Are minimal representations still representations?

Shaun Gallagher
Philosophy and Cognitive Science
University of Central Florida
gallaghr@mail.ucf.edu

Abstract
I examine the following question: Do actions require representations? Recent work by Mark Rowlands, Michael Wheeler, and Andy Clark suggests that actions may require a minimal form of representation. I argue that the various concepts of minimal representation on offer do not apply to action per se and that a non-representationalist account that focuses on dynamic systems of self-organizing continuous reciprocal causation at the subpersonal level is superior. I further recommend a scientific pragmatism regarding the concept of representation.

Representation in action

Is the concept of representation required for an account of action? In this paper I want to take the anti-representationalist argument as far as I can. My focus is on action, however, and I don't want to make any claims about the more general question of representational accounts of cognition. I also want to set aside the question of whether representation plays a role in deliberation about action, or the planning of action, or the working out of prior intentions, and so forth. The question I address is whether representation is necessary *in* action, as part of action itself.

Consider the classical concept of representation. The following list of characteristics is based on (but not identical to) one provided by Mark Rowlands (2006, 5ff), who understands the classical concept to be modeled on language, on how words work.

1. Representation is internal (image, symbol, neural configuration)
2. Representation has duration (it's a discrete identifiable thing)
3. Representation bears content that is external to itself (it refers to or is about something other than itself)
4. Representation requires interpretation -- it's meaning derives from a certain processing that takes place in the subject -- like a word or an image its meaning gets fixed in context
5. Representation is passive (it is produced, enacted, called forth by some particular situation; or we do something with it)
6. Representation is decoupleable from its current context.[1]

[1] Rowlands does not include decoupleability in as part of the classic conception, but others certainly do. As he notes: "It is often thought that for an item to be regarded as genuinely representational it must be *decouplable* from its wider environment and, in particular, from the state of affairs that it purports to

The idea of decoupleability, for example, is that one can go "off-line" and *represent* (imagine or remember) an action or context. Representation involves a form of decoupling away from action, away from the target of action, or from the current context. But this kind of offline activity may still involve aspects of action and context (I'll come back to this point). In contrast, the question about representation *in* action is whether action itself depends on representation, and in this regard, whether there can be a decoupled element within action itself.

Taking up an anti-representationalist view, Dreyfus (2002) argues that for practiced or skillful intentional action one does not require representation.

> A phenomenology of skill acquisition confirms that, as one acquires expertise, the acquired know-how is experienced as finer and finer discriminations of situations paired with the appropriate response to each. Maximal grip [Merleau-Ponty] names the body's tendency to refine its responses so as to bring the current situation closer to an optimal gestalt. Thus, successful learning and action do not require propositional mental representations. They do not require semantically interpretable brain representations either. (Dreyfus 2002, 367).

Dreyfus associates the idea of representation with a failed Cartesian philosophy -- the concept of representation (as used in AI, for example) remains context-independent and bound up with epistemic states of *knowing-that* (propositional knowledge), when everything about intelligent action and *knowing-how* depends on being-in-the-world – rather than standing back and representing the world – and on context – both background and immediate context. The limitations of representationalism can be seen in the commonsense knowledge problem and the frame problem in AI.

Representational approaches to the problem of commonsense knowledge (pictured as an interconnected system of representations) leads to "a vicious combinatorial explosion" (Wheeler 2005), as propositional (representational) knowledge of one aspect of the world presupposes propositional knowledge of other aspects, etc. etc. The frame problem – roughly, how does a system adjust itself to recognize relevant features in a changing environment? – remains unsolved in a representationalist model. To claim that the system uses appropriate representations is just to push the problem back and leads to an infinite regress -- how does the system know what representations are appropriate/relevant to the particular context?

It's important to note that what is involved here is not simply the immediate context of a well-ordered task. One might think that a connectionist model, where a system is trained in a rich environment, could generate distributed representations that are context-sensitive. But the problem is also the *background* context that informs the particular system's abilities and decisions. Whereas a background context is not required for a robot that does one pre-defined task (my household robot does a great job at vacuuming, but it can't do the dishes), the human context is not defined by narrowly circumscribed actualities; it is also defined by finite but extensive possibilities.

