



Extended cognition and the mark of the cognitive

Dr. Mark Rowlands

To cite this article: Dr. Mark Rowlands (2009) Extended cognition and the mark of the cognitive, *Philosophical Psychology*, 22:1, 1-19, DOI: [10.1080/09515080802703620](https://doi.org/10.1080/09515080802703620)

To link to this article: <https://doi.org/10.1080/09515080802703620>



Published online: 21 Feb 2009.



Submit your article to this journal [↗](#)



Article views: 865



Citing articles: 61 [View citing articles ↗](#)

Extended cognition and the mark of the cognitive

Mark Rowlands

According to the thesis of the extended mind (EM), at least some token cognitive processes extend into the cognizing subject's environment in the sense that they are (partly) composed of manipulative, exploitative, and transformative operations performed by that subject on suitable environmental structures. EM has attracted four ostensibly distinct types of objection. This paper has two goals. First, it argues that these objections all reduce to one basic sort: all the objections can be resolved by the provision of an adequate and properly motivated criterion—or mark—of the cognitive. Second, it provides such a criterion—one made up of four conditions that are sufficient for a process to count as cognitive.

Keywords: Cognition; Extended Mind; Mark of the Cognitive; Vehicle Externalism

1. The Extended Mind and its Discontents

According to the view known variously as the *extended mind* (Clark & Chalmers, 1998), *vehicle externalism* (Hurley, 1998; Rowlands, 2006) *active externalism* (Clark & Chalmers, 1998), *locational externalism* (Wilson, 2004) and *environmentalism* (Rowlands, 1999), at least some token cognitive 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, etc.

Mark Rowlands is professor of philosophy at the University of Miami.

Correspondence to: Dr. Mark Rowlands, University of Miami, Philosophy, P.O. Box 248054, Coral Gables, FL 33124. Email: mrowlands@miami.edu

- Cognitive 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 the subject with the ability to appropriately use relevant structures in its environment.

EM has attracted a number of objections, broadly divisible into four categories.

(1) *The Differences Argument*. This type of objection points to the significant differences between internal cognitive processes and the external processes that *EM* alleges are also cognitive. This casts doubts on the claim that both processes should be regarded as belonging to a single psychological kind. This sort of objection has been vigorously championed by Rupert (2004).

(2) *The Coupling-Constitution Fallacy*. This objection claims that *EM* confuses cognition with its extraneous causal accompaniments. More precisely, it confuses those structures and processes constitutive of cognition with those in which cognition is (merely) causally embedded. This type of objection has been developed by Adams and Aizawa (2001, 2009) and also, in a somewhat different way, by Rupert (2004).

(3) *The Cognitive Bloat Objection*. According to this objection, the admission of extended cognitive processes places one on a slippery slope. Once we permit such processes, where do we stop? Our conception of the cognitive will become too permissive, and we will be forced to admit into the category of the cognitive all sort sorts of structures and processes that clearly are not cognitive.

(4) *The Mark of the Cognitive Objection*. According to this objection, *EM* should be rejected on the grounds that it is incompatible with any plausible mark of the cognitive; that is, any criterion that specifies the conditions under which a process qualifies as cognitive. This objection is developed by Adams and Aizawa (2001, 2009).

This paper has two primary goals. First, it argues that the mark of the cognitive objection is most basic in that all the other objections either reduce it, or can be solved by the provision of an adequate and properly motivated criterion of the cognitive. Second, it provides such a criterion—one made up of four conditions that a process must satisfy in order to qualify as cognitive.

2. The Differences Argument: Parity and Integration in *EM*

EM is often thought to be grounded in the concept of *parity*: roughly speaking, the *similarity* between the external processes involved in cognition and internal processes that are widely accepted as cognitive.¹ *EM*'s reliance on this notion of parity is often thought to be embodied in, and demonstrated by, Clark and Chalmers' (henceforth *C&C*) deployment of what they call the *parity principle*. They write: 'If, as we confront some task, a part of the world functions as a process which, were it done in the head, we would have no hesitation in recognising as part of the cognitive

process, then that part of the world is (so we claim) part of the cognitive process.’ (Clark and Chalmers, 1998, p. 12)

Critics of *EM* have, without exception, but almost certainly mistakenly, understood the parity principle as introducing a *similarity-based criterion* of when an external process or structure is to be understood as cognitive—that is, as a genuinely cognitive part of a cognitive process. Very roughly, on this understanding of the parity principle, if an external process is sufficiently similar to an internal cognitive process, then it too is a cognitive process.

It is this interpretation of the role of the parity principle—as introducing a similarity-based criterion for when cognition can legitimately be regarded as extended—that underwrites the *differences argument*. Thus Rupert, in connection with an argument for extended memory developed by me (1999), outlines his strategy as follows:

I argue that the external portions of extended “memory” states (processes) differ so greatly from internal memories (the process of remembering) that they should be treated as distinct kinds; this quells any temptation to argue for HEC from brute analogy (*viz.* extended cognitive states are like wholly internal ones; therefore, they are of the same explanatory cognitive kind; therefore there are extended cognitive states). (Rupert, 2004, p.407)

The operative assumption is that the function of parity principle is to introduce a similarity-based criterion of when a cognitive process such as remembering can be extended into the world: *if* an external process is sufficiently similar to internal cognitive processes, *then* it too is a cognitive process. Rupert then argues that since external processes involved in memory are, in fact, *not* sufficiently similar to internal cognitive processes, then they are *not* cognitive processes. Presumably this is intended as an inductive argument: since the internal and external processes involved in cognition are not sufficiently similar, the parity principle provides no reason to regard the latter as cognitive.²

This *differences argument*, however, rests on a failure to properly understand the arguments for *EM*. *EM* does not rely on a similarity-based criterion of when a cognitive process may legitimately be regarded as extended. The notion of parity is indeed, I shall argue, an important one for *EM*. However, equally important is the notion of *integration*: the meshing of disparate types of process that, *precisely because* of their disparate character can enable a cognizing organism to accomplish tasks that it would not be able to achieve by way of either type of process alone (Menary, 2006, 2007; Sutton, 2006). From this *integrationist* perspective, the *differences* between internal and external processes are as important as, or even more important than, the similarities. The reason cognition extends into the environment is precisely because, with respect to the accomplishing of certain cognitive tasks, external processes can do things that internal processes cannot do (or, depending on how you want to understand *EM*, simply, in certain cases, *do not do*). External structures and processes possess quite different properties from internal ones; and it is precisely this difference that affords the cognitive agent the opportunity to accomplish certain

tasks that it could not, or might not, be able to accomplish purely by way of internal cognitive processes. Without these differences, the external processes would be *otiose*.

