Review Essay

Systems, Functions, and Intrinsic Natures: On Adams and Aizawa's *The Bounds of Cognition*

Robert D. Rupert

The Bounds of Cognition
Frederick Adams & Kenneth Aizawa
Oxford, England: Wiley-Blackwell, 2008
216 pages, ISBN: 1405149140 (hbk): $74.95

Correspondence to: Robert D. Rupert, Department of Philosophy, Campus Box 232, University of Colorado at Boulder, Boulder, CO

1. Introduction and Overview

Where is human cognition located? Is human cognitive processing literally constituted (at least partly) by non-neural portions of the environment? The contemporary debate about extended cognition and the extended mind focuses on these questions, among others. Frederick Adams and Kenneth Aizawa's new book, *The Bounds of Cognition* (*BC*), contributes wonderfully to this debate. The book is critical of the extended approach; but Adams and Aizawa (A&A) also work toward a positive view, one that allows, in principle, for extended cognition, while yielding very little of it when fed the empirical facts. In this review essay, I describe *BC*'s main lines of argument, raise some critical points, and contrast A&A's critique of the extended view with my own (Rupert 2004, 2006, 2008, forthcoming *a*, *b*). I agree with much of what A&A say; nevertheless, I contend that some of A&A's criticisms of the extended position, and of the arguments for it, are inconclusive and that adoption of the systems-based view I endorse bolsters A&A's critical conclusions significantly. This approach, however, entails the rejection of a

central element in A&A's positive proposal; but A&A should, for independent reasons, want to abandon the claim in question—or so I argue.

According to the extended view, cognitive processing takes place in the environment beyond the boundary of the organisms (humans, primarily) to which we normally attribute cognitive capacities. A stock example involves the use of the pencil and paper to solve complicated math problems (Rumelhart, Smolensky, McClelland, & Hinton 1986). In such a case, the organism interacts in an ongoing way with the external resources; the organism accomplishes something it could not accomplish without the use of external resources; and the external resources play the kind of role in problem-solving we would expect internal resources to play (cf. Clark and Chalmers 1998). For these reasons, and sometimes further ones, it is held that pencil and paper are literal components of the cognitive process, having the same cognitive status as neural resources are typically thought to have.

The hypothesis of extended cognition, as it's sometimes called, can be formulated in a variety of non-equivalent ways; differences in scope (all human cognition? most human cognition? all actual cognition) and modal strength (all possible cognition? some possible cognition?) differentiate these from each other. A&A address some subtleties of this sort, but their most important contributions, and those that take up the lion's share of their critical discussion, pertain to more fundamental issues. A&A argue that, across the board, advocates for the extended view appeal to the wrong kinds of consideration; at the very least, standard arguments for the extended position are so radically incomplete that, absent substantial amendments, these arguments do not support the extended position to any appreciable degree.

A&A's critique of the extended view takes the following general shape: First A&A develop a charge of incompleteness (Chapters 2, returned to in Chapter 5). Advocates for the

extended view need a theory of cognition; the author who asserts that cognition extends into the environment had better be prepared to tell the rest of us what it is that extends into the environment. Moreover, this theory of cognition had better be independently plausible. Simply saying that cognition is whatever contributes causally to intelligent behavior, for example, does not advance the discussion; it is rather an exercise in relabeling, with no scientific payoff (pp. 73, 85—all page references are to *BC* unless otherwise noted). Typically, though, proponents of the extended view offer only a sketchy, poorly supported, or implausibly liberal theory of the cognitive. Second, A&A offer their own view (Chapters 3 and 4): cognition essentially involves the application of distinctive kinds of neural processes to nonderived representations. Finally, A&A survey a variety of arguments that have been given for the extended thesis (Chapters 6-9), criticizing them in various ways: for leaving gaps where a theory of cognition should be, for resting on implicit premises about cognition that seem independently (sometimes wildly) implausible, for running afoul of A&A's own theory of cognition without offering anything in its place (these criticisms often blend into one another, for understandable reasons), or for resting on an inference to the best explanation when an equally good (or perhaps better) nonextended explanation of the same data is available. In the following section, I say a bit more about each piece.

