once thus tempted and did so characterize *ME* (Rowlands 2002, 2003). I now suspect that this was premature. Not only are *EM* and *ME* different views; it is not even clear that they are compatible views. Thus, it is noticeable that the points of similarity between *EM* and *ME* identified above all turn on the role allotted by *ME* to the probing and exploration of the world by the perceiving subject. However, on closer analysis I think we shall find that this role has been grossly overplayed in this sense: it is far from clear that *ME* assigns any essential role to this sort of activity. In the remainder of the paper, I shall argue that *ME* turns on *expectations* and *abilities* rather than exploratory *activities*. And there is no reason for thinking that either of these are extended in the sense required to make *ME* a version of *EM*. Before we turn to this, however, I want to look at what many think is an important point of divergence between *ME* and *EM* but which I think is a dead end. This concerns the stance each position bears towards *functionalism*.

4 *EM*, *ME*, and Functionalism

It is generally accepted that the arguments for *EM* presuppose functionalism. More than that, they presuppose a peculiarly *liberal* form of functionalism.[8] Indeed, there is a way of understanding functionalism according to which *EM* emerges as a straightforward, almost *trivial*, consequence.[9] In its more liberal forms, functionalism is based on a *principled* indifference to the details of the physical structures that realize mental processes. What is crucial to a mental state or process is its functional role, not its physical realization. For the liberal functionalist: if it walks like a duck, and talks like a duck, then it is a duck. *How* it

[7] This is the principal moral of the *change blindness* results discussed extensively in O'Regan and Noë (2001). The fact that subjects can, under appropriate masking conditions, fail to notice even significant changes in a visual scene suggest strongly, O'Regan and Noë argue, that they have formed no detailed or complex internal representation of this scene.

[8] Shapiro (2004) criticizes the Clark and Chalmers version of the extended mind for this reliance. Also, it underlies many of the criticisms developed by Rupert (2004) against my position.

[9] As far as I am aware, Mike Wheeler was the first person to point this out in a paper delivered at the Extended Mind II conference, University of Hertfordshire, July 2006.

manages to walk and talk like a duck is not directly relevant. To this, *EM* simply adds: neither does it matter *where* it walks and talks like a duck.

However, enactivist approaches are often understood to require the rejection of functionalism. This is because of the central role that such approaches assign to the body. Here, for example is Noë in apparently full-on body centrist mode:

> If perception is in part constituted by our possession and exercise of bodily skills … then it may also depend on our possession of the sorts of bodies that can encompass those skills, for only a creature with such a body could have those skills. To perceive like us it follows that you must have a body like ours. (Noë 2004p. 25)

Functionalism is perceived as being incompatible with this sort of body centrism, and therefore *EM*, which presupposes functionalism, is thought to be incompatible with *ME*, which requires rejection of functionalism. If this accurately sums up the reasons why *EM* and *ME* are thought to be incompatible, however, then those reasons leave a lot to be desired. Far from being incompatible with functionalism, I think the best way of understanding *ME* is as a version of it.

Functionalism is a broad church. The sort of functionalism relevant to *EM* can easily incorporate the sort of body centrist sentiments (or perhaps they are merely suspicions) expressed by Noë. According to this sort of functionalism, (1) cognitive processes are functional roles, (2) the vehicles of cognition are whatever structures realize those roles, and (3) as a matter of contingent fact, some of those structures are external to the skin of the subject of the cognitive processes. The claim that a given functional role consists in, or is constituted by, a skill that can, in fact, only be exercised by something that has a human body is one that is perfectly compatible with the functionalism that underwrites *EM*.[10] In other words we must distinguish two questions: (1) What is the functional role constitutive of a given cognitive process? (2) What sorts of things can realize the functional role constitutive of a given cognitive process? *EM* is the thesis that, as a matter of contingent fact, for some cognitive processes but not for all, the things that realize the functional role of such processes extended outside the skin and skull of the subjects of those processes. This is perfectly compatible with the claim that in the case of some processes, only things that have a human body can exercise the skills that are constitutive of the functional process in question. After all, the skill in question might involve manipulation of an external structure. Therefore, from the point of view of their respective stances towards functionalism, there is no incompatibility between *EM* and *ME*. That this is not, in general, recognized is due to a widespread tendency to confuse functionalism in general, with specific versions of functionalism, or specific theses that are entailed by those versions (such as the thesis of *body neutrality*). But functionalism itself is a much broader doctrine, and should not be confused with these specific incarnations.