---

represent" (2006, 157). Rowlands, however, does include this characteristic in his more enactive definition of representation.

A similar anti-representationalist stance is explicated by Berthoz and Petit (2006). They argue that the brain is an organ for action rather than an organ for representation. They want to move away from a philosophy that puts language in first place and that models action on language-like representation, where action is equivalent to movement plus representation.

> By applying this representational filter, everything in the external as well as in the internal world appears frozen, fixed and stabilized by the projection of the propositional form, which implicitly structures representation. (Berthoz and Petit 2006, 23).

Representations, on the propositional model, are discrete static structures, states of affairs that lack the kind of dynamics found in action. For Berthoz and Petit, action involves anticipation. As we'll see, however, on one model, anticipation itself might be considered representational.

What takes the place of representations in non-representationalist accounts of action is a form of perceptually based online intelligence which generates action "through complex causal interactions in an extended-body-environment system" (Wheeler 2005, 193). But can this sort of system do everything it needs to do without representation?

Minimal representations

Michael Wheeler (2005) is a friend of Dreyfus' anti-representationalist view, but following Andy Clark (1997, 47ff., 149ff.), and inspired by Brooksian robotics, he suggests that certain actions require "action-oriented representations" (AORs). AORs are temporary egocentric motor maps of the environment that are fully determined by the situation-specific action required of the agent (organism or robot). On this model, it is not that the AORs re-present the pre-existing world in an internal image, or that they map it out in a neuronal pattern: rather, "how the world is *is itself encoded in terms of* possibilities for action" (197). What is represented in AORs is not knowledge that the environment is x, but knowledge of how to negotiate the environment. AORs are action specific, egocentric relative to the agent, and context dependent.

But what sort of thing can count as an AOR? Is it a neural firing pattern, a motor schema, or something like a bodily movement? Before we consider Wheeler's answer to this, let's consider something that he rules out, namely bodily movement itself (2005, p. 209). According to Wheeler, bodily movements do not have representational status because in the relevant contexts of action, they play a role that can be given a fully causal and specifically non-representational explanation. Rowlands (2006), in contrast, defends the idea that certain bodily movements that are elements of action, can be representational. He argues that the classical concept of representation, modeled on the word – internal but with content externality, in need of interpretation, and passive – is not adequate to capture the concept of representation in action. To get to a more adequate action concept of representation (or AOR) he gives up some of these aspects (2006, 11). For Rowlands, representation in action includes the following characteristics.

- Representation carries information about something other than itself (x).[2]
- Representation is teleological -- it tracks or has a specific function towards x
- Representation can misrepresent x
- Representation can be combined into a more general representational framework
- Representation is decoupleable from x (x may be absent from the immediate environment) (2006, 113-14).

In order to see how this concept of representation applies to action, Rowlands distinguishes between intentional actions, sub-intentional acts, and pre-intentional acts. Sub-intentional acts (O'Shaughnessy 1980) are non-intentional movements, e.g., of tongue or fingers, of which we are not aware, for which there is no reason, and which serve no purpose connected with action. Pre-intentional acts, or "deeds" in Rowland's terminology, include such things as the positioning of fingers in catching a ball that is flying toward you at a high rate of speed, or the movement of your fingers while playing Chopin's *Fantasie Impromptu in C# Minor* on the piano. Pre-intentional acts include an array of "on-line, feedback-modulated adjustments that take place below the level of intention, but collectively promote the satisfaction of [an] antecedent intention" (103).

Rowlands provides a detailed example: Yarbus' (1967) experiments on saccadic eye movements. Yarbus presents subjects with a painting that shows six women and the arrival of a male visitor; subjects are then asked to do certain tasks.

1. View the picture at will
2. Estimate the family's wealth
3. Judge the age of the people in the painting
4. Guess what the people had been doing prior to the arrival of the visitor
5. Remember the clothing worn
6. Remember the position of the objects in the room
7. Estimate how long it had been since the visitor was last seen by the people in the painting.