Thus, for example, extended models tend to emphasize the relative *stability* of relevant external structures, and the enhanced possibilities for manipulation and exploitation that this stability engenders (Donald, 1991; Noë, 2004; O'Regan & Noë, 2001; Rowlands, 1999). These possibilities, it is argued, have little or no echo in the case of internal processes, and they underwrite the abilities of organisms to accomplish certain cognitive tasks that they could not (or, perhaps, simply do not) accomplish by way of internal processes alone. *EM* tends also to emphasize the distinctive *structure* of external items—for example, linguistic, combinatorial, etc—structure that, again, has (arguably) no echo in internal items (Donald, 1991; Hurley, 1998; Rowlands, 1999). In each case, it is precisely the *different* properties of external structures that allow the cognitive agent to accomplish things that it either could not, or in fact does not, accomplish by way of internal processes alone.

Given the central role played by the notion of integration in *EM*, one cannot predicate, as does the *differences argument*, an objection to *EM* simply by citing differences between internal and external processes. *EM*, properly understood, both *predicts* and *requires* such differences. Understood on its own terms, therefore, the differences argument fails.

However, the integrationist's emphasis on the *differences* between the internal and external processes involved in cognition does leave *EM* vulnerable to another objection. If *EM* requires significant differences between internal processes and the external processes that it regards as cognitive, what reason is there for supposing that the latter are really part of cognition rather than a merely external accompaniment to real, internal, cognitive processing? Given the integrationist's emphasis on the differences between the internal and external processes involved in cognition, it is not possible to establish the cognitive status of the latter simply by analogical extension from the former. Therefore, if *EM* is to defend the cognitive status of the extended processes, it needs to provide an adequate and properly motivated *criterion of the cognitive*: a criterion that would allow *EM* to justify the claim that the external processes involved in cognition are indeed cognitive processes. In short, the integrationist response deflects the differences argument only by leaving *EM* vulnerable to the mark of the cognitive objection.

3. The Coupling-Constitution Fallacy

The *coupling-constitution fallacy* objection can take slightly different forms. According to Adams & Aizawa (henceforth, *A&A*):

This is the most common mistake that extended mind theorists make. The fallacious pattern is to draw attention to cases, real or imagined, in which some

object or process is coupled in some fashion to some cognitive agent. From this, they slide to the conclusion that the object or process constitutes part of the agent's cognitive apparatus or cognitive processing. (2001, p. 408)

Rupert expresses a similar objection, albeit in more cautious terms. Referring, again, to my version of *EM*, he writes:

Rowlands, however, does not make clear why the use of an internally represented code applied to the contents of an external store implies HEC, rather than what it would seem to imply: HEMC (Rupert, 2004)

HEC is the *hypothesis of extended cognition*, which Rupert correctly distinguishes from HEMC, the *hypothesis of embedded cognition*. What reason, Rupert asks, do we have for regarding the external processes as part of cognition rather than simply a form of extraneous scaffolding in which *real*—internal—cognitive processes can be causally embedded?

It is, however, implausible to suppose that *EM* is guilty of simply *confusing* constitution and causal coupling. Far from *confusing* constitution and causal coupling, the most natural way of understanding the arguments for *EM* are precisely as *arguments* for reinterpreting what had traditionally been regarded as extraneous causal accompaniments to cognition as, in fact, part of cognition itself. And, in general, to *argue* for the identification of X and Y, when X and Y had hitherto been regarded as distinct types, is not to *confuse* X and Y.

Consider, for example, my arguments for extended memory cited by Rupert. I argue that, in certain cases, the external processes involved in cognition—bodily manipulation and exploitation of information bearing-structures in the cognizer's environment—possess certain abstract, general, features of processes commonly regarded as cognitive, while also differing in the sorts of concrete ways required by the integrationist underpinning of *EM*. Thus, these external processes are employed in order to accomplish cognitive tasks. They involve information processing—the manipulation and transformation of information bearing structures. This processing results in the making available to organisms of information that was previously unavailable; and so on. That is, I *argue* for the cognitive status of external processes of these sorts by trying to show that they satisfy a certain criterion of the cognitive. One can, I think, legitimately question whether I rendered this criterion sufficiently explicit, and even if it were explicit whether it is adequate. But one can hardly accuse him of *confusing* causation and constitution. And if I were in possession of an adequate and properly motivated criterion of the cognitive, and if the sorts of external processes he identifies were to satisfy this criterion, then he would have, in fact, made it clear why his view implies HEC rather than HEMC.

Thus, like the differences argument, the coupling-constitution fallacy objection is derivative on the mark of the cognitive objection. If *EM* can provide an adequate criterion of the cognitive, and demonstrate that the external processes it regards as cognitive satisfy this criterion, then there is no substance to the charge that it confuses constitution and mere coupling.

4. Cognitive Bloat

The *cognitive bloat objection* is, in essence, a slippery slope argument usually raised in connection with C&C's discussion of Otto's notebook. Otto is in the early stages of Alzheimer's and so keeps a notebook with him in which he writes down various facts that can help him with his day-to-day living. C&C argue, notoriously, that the sentences in Otto's book can constitute a subset of his beliefs. The entries in Otto's book are literally beliefs that Otto possesses because they have a functional role in Otto's truncated psychological economy that is sufficiently and relevantly similar to the role played by beliefs in the psychological economy of Otto's unimpaired friend Inga.