However, it is not simply that *ME* is compatible with functionalism, broadly construed. It actually seems to be a specific form that functionalism might take. It should not be forgotten that, when push comes to shove, *ME* is a *dispositionalist* account of visual experience. On the enactive approach, the content of a visual experience is constituted by expectations concerning the way the experience will change in the event of certain contingencies. Therefore, suppose there is a visual experience, E, possessed by subject S, that has a content C, where C consists in the object of the experience O falling under a mode of presentation p_1. Then, we might formulate a (primitive) *Ramsey sentence* along the following lines:

> E is the experience such that if O were to move from coordinates (x_1, y_1) to coordinates (x_2, y_2) while the spatial position of S remains unchanged, then the mode of presentation of O would change from p_1 to p_2; and if O were to move from coordinates (x_1, y_1) to coordinates (x_3, y_3) while the spatial position of S remains unchanged, then the mode of presentation of O would change from p_1 to p_3…

And so on and so forth to infinity in the good old fashioned Ramsey way. In other words, the central claims of the enactivist approach, at least as this is developed by Noë, seem amenable to a Ramsey style analysis—which suggests that they are broadly functionalist claims.

Therefore, if *EM* and *ME* turn out to be distinct doctrines, or even incompatible doctrines, I don't think it is functionalism that lies at the heart of this divergence.

5 *EM* and *ME*: Divergences

Recall the two claims that, I have argued, are constitutive of *ME*. Visually perceiving the world is made up of two things:

(1) Expectations about how our experience of an object will change in the event of our moving or the object of our vision moving (or some other object moving with respect to that object—for example, in front of it). This is *sensorimotor knowledge* or *knowledge of sensorimotor contingencies*.

[10] Whether this is a plausible claim is, of course, another matter; one that cannot be addressed here. Though, for what it is worth, when the claim is properly vehicle-content disambiguated, I suspect that it is not plausible at all.

(2) The ability to act on the world—i.e., to probe and explore environmental structures by way of the visual modality.

Claim (1) concerns sensorimotor *knowledge*: knowledge that consists in a related set of *expectations* about how our experience will change given the obtaining of certain environmental contingencies. Claim (2) concerns our ability to act on the world. Adjudicating the claim that *ME* yields an extended account of perception, then, amounts to answering this question: is there any reason for supposing that either our expectations or our ability to act (or both) are extended? Is there any reason for thinking that our expectations about how our experience will change in the event of certain contingencies are extended in the way that *EM* claims (some) mental processes are extended? Is our ability to probe and explore environmental structures extended into, or distributed onto, the world? If the answer to both of these questions is 'no', then we would have to reject the idea that *ME* yields an extended account of perceptual processes. I shall argue that the answer to the first question is a definite 'no'. And the answer to the second question is, in all likelihood, also 'no'. Therefore, appearances notwithstanding, *ME* probably does not yield an extended account of perception.

5.1 Claim (1): Sensorimotor Knowledge

There seems to be little reason why expectations about how our experience of an object will change in the event of our moving or the object of our vision moving should be extended. The idea that these sorts of expectations constitute our experience is one that originates in the phenomenological tradition; and they were certainly not introduced there as examples of extended mental processes. So, there is certainly no reason why these expectations *must* be extended ones. But is there any reason for supposing that they *might* be?