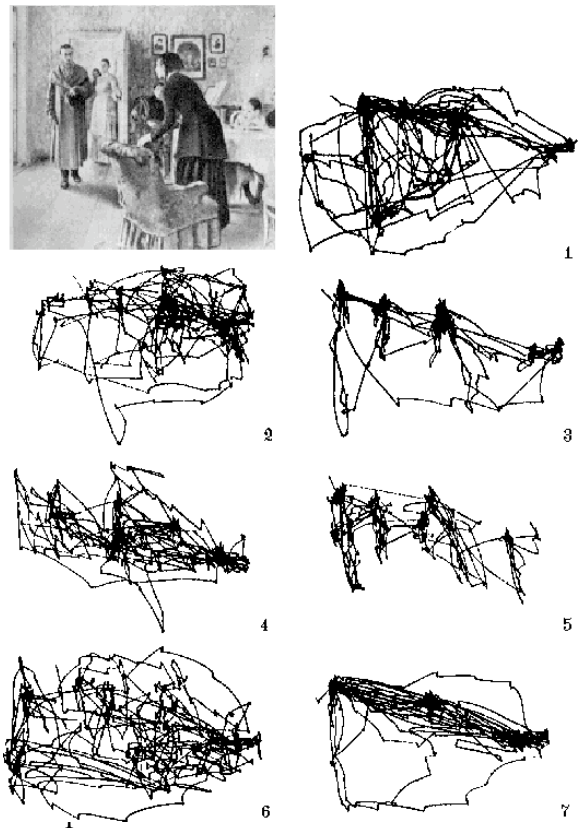

Figure 1: Seven records of eye movements by one subject (from Yarbus 1967).

[2] Although Rowlands suggests that on some alternative account of representation content may be internal to representation (2006, 11), when it comes to explaining the idea that representation carries information (pp. 115ff), in Rowland's account it is difficult to distinguish it from the traditional concept that representation bears content that is external to itself (it refers to or is about something other than itself).

Yarbus found that the scan paths varied systematically with the nature of the task. Thus, the saccades are in some way governed by the intention/task, but they are not intentional in the sense that we do not decide to use this visual tactic, and we are not conscious we are doing the saccades: they are pre-intentional acts.

Rowlands argues that such "deeds" or pre-intentional acts are representational. Although he hesitates to call them representations per se, to call them representational suggests that they involve representations at some level, and in any case they do fit his definition of representation. Pre-intentional acts

- carry information about x (the trajectory, shape, size of ball, the keyboard, a specific aspect of people in painting)
- track x or function in a way that allows the subject to accomplish something in virtue of tracking x
- can misrepresent (get it wrong)
- can be combined into a more general representational structure (I catch the ball in a way that allows me to throw it back efficiently; I continue to play the music; I continue to systematically scan a painting)
- are decouplable from x (x may be absent from the immediate environment -- e.g., I can later remember and demonstrate how I caught the ball replicating the same act; I can imagine looking at the painting and in doing so follow the same scan paths).

On some interpretations (e.g., Anscombe 1957; Merleau-Ponty 1962), "deeds" are intentional insofar as they subserve intentional action. For example, in regard to playing the piano or catching the ball, if you ask me did I intend to posture my fingers in just such a position, I could say yes in so far as I intended to play this piece, or catch the ball. On this view, intentionality reaches down into the motor elements that serve the intentional action. In regard to the Yarbus example, one might ask: "Did you mean to focus on the faces when you were answering that question?" I might respond that I wasn't conscious of doing so, or of controlling my eyes in any explicit way (and in that sense I might claim that it was unintentional). Or I might say "yes, since I was trying to answer a question about the picture, I was certainly scanning it." The issue of whether deeds are pre-intentional or in some sense share in the intention nature of the action can be set aside for our purposes. For simplicity, I'll follow Rowlands and continue to refer to them as pre-intentional.

Decoupleability and causal spread

According to both the standard definition and Rowland's definition of representation, representation is decoupleable from x (I can represent x even if x is absent from the immediate environment). But once we do decouple a pre-intentional act from x (the ball, the piano keys, the painting), I suggest that we are no longer talking about action in the same sense. Indeed, it is difficult to see how pre-intentional acts can decouple from x (the ball, the piano keys, the painting) or the context, without becoming something entirely different from an element of the action at stake, or an AOR. Off-line cognition,

imagining, remembering, or even re-enacting an action decoupled from its original context and absent x, may (or may not) require representation -- but this says nothing about representation *in action*.