The cognitive bloat objection uses this claim as a starting point. If we are willing to allow that the sentences in Otto's notebook are beliefs, why stop there? Why not the entries in the telephone directory that of which Otto also makes frequent use? Why can these not be numbered among Otto's beliefs? Indeed, why stop even there? Why does Otto not believe everything contained on the internet, given that he is able to use this in a way akin to which he uses his notebook?

C&C try to preclude this problem of bloat by advocating a *conscious endorsement* criterion on beliefs. The entries in Otto's notebook count as beliefs, whereas the entries in Otto's telephone book do not, because Otto has, at some point, consciously endorsed the former but not the latter. However, this condition is questionable: beliefs can be formed subliminally, as well as through conscious experience, and, presumably, we would not regard their mode of formation as automatically excluding them from the class of cognitive states.

I shall argue – although this argument must be postponed until later – that the cognitive bloat objection can be rebutted by way of an adequate criterion of the cognitive. The key to this rebuttal lies in the role played by the concept of *ownership* in qualifying a state or process as cognitive. The entries in the telephone directory, and the pages on the internet, are not, at least as they figure in the cognitive bloat objection, *owned* by anyone. And anything that is to count as cognitive, I shall argue, must be owned by someone or some thing. Discussion of this claim must be postponed until I have proposed and defended a criterion of the cognitive.

5. The Mark of the Cognitive

According to the *mark of the cognitive objection*, *EM* should be rejected on the grounds that it is incompatible with any plausible mark or criterion of the cognitive. The goal of the second half of this paper is to supply such a criterion, and argue that far from contradicting the claims of *EM* it actually supports those claims. Therefore, if the other objections to *EM* are indeed derivative on the mark of the cognitive objection, then provision of this criterion will serve to defuse these objections also.

Some are skeptical about the need for a mark of the cognitive. Do we need, they ask, a mark of the biological in order to do biology; or, for that matter, a mark of the physical in order to do physics? (The question is intended as rhetorical).

Underlying this attitude is the idea that science simply does what it does—identifies its laws and constructs its theories—and, as long as it can do this, has no need for any deeper understanding of what it is doing. Thus, it might be argued, we have an adequate intuitive grasp of what counts as cognitive, and this grasp is sufficient for us to adjudicate the claims of *EM*.³ There is another way of thinking about science—as a form of self-interpreting activity—according to which a gradual deepening of the understanding of what a science is doing is built into the scientific project itself.⁴ The sort of philosophy of psychology that attempts to understand what it is for a process to be cognitive is built on this alternative vision of the scientific enterprise. I don't propose to adjudicate between these two conceptions of science. However, it is worth pointing out that in the case of *EM* at least, it is not only the mark of the cognitive objection that is at stake. If the arguments of the first half of the paper are correct, then *all* of the objections to *EM* can be solved through the provision of an adequate mark of the cognitive. This might be enough to at least convince defenders of *EM* that a criterion of the cognitive is something worth investigating.

In the prevailing dialectical situation it is clear *why* *EM* needs a criterion of the cognitive; but it is not yet clear precisely *what* sort of thing this criterion should be. Questions can be raised about both the *scope* and *character* of the criterion. Consider, first, the issue of scope. The term 'cognition' can, in effect, be spelled with a big 'C' or small 'c'. With a big 'C', the sense of the term is fairly narrow. In this sense, cognition is routinely opposed to perception (and, of course, sensation). That is, in this sense, cognition is restricted to post-perceptual processing. However, there is a broader sense of the term, 'cognition' with a little 'c', where it is understood to include perceptual processing. Reference to cognition, in this paper, should be understood as reference to cognition in the more general sense; 'cognition' with a little 'c'. And the proposed criterion of 'the cognitive', is one that attempts to demarcate items that are cognitive in this broader sense from items that are not.

Consider, now, the issue of the *character* of the criterion. What sort of thing would the mark of the cognitive have to be? Here, two quite distinct projects need to be distinguished. On the one hand, there is the philosophical project of *naturalizing* the mind. If we were engaged in this project, we might take ourselves—intelligibly but, I think, mistakenly—to require a reductive definition of 'the cognitive' in terms that are entirely non-cognitive (or, initially at least, less than fully cognitive). On the other hand, there is the cognitive-scientific project of understanding how cognitive processes work. This project consists in functional decomposition of these processes into progressively simpler constituents. To engage in this project, we do not require a naturalistic reduction of the cognition—in the form of a definition of 'cognition' in non-cognitive terms. Rather we simply require a means of demarcating, with a reasonable but not necessarily infeasible level of precision, those things in which we are interested from those things in which we are not.

I shall propose and defend a criterion of this latter sort. The idea underpinning the criterion is that if we want to understand what cognitive processes are, then we had better pay close attention to the sorts of things cognitive scientists regard as cognitive. That is not to say that we must restrict ourselves to the pronouncements or

determinations of cognitive scientists, or that we should regard these as decisive, but merely that we had better be prepared to use these as our starting point. A significant part of the criterion I shall defend can be extracted from a careful examination of cognitive-scientific *practice*. When we examine such practice, I shall argue, what we find is an implicit mark of the cognitive that looks like this:

A process *P* is a *cognitive* process if and only if:

- (1) *P* involves information processing—the manipulation and transformation of information-bearing structures.
- (2) This information processing has the *proper function* of *making available* either to the *subject* or to *subsequent processing operations* information that was (or would have been) prior to (or without) this processing, unavailable.
- (3) This information is made available by way of the production, in the subject of *P*, of a *representational state*.
- (4) *P* is a process that *belongs* to the *subject* of that *representational state*.