2. Theories of Cognition and Arguments for the Extended View

The charge of incompleteness seems to me to be serious and essentially on target. Proponents of the extended view often rest their case on observations about dependence—causal, dynamical, evolutionary, metaphysical, or epistemic. Adams and Aizawa cite numerous passages of this sort (pp. 89-91 and *passim*) from a variety of authors supportive of the extended view. In response to such claims, one might reasonably wonder why *that* constitutes cognition? Why does A's

depending on B confer on B the sorts of properties we normally associate with A? Now, it is not as if proponents of the extended view offer no basis at all for their conclusions. Many of the passages cited by A&A explicitly invoke dependence of one kind or another. Missing, though, is a well-supported, theoretical account of cognition in light of which observations about the various forms of dependence are germane.

Thus, the incompleteness charge should not be understood as a mere complaint about form. After all, we can easily render valid any argument for the extended view simply by adding an appropriate premise; if nothing else, add a conditional premise of the form "if the premises originally given are true, then the extended view is true." A&A's concern cannot be allayed with such ease, however; A&A would like advocates for the extended view to make explicit *and argue convincingly for* the notion of cognition in play. To the extent that lack of explicit-*ness* is a serious charge, it is because, when one's theory of cognition goes unstated, readers—and perhaps authors themselves—find it easier to overlook the kind question suggest above: Why does *that* determine where cognition is? It does not suffice to assert, for example, that a network of causes and effects are all cognitive simply because the collective causal process produces intelligent behavior. One must also argue that any part of a network of causes and effects that produces intelligent behavior is cognitive. By pressing this demand forcefully, A&A both motivate the introduction of their own theory of cognition and prepare the reader for their criticisms of arguments for the extended view.

Having laid out the charge of incompleteness, A&A offer what they take to be a plausible account of the cognitive—a mark of the cognitive, in their terms (Adams and Aizawa 2001). According to A&A, cognitive processes are certain kinds of operation (more on which below) over nonderived representations. As A&A see things, the content of mental representations is not

merely as-if content or content in the eyes of the beholder: mental representations do not inherit their content from an outside agent or interpreter. Rather, mental representations carry externalist content in virtue of a privileged, objective relation that holds between mental representations and the objects, properties, or kinds to which mental representations are so related. A&A do not endorse a specific naturalistic semantic theory, but, to illustrate what they have in mind, they briefly survey some of the leading contenders (including Dretske's [1988] teleo-indicational view, Fodor's [1987, 1990] asymmetric dependence view, and Cummins's [1996] isomorphism-based view).

I agree with A&A that nonderived representations play a central role in cognitive processing, but it requires some finesse to shape this claim into a necessary condition for cognitive processing. As A&A concede, it is not that every aspect of every cognitive process consists solely in the manipulation of nonderived representations. Perhaps things such as parentheses or, more generally, punctuation in the language of thought play a constitutive role in cognitive processing (p. 50). (This issue is more substantive than the terminology of 'punctuation' may make it sound; many in the extended camp, broadly construed, view the entire notion of representation with suspicion—see, e.g., Thelen and Smith 1994, Keijzer 1998.) In response, A&A suggest a functional-role approach to the individuation of such units (p. 55). This, however, seems to play into the hands of the opposition; after all, if individuation is functional, why not think some external states play the functional role of parentheses or punctuation? A&A might respond by appealing to the nature of distinctively cognitive processes. So, let us move to that topic.