An advocate of representation in action could appeal to a model developed by Andy Clark and Rick Grush (1999), however. They offer a model of representation that puts decoupleability directly into action at a sub-personal level. They propose that anticipation in motor control, specifically in the working of a forward emulator, involves a decoupled representation. Since the emulator anticipates (represents) an x that is not there – a future x – or a predicted motor state, it is decoupled from x or the current movement. Thus, "emulators seem to be a nice, biologically detailed example of the sort of disengagement that Brian Cantwell Smith (1996) has recently argued to be crucial for understanding representation" (Clark and Grush 1999, 7). But it is difficult to see how an aspect of motor control that is a constitutive part of the action can be considered decoupled from x, the context, or the action itself. Isn't this kind of anticipation fully situated in the action context? Doesn't the anticipation of a future state or location of x (e.g., anticipating where the ball will be in the next second), or of the predicted motor state (anticipating where to strike the keyboard in the next measure) require reference to the present state or location of x and my hand? The idea of decoupleability seems to interfere with the concept of teleological tracking in this regard. Nor is it clear in what sense this sort of anticipatory simulation/emulation is "off-line" rather than part of the online process of action. If one does decouple the emulation process – if one disengages it from the action itself – it ceases to be part of a forward motor control mechanism, although it may turn into part of a truly off-line representation-process, that is, we may use a decoupled emulation process in memory or imagination (see, e.g., Grush and Mandik 2002, for the example of moving on an imaginary chessboard).

Of course the argument is not that the representation in action *is* decoupled from action; rather, the argument is that the representation is *decoupleable*. So one might admit that once it is decoupled it no longer has any direct function in action itself, although it may assist action planning. But it remains an open question whether in action it functions in a representational manner, or whether it takes on a representational role once it is decoupled. The claim that representations are decoupleable begs the question about how representation is supposed to function. Or to put it a different way, the traditional definition of representation understands representations to be decoupleable only because representations are always portrayed as functioning in a decoupled manner. Representations are needed because we are not directly coupled to the world or the thing that we experience.

Wheeler, for one, gives up the criterion of decoupleablity as part of the concept of a minimal (or weak) representation (219). And if we set the issue of decoupleability aside, Rowland's pre-intentional acts could be considered an example of Wheeler's Action Oriented Representations (AORs). Accordingly, the AOR is not reducible to neural firing patterns, although it does not exclude such patterns. It is clearly something that involves the body-schematic motor control system. So it is some aspect of a system that includes brain, body …. but also environment. "The vehicles of representation do not stop at the skin; they extend all the way out into the world" (Rowlands 2006, 224). Here, without further elaboration, Rowlands joins Clark and Wheeler, and some version of the extended mind hypothesis, where AORs are complex causal interactions that

involve a dynamic coupling of body and environment, and where the causality is spread around.

Wheeler here calls our attention to the "threat [to representationalism] from [non-trivial] causal spread" (2005, 200). Indeed, the commitment to some version of this idea of extended or situated cognition is what motivated anti-representationalism in the first place. On this view, the environment itself does some of the causal work, and in a way that eliminates the need for representations. Consider, for example, driving from London to Durham (see Haugeland 1995). It involves following some kind of strategy.

> Strategy 1: I have a stored inner representation of the directions
> Strategy 2: I follow a map and road signs, which are external representations
> Strategy 3: Having decided to go to Durham, I jump in the car and start off, and having done it many times before, I go on automatic pilot and allow the landscape and roads to guide me (no representations required since the actual road and landscape do the work)

This third situation depends on neither internal nor external representations and involves non-trivial causal spread. But, according to Wheeler, this does not rule out AORs. He argues that to go fully anti-representationalist in an extended cognition paradigm one also needs to understand representation as involving (1) strong instructionalism (i.e., the idea that representations provide a full and detailed description of how to achieve the outcome); and (2) the neural assumption (i.e., that neuronal processes play a central and close to exclusive role). On this understanding of representation, we can easily go anti-representational because (2) is already weakened on the extended cognition hypothesis; and we can easily give up (1) -- no need for anything like a fully-specified representation. For Wheeler, however, this does not rule out a minimal form of representation, and he still maintains that there are AORs that are distributed across brain, body, and environment.