This criterion is understood as providing a *sufficient* condition for a process to count as cognitive: if a process satisfies these conditions, it counts as cognitive. Whether it also provides a *necessary* condition is an issue that I shall leave aside (although, for what it's worth, I suspect that it does not).

Before motivating and defending this criterion, it is necessary to remark on these four conditions.

Condition (1): The idea that cognition involves information processing is now commonplace. In its classical form, cognitive science was understood to involve the postulation of internal configurations of an organism or system: configurations that carry information about extrinsic states-of affairs. The concept of information employed is, in essence, that elaborated by Claude Shannon, or a close variant thereof.⁵ According to Shannon, information is to be understood in terms of relations of conditional probability. On the version championed by Shannon, a receptor *r* carries information about a source *s* only if the probability of *s* given *r* is 1 (c.f. Dretske, 1981). Other, less sanguine, versions associate the carrying of information with an increase in conditional probability, although not necessarily to the value of 1. That is, *r* carries information about *s* only if the probability of *s* given *r* is greater than the probability of *s* given not-*r* (c.f. Lloyd, 1989).

Whichever explication of the concept of information is employed, the underlying vision of cognition is unaffected. Cognitive processes are understood as a series of transformations performed on information-bearing structures. These transformations will be effected according to certain rules or principles—principles that effectively define the character of the type of cognitive process in question.

Condition (2): The second condition relies heavily on the concept of *proper function*. I shall understand this in the etiological sense championed by Millikan, among others (Millikan, 1984, 1993). The notion of a proper function is not essentially a biological one, but rather applies to any item that has the appropriate sort of history. The appeal to proper function reminds us that the concept of cognition is, in part, a normative one. Cognitive processes can function well or badly, properly or

improperly; they are defined in terms of what they are *supposed* to do, not what they *actually* do. This important distinction can manifest itself in at least four ways:

- (i) While the proper function of a cognitive process, P, is to make available information that was previously unavailable, it might not fulfil this function in particular cases. Indeed, it might never fulfil this function, either due to faults in the mechanisms that realize it, or faults in the mechanisms that are supposed to receive the information it supplies.
- (ii) In addition to performing its proper function, P might also do all sorts of things that are not its proper function.
- (iii) P might have the proper function of making information available to a subject or subsequent operations only in the presence of environmental contingency C: and C may sometimes fail to obtain; in principle it may never obtain.
- (iv) P may have the proper function of making information available only in combination with a number—perhaps a vast number—of other processes. In such circumstances, P's fulfilling its proper function is dependent on these other processes fulfilling their proper functions.

The distinction between making information available to a *subject* and making it available to *subsequent processing operations* is also important. Roughly, and with appropriate clarifications to follow shortly, it corresponds to the distinction between *personal* and *sub-personal* cognitive processes. Processes that make information available to a subject are personal level processes, and this is true whether they also make information available to subsequent processing operations. However, processes that make information available only to subsequent processing operations are sub-personal cognitive processes.

Condition (3): This condition takes no stand on whether representation is, ultimately, a naturalistic phenomenon. Many suppose that it is possible to supply broadly naturalistic conditions of representation, and that an item *r* will qualify as representational if and/or only if it satisfies these conditions. This naturalistic assumption is neither mandated by, nor incompatible with, (3).

The notion of a representational state is not co-extensive, still less synonymous, with the notion of a *semantically evaluable* state. It would be implausible to suppose that all representational states are semantically evaluable: such states must possess *adequacy conditions*, but they need not possess *truth-conditions*. Mental models, or cognitive maps, possess adequacy conditions but not truth-conditions. The concept of truth is constitutively connected—for familiar Davidsonian reasons—with the logical connectives; and mental models do not relate to these connectives in the same way as sentences—external or internal. Thus, for example, the negation of a sentence is another *specific* sentence; but to the extent that it makes sense to speak of the negation of a map or model, this can mean nothing more than a distinct, but non-specific, map or model. Once we understand the connection between truth and the logical connectives, I take it that this point is incontrovertible, though often oddly neglected. However, nothing much will turn on it in the arguments to follow.

I shall assume that the type of representational state invoked in (3) is one that possesses *non-derived* content. Derived content is content, possessed by a given state,

that derives from the content of other representational states of a cognizing subject or from the social conventions that constitute that agent's linguistic milieu.⁶ Non-derived content is content that does not so derive. A form of content being non-derived is not equivalent to its being *sui generis*: non-derived content can, for example, be derived from, and be explained in terms of, the history or informational-carrying profile of the state that has it. It is *what* content is derived from that is crucial. Non-derived content is content that is not derived from *other* content – it is not content that is irreducible or *sui generis*.⁷

The existence of non-derived content is controversial.⁸ However, there is one very good reason why we should restrict the sort of content possessed by the representational state designated in (3) to non-derived content. The claim that non-derived content is central to cognition has been used to attack *EM*. For example, A&A attack C&C's claim that the sentences in Otto's notebook constitute a subset of his beliefs precisely on the grounds that these sentences possess only derived content.⁹ Therefore, without the restriction to non-derived content, the proposed criterion would be regarded as question-begging by the critics of *EM*. If one is chary about this sort of deployment of the concept of non-derived content, then one should recall that the proposed criterion is intended only as providing a *sufficient* condition for cognition. Accordingly, the dialectical strategy pursued in this paper is to give the critic of *EM* everything he or she could reasonably require—and then show that this is still compatible with *EM*.

Condition (4): I shall argue that anything that is to count as a cognitive process must belong to some or other representational subject (broadly construed). There are no un-owned cognitive processes. Understanding the sense in which cognitive processes have an owner is, I think, one of the hardest tasks in understanding the nature of cognition. It would not be possible to undertake that task in this paper. Here my goals are more modest. I shall argue that understanding the sense in which a cognitive process is owned by a subject is just as problematic for internalists about cognition as it is for defenders of *EM*. A problem for both is not a problem specifically for one. While this 'solution' is not ideal; it is the only strategy permitted within the constraints imposed by a paper of this length.