What about cognitive processes? A&A hold that cognitive processing is distinguished by its ("largely unknown"—p. 84) intrinsic character, that is, by the intrinsic physical nature of the

mechanisms or the states involved. A&A's argument for this view proceeds partly by analogy (pp. 58-60, 125). In other sciences, nuclear chemistry, for example, a process's kind is determined by the physical nature of the reactions, particles, or elements constituting that process. Returning to the case at hand, we know that at least some significant portion of cognition takes place in the brain. Given that the brain differs greatly in its intrinsic chemical and biological properties from such things as pencils and sheets paper, we should expect there to be a distinctive study of cognition, as a neural phenomenon. We should expect this enterprise to discover distinctive intrinsic physical properties and mechanisms and thus to deliver laws of (neuro)cognitive psychology that simply do not hold of the external things.

Here A&A's reasoning seems to me to be suspect. Cognitive modeling frequently takes place at a high level of abstraction. Simon Baron-Cohen's (1995) theory of cognitive development vis-à-vis autism, or Shaun Nichols and Stephen Stich's (2003) hybrid theory of the theory of mind, or John Anderson's (2007) most recent version of ACT-R are, first and foremost, functional-role characterizations—boxology with more or less explained in the way of mechanistic implementation. To the extent that neural considerations come into play (as they do), the neural goings-on influence formulation of, or confirm existing, proposals pitched at the functional level. The neural goings-on do not come into play in virtue of their intrinsic physical natures.

The very formulation of my concern suggests a natural way for A&A to reply. They might embrace a kind of fine-grained-functionalism, that is, accept functionalist cognitive psychology but hold that the very functions to be identified have intricate structure, the nature of which is best discovered by examining the fine-grained causal patterns in the materials of the brain (cf. pp. 69, 136). On the view I am suggesting, psychological properties are functional-role

properties, but the functions involved are, structurally speaking, so complex that only a very narrow range of materials, organized in intricate and fragile ways, stands much of chance of realizing these properties (cf. Shapiro 2004). Fine-grained functionalism thus can square what sounds like A&A's opposition to the multiple realization of cognitive kinds (p. 68) with functionalist theorizing in cognitive psychology (and cognitive science, more generally).

This tack, however, invites the following riposte. The advocate for the extended view will almost certainly claim that the external materials in question themselves instantiate fine-grained functional properties. Looking then at the neurally realized, fine-grained functional properties, the proponent of the extended position will ask, "Why is *that* cognition and not *this* [now pointing to the relevant external materials and their functional properties]?" When faced with this rejoinder, A&A seem to concede too much: they embrace a distinctive science of neural cognition, which may or may not be part of some larger cognitive science (pp. 132, 140). In doing so, they seem to abandon the mark of cognitive that was meant to support their anti-extended position; by abandoning distinctively neural processing as a necessary condition on cognitive processing, they allow that the external materials may well *be* cognitive. If the properly cognitive process includes the brain and environmental materials, then A&A's preservation of a distinctive science of the neurally instantiated portion of the cognitive process strikes me as a pyrrhic victory.

There is a way out of this jam, more on which below; but first, a few qualifications. Note that, even if my criticisms of A&A's positive view make their mark, this should provide little comfort to the extended theorist. I don't think A&A discuss any argument *for* the extended view to which their only powerful objection is that the argument's premises conflict with A&A's positive account of cognition. Moreover, A&A's emphasis on the distinctive causal profile of

human neural material—and the fine-grained functions it is suited to realize—nicely ground criticisms of what they call 'cognitive equivalence' arguments for the extended view: arguments appealing to the supposed functional equivalence in the contribution of the inner and external materials (pp. 133-43).

Let us move on, then, to reasons that have been offered in support of the extended position, the focus of *BC*'s second half. Here A&A develop cogent objections to the leading arguments for the extended view. One such argument rests on evolutionary considerations (Rowlands 1999; cf. Beer 2003, Wheeler 2005): evolved mechanisms are specially effective in the kind of environment in which they evolved; thus, the environment is part of the processes distinctive of those mechanisms (p. 147). Why, A&A ask, should facts about the evolutionary environment entail facts about the location of cognition? Why doesn't the detailing of causal interactions with structures in the evolutionary environment merely show how it is that the neural cognitive processor makes use of its noncognitive surroundings? Why doesn't the evolutionary history simply explain how it is that one particular kind of neural system for cognitive processing appeared in the natural world? A&A treat a variety of arguments—covering topics from dynamical systems and robotics to Haugeland's claims about intimacy and high-bandwidth connection—in a similarly effective manner.

