Toward a Situated, Embodied Realism[1]

Anthony Chemero
tony.chemero@fandm.edu

Center for the Ecological Study of Perception and Action
University of Connecticut
Storrs, CT 06269 USA

and

Scientific and Philosophical Studies of Mind Program
Franklin and Marshall College
Lancaster, PA 17604 USA

Abstract: Situated, embodied cognitive science is all the rage these days. Some (including the present author) have argued that situated, embodied cognitive science is incompatible with realism (metaphysical and scientific). In this paper, I argue that this is a mistake: there is no reason one cannot be both a proponent of situated, embodied cognitive science and a realist. To show this, I point to flaws in two previous arguments against realism. I also recommend a slightly modified version of Hacking's entity realism as an especially good fit for situated, embodied cognitive science.

Word Count: 8946

[1] Acknowledgements: Thanks to Colin Klein, Jared Julius and Adam Kovach for comments on an earlier draft of this paper. This research was supported by National Science Foundation grant #00-04097.

1. Introduction. Suppose everyone you know buys an old, Italian scooter. They seem really cool riding around and talking about carburetors while wiping their hands on their greasy pants. Pretty soon you decide you want one too, ideally a mid-1960s Lambretta. But then someone tells you how much those putty, two-stroke engines pollute, and you feel torn. You want to look cool, but you don't want to ruin the air more than it already is ruined. You just don't know what to do. Well, if you're in this situation, your quandary is oddly parallel to the one cognitive scientists find themselves in nowadays. Everybody, it seems, is embracing situated, embedded, embodied cognitive science. They're talking about robots and multi-limb coordination and Heidegger. Plus, they're getting papers and books published. But you've also heard that situated, embodied cognitive science leads to problems with realism, and the last thing you want to worry about is metaphysics

The purpose of this paper is to tell you that everything is OK. Not for the Lambretta—you probably shouldn't get one of those—but for the cognitive science. You can, despite a series of arguments to the contrary (Varela, Thompson and Rosch 1991, Cussins 1992, Smith 1996, Chemero 1998), join the situated, embodied in-crowd and still be a metaphysical and scientific realist. Most of my telling you that everything is OK will be showing that arguments that situated, embodied cognitive science is incompatible with realism shouldn't be trusted. The problem with that argument comes to light most obviously if one re-plays a bit of the debate between Jerry Fodor (1984, 1988) and Paul Churchland (1988) that took place in this journal. So I'll do that too. Here, then, is the plan. In section 2, I'll describe situated, embodied cognitive science. This has been done many, many times lately, so I'll be fairly brief. Then, in section 3, I'll go through two arguments that purport to show that situated, embodied cognitive science is incompatible with realism. In section 4, drawing on Churchland and Fodor, I'll show situated embodied cognitive scientists can be a realists after all. Finally, in section 5, I'll recommend a particular variety of realism, a slight modification of Hacking's entity realism (Hacking 1982, 1983), as especially appropriate for embodied cognitive science.

2. Situated, Embodied Cognitive Science. As noted above, nowadays most cognitive science is or claims to be embodied or situated[2] or both. Anything else seems so last century. Situated, embodied (SE) cognition has made its way from the cognitive sciences into marketing (Rosa and Malter 1993), nursing (Paley 2004), and economics (Boucekkine, del Rio and Licandro 2004), among many other fields. For detailed discussions and manifestoes, see Brooks 1991, 1999;

[2] The word 'embedded' is often used to describe this variety of cognitive science. It is roughly a synonym of 'situated'. (See Smith 1999.)

Varela, Thompson and Rosch 1991; Smith 1991, 1996; Wheeler 1996, to appear; van Gelder 1995; Clark 1997, 1999, 2003; Agre 1997; Clancey 1997; Lakoff and Johnson 1999; Pfeiffer and Scheier 1999; Keijzer 2001; Dourish 2001; Anderson 2003.[3] Typically, those writing about situated, embodied cognition start with the early work of Rodney Brooks (1991). To give credit where it is due, I will go back further, to the work of James Gibson (1979) and the collaborations between John Barwise and John Perry (1981, 1983).