I shall, for reasons that will become clear, motivate and defend conditions (1)–(3) separately from condition (4).

6. Defending the Criterion: Cognitive-Scientific Practice

I shall defend conditions (1)–(3) of the criterion by showing that they can be extracted, in a relatively straightforward manner, from examination of cognitive-scientific practice. The guiding thought is that if we want to identify a mark of the cognitive, then we had better pay close attention to the sorts of processes that cognitive scientists regard as cognitive, and then try and identify the general features of processes of these kinds. However, to avoid the charge that the criterion is motivated by *EM*-aforethought, the cognitive-scientific practice in question must be

internalist cognitive science—and the more paradigmatic or archetypal the form of internalism the better. Therefore, I am going to focus on David Marr's (1979) theory of vision. While many of the details of this theory are now starting to look a little quaint, the general approach adopted by Marr has dominated internalist-inspired theorizing in cognitive science.¹⁰

For Marr, visual perception begins with the formation of an informationally impoverished retinal image. The function of properly perceptual processing (which is cognitive with a small 'c' but not with a large 'C') is to transform this retinal image into, successively, the raw primal sketch, the full primal sketch, and the 2½D sketch—the culmination of properly perceptual processing. At each stage in the operation, one information-bearing structure is transformed into another. The retinal image, reputedly, contains very little information, but it does contain some. The retinal image is made up of a distribution of light intensity values across the retina. Since the distribution of intensity values is nomically dependent on the way in which light is reflected by the physical structures that the organism is viewing, the image carries some information about these structures. Then this information-bearing structure is transformed into the raw primal sketch. In the raw primal sketch, information about the edges and textures of objects has been added. Application of various grouping principles (e.g., proximity, similarity, common fate, good continuation, closure, and so on) to the raw primal sketch results in the identification of larger structures, boundaries and regions. This more refined representation is the full primal sketch. And so on.

Abstracting from the details, a very definite picture of visual perception emerges. First of all, perception involves information processing: the transformation of information-bearing structures (condition 1). The retinal image is transformed into the raw primal sketch. Further processing operations then transform this into the full primal sketch, and so on. The result of these transformations is the making available of information, to subsequent processing operations, that was previously unavailable (condition 2). Thus, in the transformation of the retinal image into the raw primal sketch new information becomes available for subsequent processing—information that was not present in the retinal image. And in the transformation of the raw primal sketch into the full primal sketch, further information becomes available, information that was not available in the raw sketch. The culmination of the perceptual process is the 2½D sketch. This sketch carries information that is available for further processing operations—the post-perceptual operations that result in the formation of *3-D object representations* (which are, in turn, available to play a further role in belief formation, etc). The general picture is clear: at each stage in the operation, it is possible to identify a new structure, one which carries novel information that is available to subsequent processing operations. Marr's theory, thus, provides a graphic illustration of (2).

Each identified stage in the operation culminates in a new representational state (condition 3). The retinal image, while impoverished, does carry some information about the environment. The goal of visual processing is to successively transform this into an item sufficiently informationally rich in content to provide the basis of visual

perception and post-perceptual judgments. Each stage of the process, therefore, culminates in a state that carries more information about the environment than its predecessor. Once we leave retinal image behind, each successive state is *normative* in at least the following, minimal, sense: given that the state is instantiated with the properties it has, the world is *supposed* to be a certain way. With the retinal image, there is no distinction between the way the world is and the way it is supposed to be: the retinal image is caused by whatever causes it. However, the raw primal sketch contains new information—information contributed by the first stage of perceptual processing. This, in effect, is the brain’s ‘guess’ about the way the world would have to be in order to have produced the retinal image being processed. As such, given the ‘guess’, the world is supposed to be a certain way; and if it is not the ‘guess’ was mistaken. At each successive stage of processing, therefore, we find a state that carries information and makes normative claims on the world. That is, we find basic representational states. Moreover, the content they carry as representational states does not derive from the content of representational states that lie outside the particular processing stream. That is, while the content of the 2½D sketch derives from that of the full primal sketch which, in turn, derives from that of the raw primal sketch, the content involved in the successive transformations that constitute this processing stream does not derive from the content of representational states that lie outside this stream. The content embodied in this particular processing stream is, therefore, *non-derived* content.

7. Extending Cognition

With this general model in mind, consider one extended account of cognition: my (1999) extended model of memory. I structure my extended account of remembering around four principles, two of which are pertinent to our concerns. These are:

- (A) An organism can process information relevant to memory task T through the manipulation of physical structures in its environment. (Rowlands, 1999, p. 122)
- (B) In certain circumstances, acting upon, or manipulating, external structures is a form of information processing. (Rowlands, 1999, 123)

In suitably fleshing out these principles we find, I shall argue, an implicit commitment to the criterion of the cognitive advanced above.

Thus, consider one of Rowlands’ principal examples: our reliance on external information storage structures in the constitution of memory. The Peruvian *kvinu* officer who employs a system of knots to store information employs his biological memory in a very different way from that of the envoy of a culture in which external forms of information storage have not been invented (Rowlands, 1999, pp. 134–6). The latter must rely on episodic memory to retain information, and must do so afresh for each item of information he needs to retain. But the *kvinu* officer need deploy his memory only in the remembering of the ‘code’ that allows his to tap into the information contained in each knot. Once he does this, a potentially unlimited

amount of information becomes available to him through his abilities to manipulate and exploit such external structures.

The knot is an information-bearing structure that exists outside the skin of the subject that deploys it: it is, in this sense, external to the subject. Deployment of knots can take various forms: one can tie them, one can modify them, one can read them, and one can use the information they contain as an aid in the construction of further knots. These are all ways in which the knots can, in the relevant sense, be manipulated or exploited. When one ties a knot, for example, or modifies a knot that one has already been tied (in order, for example, register some change in pertinent information), then one is, I argue, manipulating or transforming an information-bearing structure. Thus, the deployment of knots by the *kvinu* officer satisfies condition (1) of the criterion of the cognitive.