Noë claims that these expectations are a form of practical knowledge or *knowing how*. But, again, there is little reason for thinking that this sort of knowing how is extended. It is common, for example, to think of practical knowledge in *procedural* terms: that is, in terms of a list of instructions the following of which will, in theory, allow one to accomplish a given task. But there is no reason for thinking that these sorts of instructions are extended. Typically, that is not the way they have been understood.

If, however, we have been influenced by the Heidegger-Dreyfus-Wheeler axis, we might want to deny that sensorimotor knowledge can be reduced to procedural knowledge.[11] Our manner of relating to the world, including in this case the way in which we relate perceptually to the world, is ultimately non-propositional: propositional modes of relating to the world are always derivative on a more basic way of being-in-the-world. I have a considerable amount of sympathy for this view. However, the Heideggerian gambit is a risky one in this context because it threatens to reduce *EM* to a truism. If this is what sensorimotor knowledge is, then of course it is extended. It is so for the simple reason that *everything* is extended.

For Heidegger (1926/1996), Dasein—the being of humans—is essentially *being-in-the-world*. By this, he didn't mean that first there are humans and, in addition, there is this property of being-in-the-world that all humans possess essentially. His claim was that humans *are* being-in-the-world. That is, each of us is, in essence, a network of related practices. Each of these practices presupposes an instrumental network of related items. We might find ourselves tempted to describe this by saying that human practices are embedded in a wider system of instruments. However, this would be crucially misleading. To describe the relation as *embedding* presupposes that there is a distinction between the practices and the instrumental network that embeds them. And this is precisely what Heidegger wished to deny. The instruments are partly constitutive of the practices. *Being and Time* is the attempt to understand humans simply as a system of practices in this sense. So, each one of us incorporates both the practices and the instrumental network that is constitutive of them. But if this is the underlying vision then *everything* we do (and, indeed, are) is, in fact extended. We must eschew thinking of a human being as a biological entity with biological boundaries of the usual sort. The being of humans is simply practices, practices that take place in the instrumental networks that partly realize them. Any expectations we might possess concerning the likely trajectory of our experiences are derivative upon this more basic way of being in the world. Given this Heideggerian vision, there cannot be any special issue of whether the mental things we do are extended. The claim that the expectations constitutive of our sensorimotor knowledge are extended, therefore, would emerge as trivially true. But so too would *EM*. So, we might be able to reconcile *ME and EM* by the injection of some serious Heideggerian metaphysics but only at the cast of rendering *EM* trivial. Short of this, however, there seems to be little reason to think that the expectations constitutive of our sensorimotor knowledge are extended.

There is a further problem. While Noë's official position is that sensorimotor knowledge is a form of knowing how, all the actual examples he gives of this knowledge seems to be forms of knowing that. Recall the passage cited earlier:

> Our perceptual sense of the tomato's wholeness—of its volume and backside, and so forth—consists in

[11] See Martin Heidegger (1926); Hubert Dreyfus (1992); Michael Wheeler (2005).

> 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. (2004, p. 63; emphasis mine)

This is knowledge *that* rather then knowledge *how*. Or take the other passage cited earlier:

> When you experience an object as cubical merely on the basis of its aspect, you do so because you bring to bear, in this experience, your sensorimotor knowledge of the relation between changes in cube aspects and movement. To experience the figure as a cube, on the basis of how it looks, is to understand *how* its look changes as you move. (2004, p. 77)

Here, Noë does at least talk of understanding 'how' the look of something changes as you move. But this is such an anodyne sense of understanding how that it seems interchangeable with understanding that. After all, what is it to understand *how* the look of something changes as you move? This seems to amount to nothing more than understanding *that* if you were to move thus, then the look of the object would change in such and such a way. In other words, the grammar of Noë's claim is, here, misleading: while he appears to be talking about understanding how, he is really talking about understanding *that* (Rowlands 2006b, 2007).[12]

So, if sensorimotor knowledge were to be regarded as extended, we would have to make out the case that at least some tokens of declarative knowledge or understanding are extended. And not just any declarative knowledge: we would have to show this with regard to the declarative knowledge implicated in perception. The difficulties with this are, I think, formidable. Therefore, if there is a stronger connection between *ME* and *EM*, it will have to be found in the second constitutive feature of *ME*: the ability to act on the world—probe and explore its structures by way of the visual modality.