Where's the representation?

The idea of strong instructionalism was already given up by Dreyfus in his discussion of background knowledge -- a representation no matter how strongly instructional can never be adequate to meet the background problem or the frame problem -- it can never specify everything necessary to determine an unruly contextualized action; indeed it would lead to a paralyzed system and inaction if the system attempted to specify everything required. The context and background problems are not solvable by appeal to representations -- the more positive embodied-enactive-extended account is required.

Dreyfus (2007) thus appeals to Merleau-Ponty's work, which offers a nonrepresentational account of the way the body and the world are coupled and suggests a way of avoiding the frame problem. According to Merleau-Ponty, as an agent acquires skills, those skills are "stored," not as representations in the agent's mind, but as dispositional embodied responses to the solicitations of situations in the world. What the learner acquires through experience is not by way of representations at all but by way of more finely discriminated situations. If the situation does not clearly solicit a single

response or if the response does not produce a satisfactory result, the learner is led to further refine his discriminations, which, in turn, solicit ever more refined responses.

Note that on this model we can indeed get things wrong, but not because our representation of the world misrepresents the world. Rather, the world itself is ambiguous in the light of our particular abilities and projects. From a particular distance and perspective, or in a certain light, the mountain appears to be climbable. Once I get closer, or begin to climb, however, I can discover that the mountain is not climbable. On the representationalist view this is explained by saying that my original representation of the mountain was wrong. On the embodied, non-representationalist view, at a certain distance, in a certain light, from a certain perspective the mountain presents a certain affordance relative to my embodied skills,. Change the distance, light and/or perspective and the affordance disappears – that is, the dynamic coupling of body and environment changes. These things are physically determined factors that involve a real mountain, light conditions, and my bodily position and capabilities; they are not representational. The affordance doesn't disappear because I change the representation of my distance from the mountain – I actually have to change my distance, and when I do so, the body-mountain relation, which defines the affordance, changes.

Wheeler too appeals to a non-representational solution to the frame problem. Evolutionary and cultural contributions to the already situated subject inform a perception-based account of action. Lions, and tigers, and bears as well as other people, and specific objects in the environment, e.g., bombs, have either an evolutionary or culturally-based valence that solicits a particular response. Intentional life is keyed into relevant aspects of the environment in a way that shapes our subpersonal processes, and the latter come to serve the intentional aspects of action. In this sense, the frame problem is solved not by a network of representations, but by intuitive and emotionally informed reactions. This is why it is still a problem for AI and robotics. We try to get the robot to recognize the bomb as a threat by providing it with a cold propositional algorithm to that effect. The robot thus has to represent the bomb as a threat and then has to represent what action to take and then has to represent what parts of its mechanisms to activate. In contrast, we non-robots *see* that the bomb is a threat and we run for our lives.

Within such embodied-embedded-extended approaches, what role does a minimal representation play? Wheeler defends representation (AOR) as a perception-based, short-lived, egocentric (spatial) mapping of the environment calibrated strictly in terms of possible actions. Clark and Grush suggest that the anticipation that is built into a forward emulator for online motor control is representational. Rowlands argues that pre-intentional movement that is governed by an intentional action is representational. When we consider these aspects of action together we should notice that they reflect nothing more nor less than the dynamic temporal structure of action experience itself. On a phenomenological, non-representational model of this temporal structure the short-term mapping of the environment is a function of (1) a pragmatic retentional maintenance of the relevant aspects of the environment that has just been experienced (a holding in perceptual presence of those factors relevant to possible actions); (2) the anticipation that is essential to motor control, a protentional aspect that is an implicit characteristic of my immediate project-determined coupling with the environment, and (3) the current movement that contributes to the very structure of the action. In short, action involves an

ongoing retentional-protentional, short-lived, egocentric orientation to the environment calibrated in terms of possible actions.[3]