The result of this manipulation or transformation of knots is the making available to the cognitive subject of information that was previously unavailable. Indeed, not only does the knot do this: it is its proper function to do this. We are, of course, dealing with a case of remembering, rather than the transformation of novel information to a distinct individual. So, the relevant scenario would look something like this. The person who ties the knot would otherwise have forgotten the information that the knot contains, due to, let us suppose, other demands on his episodic memory resources. But when he picks the knot up again—the next day, for example—the information it contains is once again available to him. In this case, the knot has the proper function of making available to the subject information that would have been, without the tying of the knot, unavailable. It thus satisfies condition (2) of the proposed criterion of the cognitive.

The way in which this process of manipulating or exploiting knots makes information available to the subject is in the form of the production in that subject of a representational state: a perception of the knot, and subsequent doxastic representations of the informational content contained therein. It is, of course, no part of *EM* to claim that processes entirely external to (i.e., outside the skin of) a cognizing subject can count as cognitive. For *EM*, cognitive processes are entirely internal or are coupled wholes composed of operations occurring both inside and outside the subject's skin. That is, according to *EM*, cognitive processes always contain a non-eliminable internal element. It is here that we find representational states that possess non-derived content. Manipulation of the knot, an external information-bearing structure, thus makes information available to the subject by way of its production in that subject of representational states that possess non-derived content. Therefore, this manipulation of an external information-bearing structure satisfies condition (3) of the proposed criterion of the cognitive.

Therefore these kinds of manipulation, exploitation and transformation of external information-bearing structures satisfy conditions (1)–(3) of the criterion of the cognitive in precisely the same way as the manipulation of internal information-bearing structures described by Marr for the process of perceiving. Therefore, if (1)–(3) do indeed partly delineate what it is for a process to qualify as cognitive, then with respect to these conditions at least, external operations of the sort invoked by *EM*

seem to qualify as cognitive in the same way that classical internal operations qualify as cognitive.

It is obviously not possible, in a paper of this sort, to discuss exhaustively all possible internalist and extended models of various cognitive processes to establish that they satisfy conditions (1)–(3). Indeed, it may well be that not all versions of *EM* can satisfy these conditions. However, I hope to have shown that *EM* can at least make a *prima facie* case for the idea that the sorts of processes it claims are cognitive are cognitive precisely because they satisfy a plausible criterion of the cognitive. This is a criterion of the cognitive implicated in paradigmatic examples of internalist cognitive theorizing. And all of the principal objections to *EM* are based on its purported failure to satisfy such a criterion.

However, to complete this case, there is one further condition that needs to be discussed: the *ownership* condition.

8. Ownership and the Problem of Bloat

Suppose I am using a telescope.¹¹ The telescope is, let us suppose, a *reflector*, and therefore works by transforming one mirror image into another. Mirror images are information-bearing structures—their properties are systematically determined, by a mapping function determined by the specific properties of the mirror and the properties of the visual environment. Therefore, the operation of the telescope is based on the transformation of information bearing structures, and so it satisfies condition (1). The processes occurring inside the telescope are, in combination with other processes, of the sort normally capable of yielding a representational state. This is true even when the content of this state is non-derived. Thus, in combination with other processes—ones occurring inside *me*—the processes can yield a representational state; for example, my visual perception of Saturn’s rings. This is a representational state with non-derived content. Therefore, the processes occurring inside the telescope satisfy condition (3). And the proper function of the processes occurring inside the telescope is making information available, both to me and to subsequent processing operations within me (for example, processes of inference), of information that was previously unavailable (for example, the current orientation relative to earth of Saturn’s rings), and thus satisfy condition (2). Therefore, the processes occurring inside the telescope satisfy conditions (1)–(3) of our criterion of cognition. If these conditions were sufficient for cognition, then the intra-telescopic processes would have to be classified as cognitive.

Relevantly similar examples can be easily generated. How can we rule out, for example, processes occurring inside my calculator, or my computer, from counting as cognitive? They, (1), seem to involve the transformation of information bearing structures. They are also, (3), the sort of process that, when combined with other processes, can produce representational states. Thus, when combined with operations occurring inside my brain, the processes occurring inside the calculator can produce a representational state in me—for example, when I read off the result of their

operations from the screen. And, (2), the proper function of these processes is, it seems, to make available to me information to me, information that I might subsequently employ in further processing operations, where this information was not available prior to the operations of the calculator or computer. Therefore, the processes occurring inside the calculator and computer seem to satisfy conditions (1)–(3), and, without further constraints, would therefore count as cognitive.

This, in essence, is the problem of *cognitive bloat*; but this time it is applied to cognitive processes rather than, as in the case of *C&C*, cognitive states. Cognitive bloat is a common and important objection to *EM*. However, for our purposes, what is even more important is the intuition upon which the problem is predicated. The intuition is that there are no *subjectless* cognitive processes: they always have an *owner*, and this owner is an *individual* of some form.¹² Without this intuition, the defender of *EM* could simply reply: the intra-telescopic, and intra-computer processes are, indeed, cognitive ones.

If this is true, then *EM* has an obligation to demonstrate the sense in which the extended processes it invokes—manipulation, transformation, and exploitation of relevant environmental structures—can be owned, or in an appropriate way belong to, a cognitive subject. That is, it needs to be able demarcate external processes that properly belong to a subject from those—like those occurring inside a telescope or computer—that do not. Conditions (1)–(3) of the proposed criterion are inadequate for this purpose. Conditions (1)–(3) of our criterion of the cognitive concern what we might call the *why* and the *how* aspects of cognition: how cognitive processes do what they are supposed to do, and why they do it. The *why* of cognition—its function, broadly understood—is to make previously unavailable information available, either to subsequent processing operations or to the cognizing subject itself. They do this—the *how* of cognition—by *normally*, and perhaps in combination with other processes of the same general sort, producing representational states. And the production of such a state is the result of information processing operations—the manipulation and transformation of information-bearing structures.