5.2 Claim (2): The Ability to Act on the World

In assessing this claim, we need to draw a familiar distinction between *ability* and the *exercise* of that ability. There are two ways of understanding claim (2): one much stronger than the other. According to the weak version, visually perceiving the world only requires the *ability* to probe and explore the world by way of the visual modality. It does not require the actual *exercise* of that ability—it does not require the actual probing and exploring of the world. On the stronger version of the claim, visually perceiving the world requires not only the ability to probe and explore the world by way of the visual modality; it also requires exercise of that ability.

Consider, first, the weaker claim. Is there any reason for thinking that abilities to probe and explore the world are extended? Given the distinction between ability and the exercise of that ability, there does not seem to be. My playing of the piano is a spatially and temporally extended process that centrally involves, as one of its constituents, the keys of the piano itself. But I can have the ability to play the piano even if I never come across another piano in my life and so never have the chance to exercise that ability. The distinction between the possession and the exercise of ability can be applied to abilities of all kinds—human or not. The fertilization of an egg by a sperm is a process that incorporates, as constituents, both sperm and egg. But, the sperm has this ability even if, due to the vicissitudes of fortune, it never finds itself in the right place at the right time. The obvious moral seems to be that while the exercise of ability might be an extended process, the same does not hold for the ability itself. Abilities are not extended in the sense required by *EM*.

It is true, of course, that some abilities might be *embodied*.[13] Here is John Haugeland discussing the ability to type:

> [T]hat some particular pulse pattern [in my brain], on some occasion, should result in my typing an 'A' depends on many contingencies, over and above just which pattern of pulses it happens to be. In the first place, it depends on the lengths of my fingers, the strengths and quicknesses of my muscles, the shape of my joints, and the like. Of course, whatever else I might do with my hands, from typing the rest of the alphabet to tying my shoes, would likewise depend *simultaneously* on particular pulse patterns and these other concrete contingencies. But there need be no way to 'factor out' the respective contributions of these different dependencies, such that contents could consistently be assigned to pulse patterns independent of which fingers they're destined for. (1995, p. 253)

I think one should agree with Haugeland on this point. Many abilities are embodied in the sense that whether or not you have them is a matter not just of what is going on in your brain but also dispositions built into your body

[12] I am not, here, rehearsing the Stanley and Williamson (2001) claim that there is no distinction between knowing how and knowing that. On the contrary, I think Stanley and Williamson are clearly mistaken. There is a legitimate distinction, but Noë fails to draw it. In particular, on his account, the expectations constitutive of sensorimotor knowledge are expectations *that*.

[13] Not all of them of course. My ability to mentally picture and count the number of windows in my house when I am sitting miles away in my office is an ability that is not composed of wider bodily structures and processes. The possession of this ability seems to depend purely on what is going on in my brain.

whether through training or biological endowment. My ability to surf is not simply a matter of my brain encoding the relevant form of practical knowledge but also of my body having acquired, through a long process of training, the necessary bodily dispositions or tendencies. Without these dispositions, what is going on in the brain would not add up to the ability to surf. While not all abilities are embodied, it seems undeniable that some of them are. However, as we have seen, *EM* is a quite distinct thesis from the claim that mental processes are embodied. *ME*'s appeal to abilities to probe and explore the world by way of the visual modality might point us in the direction of an *embodied* view of perception. But one cannot move from this to an *extended* account without much further argument.