This retentional-protentional structure, which is characteristic of consciousness as well as action, is fully online in what Husserl calls the 'living present'. The retention of the just past action experience (including the action-oriented experience of the environment) in the present moment is neither a recollection (memory) nor a representation of that experience that one would somehow add to the present experience; it is rather part of what constitutes the action in its ongoing directionality. The protentional anticipation of action experience points to where the action is heading, not in the sense that it represents where it is heading, but in the sense that the ongoing action has a heading – that the action is already going in a certain direction, heading toward a completion (that it may or may not attain). This dynamic structure of being-in-process is reflected in action experience. An action is not a momentary or frozen snapshot supplemented by representations of past and future movements; it has a unity over time that is accounted for in the intentional structure of the action itself. Nothing in this dynamically dissipating process amounts to a representation, if we take representation to involve

- an internal image or symbol or sign
- a discrete duration
- decoupleability

A representationalist might reply that even if this describes the absence of mental representations in the experience of action, there are certainly sub-personal representations that underlie this retentional-protentional structure. The action is itself dependent on certain neural representations or body schemata that operate at a sub-personal level, outside of or below the threshold of experience. There is no doubt that there are neural processes involved in body-schematic motor control, but these sub-personal body-schematic processes operate on the same dynamic model reflected in experience. The ubiquity of anticipatory mechanisms in the sensory-motor system (sometimes specified in terms of efferent 'anticipation for the consequences of the action' [Georgieff and Jeannerod 1998]) are well known (Berthoz 2000; Berthoz and Petit 2006). The notion of body schema has been characterized not as a static representation, but as a process that includes a retentional component that dynamically organizes sensory-motor feedback in such a way that the current motor state is "charged with a relation to something that has happened before" (Head 1920, 606). Neither the relations nor the relata, however, are discrete or decoupleable. Neither body schemata nor neuronal patterns are discrete (or have discrete durations). Neurons, of course, are part of the highly connected complex system of the brain in which connections are effected, not by representation but causally. If body schemas were reducible to neuronal firing patterns, they would be characterized in the same way. Alternatively, if body schemas are complex processes that extend over brain and body, and are in all cases specified by environmental contexts (see Gallagher 2005), they are not constituted by one part

[3] For the interpretation of Husserl's analysis of the retentional-protentional temporal structure of experience in terms of non-representationalist dynamic systems theory, and its application to motor control, see Thompson 2007; Van Gelder 1999; Varela 1999; and Gallagher and Varela 2002.

representing another part, or by one part interpreting the other as a representation. In such neuronal or body-schematic systems what could count as representation would be purely a matter of interpretation, not by the subject or the system, but by the scientist abstracting (drawing discrete lines between one schema and the next; claiming decoupleability, etc.) from the system.[4]

Wheeler, as we noted, gives up the criterion of decoupleablity in his characterization of a minimal (or weak) representation (2004, 219). On his account, a minimal representation (1) is richly adaptive, (2) is "arbitrary" or ad hoc – in the sense that it is not predefined, but processes current information about the world, and (3) employs a homuncular mechanism, i.e., a mechanism that is hierarchically compartmentalized but contributes to a collective achievement. With the idea of the homuncular mechanism Wheeler attempts to preserve the criterion of interpretability within the system itself. Representational interpretation can be conceived of as involving modularity -- processing in one module independent of processing in another, but each communicating results to (and mutually interpreting) another module. The homuncular mechanism thus takes some information "off-line" (but, according to Wheeler, without decoupling it from the action itself) and manipulates it to anticipate possible actions. This seems to be nothing more than Clark and Grush's emulator *sans* decoupleability. At the same time it is not clear what "off-line" but not decoupled means, or, as I indicated above, how an anticipation of possible action can be formed without reference or intricate connections to the current situation. In fact, in the case of action, modularity can be given up for the dynamic systems concept of a self-organizing continuous reciprocal causation (Varela, Clark) which Wheeler himself favors in most instances. On-line sensory-motor processes that are serving intentional action and are temporally structured in dynamic relation to the environment are in fact richly adaptive and arbitrary in the relevant sense, but are not homuncular, which means they involve no interpretational element. The dynamical process (more causal than communicative) does not require the idea that one discrete part of the mechanism interprets in isolation (or off-line) the information presented by another part. Rather, action itself, on the dynamic model, is characterized by an anticipatory, protentional aspect that functions only in relation to the on-going, online, project-determined coupling with the environment.