These conditions, however, say nothing at all about the *who* of cognition. In addition to the *why* and the *how* aspects of cognition, there is also the *who* aspect. Whatever else is true of cognitive processes, whatever the specific details of their form and function, such processes always belong to someone or something. This is why we need condition (4).

The problem of bloat is, therefore, a problem because of antecedent commitment to the ownership condition. In particular, the ownership condition both explains why we have the intuition that intra-telescopic processes are not cognitive, and, more importantly, justifies this intuition. This is why the criterion of cognition includes condition (4).

Explaining the sense in which a subject owns his, her, or perhaps its, cognitive processes is, I think, far more difficult than it seems. It is not the sort of project that can be addressed in a paper of this length. Accordingly, in the rest of the paper I want to pursue only a more modest aim. I shall argue that explaining ownership of cognitive processes is just as much of a problem for internalist accounts of cognition

as it is for extended accounts. There is, therefore, no *special* problem faced by extended accounts in accounting for the ownership of cognitive processes. Therefore, the presence of condition (4) in the criterion of cognition would count against extended accounts no more than it does against internalist models.

Traditionally, cognitive processes are thought of as spatially contained within the boundaries of a cognizing subject. Extended cognitive processes are not thus contained. Therefore, one might suppose that the traditional conception is at a distinct advantage over the extended alternative with respect to explaining ownership of such processes: we explain their ownership by a subject in terms of their spatial containment within that subject. A cognitive process *P* belongs to subject *S* just in case *P* occurs inside *S*. However, this, I shall argue, will not work. Indeed, I doubt that spatial containment is plausible criterion of ownership for *any* of the primary bodily processes we undergo; *a fortiori* this is true of cognitive processes

To see this, consider, as an example of a non-cognitive biological process, *digestion*. It may seem obvious that what makes a digestive process mine, and not anyone else's, is the fact that it occurs inside of me and not anyone else. This claim, however, should be resisted: spatial containment is only a fallible guide to the ownership of digestive processes. To begin with, it is possible to imagine a case whereby one's digestive processes become *externalized*. Suppose, for example, one cannot produce enough of the relevant enzymes in one's digestive tract. The solution, drastic and implausible but nonetheless a solution, is to reroute one's tract into an external device where the relevant enzymes are added, before routing the tract back into one's body where it finishes its work in the usual way. The most natural way of understanding this scenario is, I think, as a case where *my* digestive processes pass outside my body and receive the required external aid. The processes do not stop being mine just because they are, for a time, located outside my body. Underlying this intuition is (i) the idea that digestive processes are defined by way of their proper function, and (ii) that what makes a digestive process mine is that that it fulfills this proper function with respect to *me*. Thus, the proper function of digestive processes is to break down food and release it into the body in the form of energy. And, roughly, a digestive process is *mine* if it breaks down food *I* have ingested and releases energy into *my* body.

If this is correct, then the specific character of the external device is largely irrelevant—as long as it permits this proper function to be realized. Suppose, for example, that the external device were the body of someone else. That is, suppose my digestive tract were temporarily rerouted through the tract of someone else. We will suppose that the food contents of each tract are kept separate, but that the coupling allows the other person's digestive enzymes to pass over into my tract, thus aiding in the digestion of food that I ingested. Afterwards, my tract passes back into my body where it culminates in the usual manner. This would appear to be a case where *my* digestive processes are spatially located inside someone else's body, and indeed make use of someone else's digestive enzymes. And, once again, what underwrites this intuition is the idea that a digestive process is mine is the proper function of digestion being fulfilled with respect to *me*: it is *my* food that is being broken down,

and energy is being released into *my* body. The connection between ownership and location of a digestive process is merely contingent.

Ownership of digestive processes seems to be determined not by spatial containment but by a certain kind of *integration*. For a digestive process to be mine, it is necessary and sufficient for it to be integrated into my other biological processes in the right way. For example, it is integrated into my *ingestive* processes to the extent that it consists in the breaking down of food that I have taken in. And it is integrated into my other *respiratory* processes to the extent that it releases energy that enables those processes to continue. The criterion of appropriate integration is determined by proper function: a digestive process is appropriately integrated into my other biological processes when it is fulfilling its proper function with respect to those processes.

This, I think, provides the right model for understanding ownership of cognitive processes. Ownership is to be understood in terms of the appropriate sort of integration into the life—and in particular, the psychological life—of a subject. However, in the case of cognitive processes, specification of the appropriate sort of integration will be complicated by the possibility of a distinction that has no real echo in the case of biological processes such as digestion: the distinction between *personal* and *sub-personal* cognitive processes.¹³ Very roughly, this is the distinction between processes that are conscious, or under the conscious control of the subject, and those that are not. This distinction is reflected in condition (2) of the criterion of cognition. My suspicion, although this is not something I can defend here, nor is it something upon which the arguments developed here depend, is that ownership of sub-personal cognitive processes will prove to be derivative upon ownership of personal level cognitive processes: for cognitive processes, integration is ultimately integration into the conscious life of a subject.¹⁴

However, whether or not this turns out to be the case, the important point is that there is no reason for thinking that explaining the relevant notion of integration will be any more problematic for *EM* than it will be for traditional internalist accounts. That is, there is no reason for supposing that a subject's use of external resources—its manipulation, transformation, and exploitation of external information bearing-structures—will be any more difficult to integrate into that subject's overall psychological economy than will be cognitive processes in the traditional internal sense. The only reason for supposing that there will be an additional problem for *EM* is residual commitment to the containment criterion of ownership, and that criterion is, I think, untenable.