The same sorts of considerations also point to the conclusion that some abilities are environmentally *embedded*. The bodily dispositions I have acquired in the course of learning to surf themselves have to be tailored to specific environmental contingencies. For example, the ability to surf on a 7′11″ mini-mal does not translate into the ability to surf on a 5′11″ thruster. However, as we saw earlier in this paper, *EM* is distinct from, and considerably stronger than, the claims that mental processes are embodied and embedded. The claim that mental processes have environmental constituents is a much more striking claim than merely that they have *bodily* constituents, or that they are *dependent* on the wider environment. One can accept that many abilities—though by no means all—seem to be complex constructions out of brain activity, acquired or innate bodily dispositions, and environmental feedback. This still does not give you an extended account of abilities.

According to the stronger interpretation of claim (2), visually perceiving the world requires not only the ability to probe and explore the world by way of the visual modality; it also requires exercise of that ability. It goes without saying, of course, that the exercise of many abilities consists in processes that are extended into the world and include items in the world among their constituents. So, the stronger interpretation of (2) might certainly entail *EM*. The problem, however, is that this stronger interpretation seems grossly implausible.

The immediate problem, of course, lies in accounting for *novel* visual phenomena. Suppose you encounter—to return to Noë's example—a tomato that you have never seen before. According to condition (1) of *ME*, perceiving the shape of the tomato consists in grasping the relevant sensorimotor contingencies. That is, it involves understanding how your visual experience will change contingent on your moving relative to the tomato, or the tomato moving relative to you, or an object occluding the tomato, etc. But suppose we now add on the stronger version of condition (2): perceiving the shape of the tomato involves the actual exercise of the ability to probe and explore the world by way of the visual modality. But this entails that prior to exercising the ability one does not see the shape of the tomato.

The obvious response for the defender of *ME* to make is to appeal to prior experience. You do not need to actually exercise the ability to probe and explore the environment because while you might not have seen this particular tomato before, you have seen tomatoes of a similar shape. Therefore, on the basis of this prior experience, you can anticipate how your experience would change contingent upon certain events, such as your moving relative to the tomato.

This response, however, faces two problems: the first intrinsic to it, the second pertaining to the possibility of regarding *ME* as yielding an extended account of perception. The first problem concerns the possibility of perceiving novel visual shapes. For any object with a shape that you have hitherto not encountered, the stronger version of (2) entails that you do not actually perceive that shape until you have acted on it—visually probed and explored it—and witnessed how your experience changed as a result. Failing this, you will fail to perceive the novel shape. The same, according to this strong interpretation of the enactive account, is true of any novel visual property of an object.

The worry here, of course, is that *ME* is confusing perception with subsequent cognitive operations. In essence, the worry is that *ME* runs together the distinction between *perception* and *judgment*. It certainly seems that something *in the vicinity* of seeing must be going on prior to the probing and exploratory activity. There is no probing and exploratory activity *simpliciter*. That is, probing and exploratory activity is not something one does willy-nilly.[14] On the contrary, the activity is guided by some or other visually salient feature of the situation. So, when we explore the visual potential of a novel shape, for example, what is it that guides our exploration? The obvious response is that what guides our exploration is our perception of the shape. We certainly see something, and the most natural candidate for what we see is the shape. We may not know exactly what shape it is; that is what the subsequent exploration is to tell us. But this latter issue is a matter of judgment not perception.

The second problem is more germane to our concerns. Noë does seem to endorse the stronger interpretation of claim (2). In a passage cited earlier, for example, he talks of perception being 'constituted by our possession *and* exercise of bodily skills' (2004, p. 25).[15] However, sometimes,

[14] Ironically, that would make the probing and exploratory activities involved in visual perception the equivalent of 'closing one's eyes and swinging'.

[15] It is true that he puts this in interrogative form. But it is clear from context that this is a claim he wishes to endorse.

his claims seem to suggest that the actual exercise of a sensorimotor ability is required only during the process of learning to perceive a visual property. Thus: 'only through self-movement can one test and so learn the relevant patterns of sensorimotor dependence.' (2004, p. 13).