What's left of the idea of representation in action?

At this point, however, we can surely ask, what's the point in retaining the term 'representation' in the case of action? What work does the concept of representation

[4] Borrowing on Menary's (2007) reading of Peirce's semiotics, one simply way to put this is that one of the triadic elements of the representational process is missing in the case of neuronal events or subpersonal processes more generally. For Peirce, "representation necessarily involves a genuine triad" (1931, 1.480). It involves a vehicle (sign) mediating between an object and an interpreter. A neuronal pattern or event might be considered a representational vehicle, but only in connection with an object (some event in the environment, perhaps) and a consumer or interpreter (to produce an interpretant or meaning). The missing element is the consumer (interpreter). The experiencing subject is not an interpreter of its own brain events, but neither is the brain itself, unless one is willing to say that one process in the brain interprets another process in the brain as a sign of something happening in the environment (see the next paragraph). On the Peircean model, if one of the elements is missing, there is no representation.

really do since nothing is being re-presented to the subject; since it is not consistent with the classical notion of representation; and since in working out the justification one is already explaining action in non-representational terms of perception-based complex causal interactions in an extended-body-environment system. A facetious economic argument against representationalism would suggest that the explanatory work that the concept of representation does is less than the work it takes to justify the use of the term 'representation'.

Here is a detailed negative characterization of minimal representation.

1. Minimal representation is not internal - it extends to include embodied-environmental aspects and is only "weakly" neuronal.
2. Minimal representation is not a discrete identifiable enduring thing -- it's more like a temporal, dynamic, and distributed process.
3. Minimal representation is not passive -- it's pragmatically enactive - proactively contributing to the adaptability of the system.
4. Minimal representation is not decoupleable -- indeed, if it is to remain teleological, it has to continue tracking x or it has to involve a continuing and online anticipation or protention of a predicted motor state.
5. Minimal representation is not strongly instructional -- even if it can be combined into a more general representational framework, it is never sufficiently strong enough to solve the commonsense knowledge or the frame problem.
6. Minimal representation is not homuncular and does not involve interpretation.

In effect, the idea of a minimal representation no longer conforms to the criteria that would make it a representation. Actions do involve processes that are *intentional*, certainly at the personal level, and in a way that contributes to the organization of the sub-personal processes that support the intentional action (motor control processes and pre-intentional acts that contribute to the accomplishment of actions). But if representation is one form of intentionality; not all intentionality is representational. The kind of "motor intentionality" described in terms of body schematic processes by Merleau-Ponty (1962), for example, is a non-representational dynamic process. Actions also involve *teleological* functions insofar as they sometimes require tracking something in the world. This is a perceptual tracking -- I see the ball that I want to catch, and I bodily respond to its trajectory. Actions are also *fallible* -- action can fail, not because of a misrepresentation, but because perception is finite: things look climbable but turn out not to be; things look catchable, but often turn out not to be, etc.

There is thus an intentionality of the body-in-action that is not characterized as internal, decoupleable, or instructional; that does not involve interpretation in the relevant sense; and that is accordingly non-representational. This kind of intentionality is dynamically linked with the environment in a way that reflects a specific temporal structure at the subpersonal level. As Wheeler puts it:

> […] as the brain becomes ever more bound up in complex distributed [and extended, non-neural] causal interchanges with the non-neural body and the

> wider physical environment, it seems likely that the temporal character of those interchanges will become increasingly rich (Wheeler 2005, 244).

Action involves temporal processes that can be better explained in terms of dynamic systems of self-organizing continuous reciprocal causation at the subpersonal level. Action involves a retentional-protentional, short-lived, egocentric orientation to the environment calibrated in terms of possible actions (as Wheeler would have his AORs), and this is dynamically manifested in current pre-intentional movements ("deeds") that serve the intentional action.

Scientific pragmatism about representations

Do representational accounts provide a helpful short-cut for explaining action? Or do they explain anything about action at all? At best, representationalism is just one way -- a scientifically abstract way -- of explaining the action process. But a representation is not an *explanans* that does any work itself. It's a concept under which one still needs all the explanation to be made. The risk is that representational accounts come with ontological claims -- there really are discrete representations in the system and they are something more than what a motor control system does as part of the action itself.