One criticism, for which A&A are well known, concerns the causal-constitution (or coupling-constitution) fallacy (Chapters 6 and 7). Many arguments for the extended view appeal primarily to the subject's causal interaction with the environment, perhaps of an ongoing, reciprocal sort that amounts to organism-world coupling (van Gelder 1995, Clark 1997), during problem-solving. Why, though, are all of the contributors to such causal processes cognitive? Why not model these processes instead as causal interactions between the brain-bound cognitive

system and its environment? To the extent that authors fail to address this question effectively, they commit the causal-constitution or the coupling-constitution fallacy.

One incarnation of the fallacy deserves special attention in the present context. This style of argument assumes a liberal view of cognitive *systems*, finds that cognitive systems frequently have components in the environment beyond the boundary of the organism, and concludes that cognition extends into the environment. A&A draw attention to the missing premise, *viz.* that cognitive processing (equivalent to cognition) permeates every part of a cognitive system. Our scientifically informed reactions to analogous cases speak against this premise, however. Only certain parts of an air conditioning *system* actually condition the air. The ducts carrying the air certainly do not condition the air. Thus, in general, not every portion of an *X*-system does *X* (pp. 107, 117-118, 125). Why, then, should we think that cognition permeates a cognitive system, especially if we accept a very liberal standard for something's being part of a system?

A&A's fundamental point is incontrovertible; so long as cognition amounts to something other than all activity in a cognitive system, it is possible that some process be part of the activity of a cognitive system without being genuinely cognitive. Nevertheless, one should wonder whether this possibility is actualized in the cases central to the debate over extended cognition. A&A offer two kinds of consideration in support of their skeptical position. First, A&A rely on their theory of cognitive processes, appealing in particular to intrinsic physical natures. Assume it is necessary that a cognitive process exhibit certain neural properties or fine-grained functional properties that only neural matter is likely to realize. Noncognitive parts of a cognitive system lack these properties. In the actual world, these bits are the organismically external parts of the relevant cognitive systems, as well, most likely, as whatever non-neural parts of the human body partly constitute such cognitive systems.

I find this way of fleshing out A&A's concern implausible, for reasons discussed in connection with A&A's theory of cognition. Notwithstanding A&A's repeated invocation of cognitive psychology (pp. 57, 60, 68), the reference to intrinsic physical character seems out of step with cognitive psychological practice. Groups of neurons are treated as cognitively important because of what they do, because of the way in which they causally contribute to the production of intelligent behavior, not because of their intrinsic character. Thus, we should expect cognitive psychology to deliver a functional criterion distinguishing the genuinely cognitive from the merely causal contributors to cognitively driven forms of behavior; and it is not clear what functional criterion will do the job for A&A. It does not suffice to observe that the fine-grained functional properties of the brain are different from the fine-grained functional properties of external stuff. What makes the former cognitive and the latter not? As noted above, the differences at issue help to defuse cognitive equivalence arguments for the extended view (as well as natural- or explanatory kinds based arguments [Rupert 2004]), but they do not effectively ground a positive, nonextended view of cognition.

Second, A&A frequently take aim at the low standard sometimes set for something's being cognitive (pp. 76, 86) and some of their remarks about cognitive systems (pp. 106, 130) suggest a similar criticism might apply to the extended theorist's inference from systems and processes. The underlying concern seems to be that cognitive systems come too cheaply if one sets such a low bar, in the sense that there is no scientific value to individuating cognitive systems (or cognition) in this liberal way. This criticism will strike the extended theorist as somewhat weak, though. It erects no principled barrier to an extended approach; it merely demands that the extended cognition research program put up or shut up, so to speak, and I suspect that advocates for the extended view take themselves to be pursuing the former