Even when restricted to learning to perceive, the condition of self-movement does seem to be very strong. Following Schellenberg (2007), imagine a *sentient statue*. The statue cannot initiate its movements, but is moved around by some external agency. Why could the statue not learn to identify the sensorimotor contingencies associated with various objects? Why could a statue thus moved not learn to perceive the shape of a tomato or cube, for example? *ME* can respond by pointing out that sensorimotor probing can be a lot more subtle than simply gross movements of the body. Saccadic eye movements, for example, are one way in which sensorimotor contingencies can be identified (Rowlands 2006a). However, this would simply push back the problem a stage. We now imagine a sentient statue whose saccadic eye movement are induced by some external agency.

However, the issue of self-movement is only tangential to our concerns. If the exercise of one's ability to probe and explore the environment is only required for *learning* how to perceive a visual property, whereas simply the ability will suffice for actually perceiving a property one has previously encountered, then this means that only learning how to perceive a visual property will be an extended process. Perceiving an already encountered property will require only the relevant expectations concerning how one's experience will change given certain contingencies, and the ability to probe and explore the relevant portion of the environment. And if the arguments developed here are correct, there is no reason for thinking that either of these is an extended process.

We can represent the situation in the form of a dilemma. If *ME* claims that the actual exercise of one's ability to visually probe the world is required for perception, then *ME* is implausible. If, on the other hand, it claims that exercise of this ability is required only during the learning phase, then it yields only an extremely attenuated version of *EM*: an extended account of learning to perceive, but not of perception itself. As far as its account of the latter goes, *ME* supplies us with a solidly internalist account oriented around the possession of expectations and abilities. This is, of course, not necessarily a bad thing. Many would regard this anodyne internalist interpretation of *ME* as counting in its favor rather than as a strike against it. However, if one would like *ME* to live up to the sort of radical billing often associated with it, one will have to provide further argument.

One possibility would be to try and undermine the distinction between perceiving and learning to perceive; and it has to do this in such a way that perception turns out to be a lot more like learning than learning is like perception. Interestingly, an attempt to undermine the distinction between learning and perception can be found in both Hurley and Noë (2003) and Hurley (forthcoming). Here, the learning/perception distinction is represented as the training/post-training distinction. The attempt is based on trying to shift attention away from what Hurley calls the 'sufficiency question' to what she calls the 'explanatory question'. With regard to perceptual experience, the sufficiency question would be: 'What in the system suffices for a visual experience, P, with a given content?' The corresponding explanatory question would be: 'Why is this neural state the neural correlate of the visual experience P?' Thus, Hurley proposes that we switch focus from the issue of the *most local mechanism* of perceptual experience to the issue of what provides the *best explanation* of the quality and character of the experience. While the local mechanism that suffices for a perceptual experience may be internal to the perceiving subject, Hurley argues, the best explanation of the quality and character of the experience will have to advert to 'a characteristic extended dynamic'. That is, the best explanation of the quality and character of the perceptual experience will advert to a distributed process incorporating brain, body and the active probing or exploration of the world.

Unfortunately, however, this attempt does not work. To see why, consider a distinction I drew in previous work (2003): the distinction between *possession* of a property and the *location* of things that possess that property. Consider for example, the property of being a planet. Possession of this property by an object requires that it stand in a certain relation to things outside it—a sun which it orbits, for example. It is standing in this relation that makes something a planet; and an explanation of why something is a planet would, therefore, have to refer to these things. But it does not follow from this that a planet is located wherever its central sun is located. Issues of property *possession*, and explanation of that property possession, do not translate into issues concerning the *location* of the token items that possess this property.[16] Thus, we might agree with Hurley that the best explanation of the quality and character of an experience might appeal to a characteristic extended dynamic. And we might agree that this is the best explanation because this dynamic is indeed responsible for the possession by the experience of this quality or character. However, it does not follow from this that the experience is extended.[17] In other words, while we

[16] This point originally goes back to Davidson (1987). For something to be sunburn, it must stand in a certain relation to solar radiation. But it does not follow that the sunburn must 'extend' into the solar radiation.