9. Conclusion

I have argued that all of the principal objections to *EM* reduce, in one way or another, to the mark of the cognitive objection. However, far from falling foul of any plausible criterion of the cognitive, I have argued that *EM* is compatible with a criterion of the cognitive that emerges in a fairly straightforward way from analysis of standard

internalist cognitive-scientific practice. Indeed, not only is *EM* compatible with such a criterion, this criterion actually seems to be implicit in at least one influential development of *EM*. There is every reason to suppose that *EM* satisfies conditions (1)–(3) of the criterion of the cognitive proposed here. And there is no reason for supposing that it will have any more difficulty satisfying condition (4) than will internalist accounts of cognitive processes.

Notes

- [1] In the context of *EM*, the terms ‘internal’ and ‘external’ have engendered a surprising amount of confusion. Internal processes are ones occurring inside or at the spatial boundary of the organism (typically, the skin). External processes are ones located spatially outside this boundary. *EM* should not, of course, be understood as the claim that structures and processes entirely external to the cognizing organism could, in certain circumstances, count as cognitive. The claim is that structures and processes external to the organism could, when appropriately coupled with internal cognitive structures and processes, partly constitute cognitive processes. For *EM* cognitive processes are always either internal or hybrid combinations of internal and external. They are never purely external.
- [2] Understood deductively, of course, the argument would be a version of the *modus tollendo ponens* fallacy.
- [3] I am grateful for conversations with Andy Clark for this point.
- [4] See, for example, Heidegger (1926).
- [5] One should not overlook the deeply mysterious character of the concept of information that lies at the heart of this tradition in cognitive science. Shannon’s theory, notoriously, only provides a *criterion* for when one item carries information about another. It does not say what information *is*.
- [6] Adherence to such conventions might be taken to involve some kind of intentional action on the part of the agent, and therefore also invoke the content of other representational states, merely at one step removed.
- [7] In their invocation of non-derived content, *A&A* sometimes refer to it as ‘intrinsic content’. This is a deeply unfortunate locution, even more misleading than ‘non-derived content’. No content is intrinsic. As Dretske once quipped, one might as well talk of an *intrinsic grandmother*.
- [8] See, for example, Dennett (1987).
- [9] Adams and Aizawa (2001).
- [10] Marr’s theory has also been picked because of its utter familiarity. The rationale for extracting an implicit mark of the cognitive from cognitive-scientific practice is, like Plato’s idea of *anamnesis*, to convince you of something you already (implicitly) know.
- [11] My thanks to a conversation with Richard Samuels for this example.
- [12] To say that they their owner is an individual is, of course, *not* to say, necessarily, that this is a person.
- [13] In the case of digestive processes, it is undoubtedly possible to draw a distinction between personal and sub-personal processes. But this amounts to nothing more than the distinction between the digestive process as a whole, and its constituent parts. Digestion, as a whole, is something the organism does. The various components of digestion—peristalsis, the release of enzymes, etc—are processes performed by sub-systems of the organism. This is a legitimate distinction, although it may be difficult to apply with precision in particular cases. However, as we shall see, it is not the same as the personal/sub-personal distinction as this applied to cognitive processes.
- [14] This sort of claim is defended by McDowell (1994).

References

- Adams, F., & Aizawa, K. (2001). The bounds of cognition. *Philosophical Psychology*, 14, 43–64.
- Adams, F., & Aizawa, K. (2009). Defending the bounds of cognition. In R. Menary (Ed.), *The extended mind*. Cambridge, MA: MIT Press.
- Clark, A. (1997). *Being-there: Putting brain, body and world back together again*. Cambridge, MA: MIT Press.
- Clark, A. (2008). *Supersizing the mind*. New York: Oxford University Press.
- Clark, A., & Chalmers, D. (1998). The extended mind. *Analysis*, 58, 7–19.
- Donald, M. (1991). *Origins of the modern mind*. Cambridge, MA: Harvard University Press.
- Dretske, F. (1981). *Knowledge and the flow of information*. Cambridge, MA: MIT Press.
- Heidegger, M. (1962). *Being and Time* (Macquarie & Robinson, Trans.). Oxford: Blackwell. (Original work published in 1926).
- Hurley, S. (1998). *Consciousness in action*. Cambridge, MA: Harvard University Press.
- Lloyd, D. (1989). *Simple minds*. Cambridge, MA: MIT Press.
- Marr, D. (1979). *Vision*. San Francisco, CA: W.H. Freeman.
- McDowell, J. (1994). The content of perceptual experience. *Philosophical Quarterly*, 94, 190–205.
- Menary, R. (2006). Attacking the bounds of cognition. *Philosophical Psychology*, 19, 329–44.
- Menary, R. (2007). *Cognitive integration: Attacking the bounds of cognition*. Basingstoke: Palgrave.
- Menary, R. (2009). *The extended mind*. Cambridge, MA: MIT Press.
- Millikan, R. (1984). *Language, thought and other biological categories*. Cambridge, MA: MIT Press.
- Millikan, R. (1993). *White queen psychology, and other essays for alicia*. Cambridge, MA: MIT Press.
- Noë, A. (2004). *Action in perception*. Cambridge, MA: MIT Press.
- O'Regan, K., & Noë, A. (2001). A sensorimotor account of vision and visual consciousness. *Behavioral and Brain Sciences*, 24(5), 939–60.
- Rowlands, M. (1995). Against methodological solipsism: The ecological approach. *Philosophical Psychology*, 8, 5–24.
- Rowlands, M. (1999). *The body in mind: Understanding cognitive processes*. Cambridge: Cambridge University Press.
- Rowlands, M. (2003). *Externalism*. London: Acumen.
- Rowlands, M. (2006). *Body language: Representation in action*. Cambridge, MA: MIT Press.
- Rupert, R. (2004). Challenges to the hypothesis of extended cognition. *Journal of Philosophy*, 101, 389–428.
- Sutton, J. (2006). Exograms and interdisciplinarity: History, the extended mind, and the civilizing process. In R. Menary (2008b).
- Wilson, R. (2004). *Boundaries of the mind*. New York: Cambridge University Press.