alternative (the putting up). A&A concede that their opponents might be offering what amount to competing empirical hypotheses concerning the theoretically relevant system the operation of which explains intelligent behavior (p. 112), but do not engage sufficiently with this possibility. I think this empirically oriented way of casting the debate is well advised. Rather than focusing on the distinction between genuinely cognitive portions of the cognitive system and those that are not, we should try to differentiate the theoretically important system, from the standpoint of successful cognitive psychology, from further causal contributors to the production of intelligent behavior. On a functional approach, identifying the privileged system and its theoretically relevant properties would seem—at a first pass, anyway—to be all there is to identifying cognitive states and processes.

3. The Systems-based Approach

An enormous amount of success in cognitive psychology results from treating human bodies as if they contain some kind of integrated, persisting system that plays a privileged role in the production of intelligent behavior (Rupert 2004, 2008, forthcoming *a*, *b*; cf. Clark 2007 on the hypothesis of organism-centered cognition). This is *not* to say that cognitive psychology has been uniformly successful or that the cognitive system is body-bound *simply because* there has been a significant measure of success in orthodox cognitive science. I am not begging the question against those who would rather focus on orthodox cognitive psychology's shortcomings. Rather, I argue that cognitive psychology has identified one candidate for the genuinely cognitive system and that no other principled candidate is on the table. It serves no purpose to attribute cognitive status to everything that causally contributes in a significant way to the production of the explananda of cognitive science; the sun surely contributes to visually guided cognitive behavior, but I see no point in attributing cognitive status to the sun. The causal

contribution of light has long (always?) been recognized by even the most orthodox of vision theorists, so it is an empty exercise in relabeling to assert, with revolutionary tone, that the sun is a literal component of the cognitive process.

Orthodox cognitive psychology thus grounds a candidate principle for demarcating cognition: a state is cognitive if and only if it involves the activation or operation of a component of the integrated cognitive system (a component that distinctively and casually contributes to the production of a cognitive outcome). Here is not the place to give a detailed account of integration (see Rupert forthcoming *a*), but the basic idea is this: the cognitive system consists in the collection of mechanism or capacities the contributions of which are highly correlated and which (unlike, say, air pressure) contribute distinctively to cognitive outcomes; this is the collection of capacities (or underlying physical mechanisms) each element of which contributes, as a member of overlapping subsets of mechanisms, to a wide range of cognitive outcomes. For convenience, think of the system as constituting a cognitive architecture.

Notice that this way individuating the cognitive system does not exclude modules (Fodor 1983), be there such things, so long as they contribute causally and distinctively to a wide range of cognitive explananda and their contribution is highly correlated with the contributions of other mechanisms of similar standing. Notice, too, that the general form of my argument can proceed without appealing to a detailed account of integration. Psychology offers a candidate system the boundaries of which might demarcate cognition. This is the system the functioning of which is fruitfully probed by placing humans in a wide variety of circumstances and accounting for systematic variations in human behavior in terms of the organismically bounded system's principle-governed interactions with varying aspects of those experimental circumstances. No other plausible candidate for demarcation—where plausibility partly requires the making of a

substantive contribution to cognitive scientific results or methodology—are on the table. The kind of system in question (the obligate system, as Wilson [2002] calls it) is not likely to extend beyond the boundary of the organism. Therefore, no significant portion of human cognition is extended.

In a nutshell, then, I recommend that we treat the issue as a question in the philosophy of cognitive science. What is cognitive? Whatever plays a privileged role in cognitive scientific theories, models, and explanations. Does this amount to everything that causally contributes to the production of intelligent behavior (language use, reasoning, etc.)? No, that is relabeling without progress. Does any proper subset of causal contributors play a privileged role in cognitive science? Yes, the set of integrated mechanisms or capacities—the cognitive architecture or obligate system. What is cognition, then? The relevant states of this integrated, obligate system, and the processes that consist in causally connected series of such states.