Gibson's ecological theory of vision (1979) was intended as a direct response to the increasing dominance of computational theories of mind, according to which perception and thought are rule-governed manipulations of internal representations. Gibson's ecological approach to perception has three major tenets. First, perception is direct, which is to say that it does not involve computation or mental representations. That is, Gibson thought that perception was not a matter of internally adding information to sensations. Second, perception is primarily for the guidance of action, and not for action-neutral information gathering. We perceive the environment in order to do things. The third tenet follows from the first two. Because perception is does not involve mental addition of information to stimuli, yet is able to guide behavior adaptively, all the information necessary for guiding adaptive behavior must be available in the environment to be perceived. Thus the third tenet of Gibson's ecological approach is that perception is of *affordances*, i.e., directly-perceivable, environmental opportunities for behavior. Affordances, as Gibson was well aware, are ontologically peculiar:

> [A]n affordance is neither an objective property nor a subjective property; or it is both if you like. An affordance cuts across the dichotomy of subjective-objective and helps us to understand its inadequacy. It is equally a fact of the environment and a fact of behavior. It is both physical and psychical, yet neither. An affordance points both ways, to the environment and to the observer. (1979, 129)

Despite this ontological peculiarity and the controversy over how to best understand affordances (Turvey 1992, Reed 1996, Chemero 2003, Scarantino 2004), the idea of affordances—divorced of their relation to direct perception—is the one aspect of Gibson's theory that gained significant attention from the beginning, e.g., from designers (see Norman 1986). The rest of Gibson's ideas were not widely accepted by cognitive scientists upon their appearance. They were, however, widely discussed (see Fodor and Pylyshyn 1981; Turvey, Shaw, Mace and Reed 1981; Ullman 1981), and did attract a small, solid-core of devotees. More recently, Gibson has

[3] Dear Reader: if your work is cited in this list, feel free to skip ahead to section 3. You won't miss anything you don't already know.

become one of the heroes of SE cognitive science, which has adopted these views (slightly softened) as its own.

Moving slightly closer to the present, we can trace the origins of the *situated* portion of SE cognition to situation semantics, the work in the philosophy of mind and language initiated in the 1980s by John Barwise and John Perry (Barwise and Perry 1981, 1983). Taking themselves to be providing a semantics for Gibsonian psychology,[4] Barwise and Perry argued that we can't understand meaning or cognition without taking into account that thinkers are spatially located (i.e., situated) and so have only incomplete, locally available information at their disposal. Every thinker and speaker is someone, who is somewhere, and is aware of only certain things. One major upshot of this is that indexicals move from the periphery of accounts of cognition to the center. The idea is that because we are situated in the environment, thoughts about "here", "there", "now", "me" are ubiquitous. This focus on indexicals, we will see, is a crucial but almost incidental feature of the embodied movement in cognitive science. A second important feature of Barwise and Perry's was derived directly from its Gibsonian motivation. Barwise and Perry developed their situation semantics in order to account for meaning without reference to mental representations. In their non-representational account, having meaningful thoughts (perceptions, utterances) has nothing to do with having mental representations, or indeed with anything that might be called epistemic. The meaning of thoughts and sentences are a matter of the relationship between thinkers/speakers and information in their environments.

It is this latter aspect of Barwise and Perry's situation semantics that Rodney Brooks (1991) picks out when he uses the word "situated" to describe his robots. When Brooks says that his simple, mobile robots are situated, he means that, because they are in the midst of a changing world, they do not need to use representations of the world to plan or guide their behavior. Instead they interact with the world itself. The idea is that there is no need to store information on board, and make predictions about how things will change during an action, when you can just act and check again. Brooks sums this up this anti-representationalism with the slogan "The world is its own best model." This Gibson-like[5] skepticism about mental representations is perhaps the most (in)famous aspect of Brooks's early work, but it is not his anti-representationalism that makes Brooks the model for situated, embodied cognitive science. Instead it is his insistence that intelligence is necessarily embodied. Brooks argues that it is real interaction with the real world, not cleverly manipulating representations, that is the mark of

[4] See footnote 1 of Barwise and Perry 1981.
[5] Perhaps inadvertently so. Brooks does not cite Gibson, but Kirsh (1991) argues convincingly that the theory of action Brooks outlines just is Gibson's.

intelligence. In effect, Brooks sees Barwise and Perry and raises them: for Brooks it is not just a thinker's setting, but also its physical constitution, that is essential for understanding it as intelligent, thinking, etc. And, of course, since having a physical constitution is essential to intelligent behavior, so too is being situated in a physical (not to mention social) environment. Embodied cognition is always situated.

The current work in SE cognition arises from these sources (among others, of course) is a broad-based movement, incorporating work in robotics (Beer 1995, 2003; Harvey, Husbands and Cliff 1994; Husbands, Harvey and Cliff 1995**;** Nolfi and Floreano 2000; Slocum et al 2000), simulated evolution (Nolfi and Floreano 2000**;** Husbands et al 1995; Harvey et al 1994), developmental psychology (Thelen and Smith 1994, Thelen 1995), perception (Turvey and Carello 2004; Turvey 1996; Bingham 1995; Bingham, Bradley, Bailey and Vinner 2001), motor control (Kelso 1995, Turvey 1990), cognitive artifacts (Hutchins 1995, Dourish 2001), and, of course, theoretical manifestoes (see above). Given this variety of subject matter, there is also variety in theoretical approach. The following tenets, though, are more or less universally held among SE cognitive scientists.