[17] A similar claim is endorsed by Clark (submitted).

might justifiably conclude that the experience is environmentally *embedded*, it does not follow from this that it is environmentally *extended*.[18]

References

Clark A (2008) Pressing the flesh: a tension in the study of the embodied, embedded mind. Phil Phenomol Res 76(1):37–59

Clark A (submitted) Spreading the joy: why the machinery of consciousness is (probably) still in the head

Clark A, Chalmers D (1998) The extended mind. Analysis 58:7–19

Damasio A (1994) Descartes' error. Grosset/Putnam, New York

Davidson D (1987) Knowing one's own mind. Proc Am Philos Assoc 60:441–458

Dreyfus H (1992) Being-in-the-world: an interpretation of Heidegger's being and time, division 1. MIT Press, Cambridge, MA

Haugeland J (1995) Mind embodied and embedded. In: Haaparanta L, Heinämaa S (eds) Mind and cognition: philosophical perspectives on cognitive science and artificial intelligence. Acta Philos Fennica 58:233–267

Heidegger M (1926/1996) Being and time (trans. J. Stambaugh). SUNY press, New York

Hurley S (1998) Consciousness in action. Harvard University Press, Cambridge, MA

Hurley S (forthcoming) Varieties of externalism. In: Menary R (ed) The extended mind. MIT Press, Cambridge, MA

Hurley S, Noë A (2003) Neural plasticity and consciousness. Biol Philos 18:131–168

Husserl E (1913) Ideas pertaining to a pure phenomenology and phenomenological philosophy, vol 1. Martinus Nijhoff Publishers, The Hague

Mackay D (1967) Ways of looking at perception. In: Watthen-Dunn W (ed) Models for the perception of speech and visual form. MIT Press, Cambridge, MA

Noë A (2004) Action in perception. MIT Press, Cambridge, MA

O'Regan K, Noë A (2001) A sensorimotor account of vision and visual consciousness. Behav Brain Sci 23:939–973

Rowlands M (1999) The body in mind: understanding cognitive processes. Cambridge University Press, Cambridge

Rowlands M (2002) Two dogmas of consciousness. In: Noë A (ed) Is the visual world a grand illusion? Special edition of J Conscious Stud 9(5–6):158–180

Rowlands M (2003) Externalism: putting mind and world back together again. Acumen, London

Rowlands M (2006a) Body language: representation in action. MIT Press, Cambridge, MA

Rowlands M (2006b) Sensorimotor activity. Psyche 12(1). Special edition on Alva Noë's Action in perception

Rowlands M (2007) Understanding the active in enactive. Phenomenol Cogn Sci 6:427–443

Rupert R (2004) Challenges to the hypothesis of extended cognition. J Phil 101(8):389–428

Schellenberg S (2007) Action and self-location in perception. Mind 116(463):603–631

Shapiro L (2004) The mind incarnate. MIT Press, Cambridge, MA

Stanley J, Williamson T (2001) Knowing how. J Philos 98:411–444

Thompson E (2007) Mind in life: biology, phenomenology and the sciences of the mind. Harvard University Press, Cambridge, MA

Wheeler M (2005) Reconstructing the cognitive world: the next step. MIT Press, Cambridge, MA

Wilson R (2004) Boundaries of the mind: the individual in the fragile sciences. Cambridge University Press, New York

[18] Indeed, it is not clear that the claim of environmental embedding is justifiable. As the example of the planet makes clear, one can be an internalist about experience and accept with equanimity the claim that the possession of a given property by an experience depends on a 'characteristic extended dynamic.' Further adjudication of this point requires detailed working out of the concept of embedding, and that is beyond the scope of this paper.