The problem is that a majority of cognitive scientists continue to use the R-word and do so in ways that are not often clear. I suggest that, in the case of action, it is sometimes nothing more than a handy, but often confused and misleading term that is nothing other than a place-holder for an explanation that needs to be cast in dynamical terms of an embodied, environmentally embedded, and enactive model.

In this regard, however, even if you think that the concept of representation does do some explanatory work, what I identified as the facetious economic argument against representationalism is really more pragmatic than facetious. It may take more energy to define and distinguish any legitimate sense of representation from amongst the plethora of uses of that term, and to justify it's use, than it would take to explain the phenomenon in non-representationalist terms. And if one can explain the phenomenon in non-representationalist terms, then the concept of representation is at best redundant.

References

Anscombe, G. E. M. 1957. *Intention*. Oxford: Blackwell.

Berthoz, A. 2000. *The Brain's Sense of Movement.* Cambridge, MA: Harvard University Press.

Berthoz, A. and Petit, J-L. 2006. *Phénoménologie et physiologie de l'action*. Paris: Odile Jacob.

Clark, A. and Grush, R. 1999. Towards a cognitive robotics. *Adaptive Behavior* 7 (1): 5-16.

Dreyfus, H. 2007. Why Heideggerian AI Failed and How Fixing it Would Require Making it More Heideggerian. *Philosophical Psychology* 20 (2): 247–268

Dreyfus, H. 2002. Intelligence without representation: Merleau-Ponty's critique of mental representation. *Phenomenology and the Cognitive Sciences* 1 (4): 367-383.

Gallagher, S. and Francisco Varela. 2003. Redrawing the map and resetting the time: Phenomenology and the cognitive sciences. *Canadian Journal of Philosophy*. Supplementary Volume 29: 93-132.

Georgieff, N. and Jeannerod, M. 1998. Beyond consciousness of external events: A 'Who' system for consciousness of action and self-consciousness. *Consciousness and Cognition*, 7: 465-77.

Grush, R. and Mandick, P. 2002. Representational parts. *Phenomenology and the Cognitive Sciences* 1 (4): 389-94.

Haugeland, J. (1998). Mind embodied and embedded. *Having thought: Essays in the metaphysics of mind* (pp. 207–237). Cambridge, MA: Harvard University Press.

Head, H. 1920. *Studies in Neurology*. Vol 2. London: Oxford University Press.

Menary, R. 2007. *Cognitive Integration: Mind and Cognition Unbounded*. London: Palgrave-Macmillan.

Merleau-Ponty, M. (1962). *Phenomenology of perception* (C. Smith, Trans.). London: Routledge & Kegan Paul.

Merleau-Ponty, M. (1966). *The structure of behavior* (A. L. Fisher, Trans., 2nd ed.). Boston: Beacon Press.

Peirce, C. S. 1931. *Collected Papers of Charles Sanders Peirce*, vol. 1. (eds. C. Hartshorne and P. Weiss). Cambridge, MA: Harvard University Press.

O'Shaughnessy, B. 1980. *The Will*, 2 vols. Cambridge: Cambridge University Press.

Rowlands, M. 2006. *Body Language*. Cambridge, MA: MIT Press.

Thompson, E. 2007. *Mind in Life: Biology, Phenomenology and the Sciences of Mind*. Cambridge, MA: Harvard University Press.

van Gelder, T. 1999. Wooden iron? Husserlian phenomenology meets cognitive science. In J. Petitot, F. J. Varela, J.-M. Roy, and B. Pachoud (eds.), *Naturalizing Phenomenology: Issues in Contemporary Phenomenology and Cognitive Science*. Stanford: Stanford University Press.

Varela, F. J. 1999a. The specious present: A neurophenomenology of time consciousness. In J. Petitot, F. J. Varela, B. Pachoud, and J.-M. Roy (ed.), *Naturalizing Phenomenology: Issues in Contemporary Phenomenology and Cognitive Science* (pp. 266-314). Stanford: Stanford University Press.

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

Yarbus, A. 1967. *Eye Movements and Vision*. New York: Plenum Press.