What are the advantages of this approach with respect to *BC*'s contents. The systems-based approach resolves the question, raised above, about multiple realization and the extended theorist's likely *tu quoque*. Yes, most things in the universe instantiate fine-grained functional properties. Why are the properties supported by the brain of any special importance? Because they play a role in the integrated, persisting system's production of cognitive behavior. The systems-based view solves the problem of punctuation and parentheses in the same way. A state can be a state of the cognitive system—and thereby instantiate a cognitive property—in virtue of being a state of the persisting, integrated system, even if the state in question does not constitute a nonderived representation.

Moreover, the existence of an integrated, persisting system seems presupposed by most theories of nonderived content. For there to be representations at all on Dretske's view (Dretske

1988), there must be learning; but learning, plausibly enough, amounts to changes within an integrated, persisting system. Similar comments apply to the many other theories of content (Millikan 1984, Cummins 1996, Rupert 1999, Ryder 2004) that require a substantive history or a coherent system in order that content be determinately fixed.

4. Conclusion: Compare and Contrast

Does the systems-based view provide a viable alternative route to A&A's critical conclusion, or does the systems-based approach yield a significant amount of extended cognition after all? The work that can be done by focusing on systems is of no use to A&A if I have bargained away their ultimate conclusion. I suspect that, at least with respect to most humans for most of their lives, the systems-based view offers a friendly verdict: there is little or no cognitive extension into the environment. Most of the external props that appear in examples of supposedly extended cognitive processing (external rotation on the Tetris screen, for example—Clark and Chalmers 1998) are special-purpose devices, not contributing in varied and overlapping arrangements with other mechanisms of similar standing to the production of a wide range of cognitive explananda. Moreover, the successful psychological methodology that motivates the attempt to characterize integration focuses on things (the experimental subjects) with bodily boundaries; so the system searched for, the operation of which should account for interaction with the various *ad hoc* contributors, is not likely to extend beyond the body. Perhaps when the persisting system is successfully characterized, and that characterization is applied to newly arriving empirical data, it will allow for some extended cognition. As things stand, though, the systems-based approach appears to deliver essentially the result A&A would like with regard to what I call 'extended cognition'—cognition beyond the boundary of the body.

I am more likely to part ways with A&A when it comes to embodied cognition—that is, cognitive processing that takes place in the nonneural biological body. Even here, however, a lot depends on unsettled empirical matters. The success of orthodox cognitive psychological research suggests that the obligate cognitive system is within the boundaries of the organism, but of course, the brain is within said boundaries; it travels with the body from experimental set-up to experimental set-up.

Note, too, that the systems-based view could be constrained in principled ways. Consider this compromise between A&A's view and my own, a compromise that may yield a rabidly neural view of cognition: take the cognitive system as I am inclined to characterize it, and identify just the mechanisms that operate on nonderived representations; these, and only these, are cognitive mechanisms or processes. Of course, for all we know about representation, this compromise may yet yield a significant amount of embodied cognition (Rowlands 2006). In the end, I think we should do our best to extract principled distinctions from successful experimental, theoretical, and explanatory practice in cognitive science; and if the resulting principles yields embodied cognition, there will be our answer.

One final point of potential disagreement. A&A criticize (pp. 79-83) what they take to be a latent operationalism in the literature on extended cognition. For example, they criticize Clark for wishing "to explicate the notion of a cognitive process by appeal to cognitive tasks, just the opposite order of our approach" (p. 80). If Clark holds the view that anything contributing to the production of cognitive explananda is thereby cognitive, then I too reject his view as excessively operationalist. Nevertheless, there is a perfectly respectable approach in this vicinity. A science gets going because there appear to be distinctive phenomena in need of explanation. Sometimes, upon further investigation, the processes producing those phenomena turn out to have no

theoretical unity (here we face a kind of eliminativism); in other cases, there turns out to be enough unity to ground a distinct science (with revisions allowed, of course; there is an ongoing interplay between our intuitive categorization of the explananda and theoretical accounts of what produces those phenomena). As regards cognition, I argue that the theoretically unifying construct in cognitive psychology is the obligate system. If this be operationalism, then, with all due respect to A&A, I endorse it; but it is no big bullet to bite, for this form of operationalism is standard scientific methodology and should be embraced by naturalistically minded philosophers.

References

Adams, F., and Aizawa, K. (2001). The bounds of cognition. *Philosophical Psychology*, 14, 43–64.

Anderson, J.R. (2007). *How can the human mind occur in the physical universe?* Oxford: Oxford University Press.

Baron-Cohen, S. (1995). *Mindblindness: An essay on autism and theory of mind*. Oxford: Oxford University Press.

Beer, R. (2003). The dynamics of active categorical perception in an evolved model agent. *Adaptive Behavior*, 11, 209–43.

Clark A. (1997). *Being there: Putting brain, body, and world together again*. Cambridge, MA: MIT Press.

Clark, A. (2007). Curing cognitive hiccups: A defense of the extended mind. *Journal of Philosophy* , 104, 163–192.

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

Cummins, R. (1996). *Representations, targets, and attitudes*. Cambridge, MA: MIT Press.

Dretske, F. (1988). *Explaining behavior: Reasons in a world of causes*. Cambridge, MA: MIT Press.

Fodor, J. (1983). *The modularity of mind.* Cambridge, MA: MIT Press.

Fodor, J. (1987). *Psychosemantics: The problem of meaning in the philosophy of mind*. Cambridge: MIT Press.

Fodor, J. (1990). *A theory of content and other essays*. Cambridge, MA: MIT Press.

Keijzer, F. (1998). Doing without representations which specify what to do. *Philosophical Psychology*, 11, 269–302.

Millikan, R.G. (1984). *Language, thought, and other biological categories: New foundations for realism*. Cambridge, MA: MIT Press.

Nichols, S., and Stich, S. (2003). *Mindreading: An integrated account of pretence, self-awareness, and understanding other minds*. Oxford: Oxford University Press.

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

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

Rumelhart, D.E., McClelland, J.L., Smolensky, P., Hinton, G.E. (1986). Schemata and sequential thought processes in PDP models. In D.E. Rumelhart, J.L. McClelland, and the PDP Research Group (Eds.), *Parallel distributed processing: Explorations in the microstructure of Cognition*, Vol. II., *Psychological and biological models* (pp. 7–58). Cambridge, MA: MIT Press.

Rupert, R.D. (1999). The best test theory of extension: First principle(s). *Mind & Language*, 14, 321–55

Rupert, R.D. (2004). Challenges to the hypothesis of extended cognition. *Journal of Philosophy*, 101, 389–428.

Rupert, R.D. (2006). Review of Raymond Gibbs, *Embodiment and Cognitive Science* (Cambridge UP, 2006). *Notre Dame Philosophical Reviews*, 2006.08.20 http://ndpr.nd.edu/review.cfm?id=7443

Rupert, R.D. (2008). Innateness and the situated mind. In P. Robbins and M. Aydede (Eds.), *Cambridge handbook of situated cognition* (pp. 96–116). Cambridge: Cambridge University Press.

Rupert, R.D. forthcoming *a*. *Cognitive systems and the extended mind*. New York: Oxford University Press.

Rupert, R.D. forthcoming *b*. Representation in extended cognitive systems: Does the scaffolding of language extend the mind? To appear in R. Menary (Ed.), *The extended mind*. Cambridge, MA: MIT Press.

Ryder, D. (2004). SINBAD Neurosemantics: A theory of mental representation. *Mind & Language*, 19, 211-40.

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

Thelen, E., and Smith, L.B. (1994). A dynamic systems approach to the development of cognition and action. Cambridge, MA: MIT Press.

van Gelder, T. (1995). What might cognition be, if not computation? *Journal of Philosophy*, 92, 345–81.

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

Wilson, M. (2002). Six views of embodied cognition. *Psychonomic Bulletin & Review*, 9, 625–36.