Skepticism about mental representations. Although with few exceptions (Keijzer 2001, Chemero 2000), SE cognitive scientists reject the strongest claims made by Brooks, Gibson and Barwise and Perry about mental representations, they agree that a good deal of the behavior of humans and other animals does not involve mental representations. Most intelligent behavior, indeed most cognition, is real-time interaction with the environment in which the world can serve as its own model. Furthermore, it is agreed that when representations are involved in cognition or behavioral control, the representations are not the sentence-like representations of classical cognitive science. Instead, they are indexical, context-dependent, action-oriented, etc. Typically, they are either representations of affordances (Clark 1997) or emulations of ongoing actions (Grush 1997).

Intelligent Bodies, Scaffolded Environments, Fuzzy Borders. Given this minimization of importance of complex, internal mentation, it is a challenge to explain complex behavior. In SE cognitive science, some of the intelligence is "off-loaded" from the brain to the body and environment. On this view, our bodies are well-designed tools, making them easy for our brains to control. For example, our kneecaps limit the degrees of motion possible with our legs, making balance and locomotion much easier. It is only a small exaggeration to say that learning to walk is easy for humans because our legs already know how. (See Thelen and Smith 1994, Thelen 1995.) This off-loading goes beyond the boundaries of our skin. The natural environment is already

rich with affordances and information that can guide behavior. As when beavers build dams, in interacting with and altering the environment, animals enhance these affordances. Kirsh and Maglio (1994; see also Kirsh 1995) show that manipulating the environment is often an aid to problem-solving. Their justly-famous example is of Tetris players rotating zoids on-screen, saving themselves a complicated mental rotation. Hutchins (1995) shows that social structures and well-designed tools allow humans to easily accomplish tasks that would otherwise be too complex. This leads many to believe that cognitive systems are not confined to the brain or body, but include aspects of the environment (Clark 1997). Clark (2003) takes this further, arguing that external tools (including phones, computers, language, etc.) are so crucial to human life that we are literally cyborgs, partly constituted by technologies.

Interactive Explanation and Dynamical Systems. Taking cognitive systems to include aspects of the body and environment requires an explanatory tool that can span the agent-environment border. Many situated, embodied cognitive scientists have turned to dynamical systems theory (van Gelder 1995; Port and van Gelder 1995; Kelso 1995). That is, many (though not all) proponents of situated, embodied cognitive science take thinkers to be dynamical systems, best explained using the tools of dynamical systems theory. Dynamical systems theory is especially appropriate for explaining cognition as interaction with the environment because single dynamical systems can have parameters on each side of the skin. That is, we might explain the behavior of the agent in its environment over time as coupled dynamical systems, using something like the following equations, from Beer (1995, 1999):

$$\dot{x}_A = \mathsf{A}(x_A; S(x_E))$$
$$\dot{x}_E = \mathsf{E}(x_E; M(x_A))$$

where A and E are continuous-time dynamical systems, modeling the organism and its environment, respectively, and $S(x_E)$ and $M(x_A)$ are coupling functions from environmental variables to organismic parameters and from organismic variables to environmental parameters, respectively. It is only for convenience (and from habit) that we think of the organism and environment as separate; in fact, they are best thought of as comprising just one system, U. Rather than describing the way external (and internal) factors cause changes in the organism's behavior, such a model would explain the way U, the system as a whole, unfolds over time. It is also worth pointing that dynamical systems theory is neutral over whether to consider parameter or variable values as representations. Some dynamical modelers are representationalists (e.g., Petitot 1995), some are not (e.g., Beer 1995).

3. Situated, Embodied Cognition and Realism. Ever since Rodney Brooks kick-started the situated, embodied cognition movement, there have been arguments that the approach is somehow inconsistent with realism (Varela et al 1991; Cussins 1992; Smith 1996; Chemero 1998). Unlike the three tenets just discussed, the inference from SE cognition to irrealism is highly controversial even within the movement. Many proponents of SE cognition are steadfast realists (Mandik and Clark 2002) or don't want to be bothered by metaphysical concerns in the first place (Clark 1997). Yet others have argued that having cognition necessarily embodied, and hence limited, profoundly affects our abilities to know about and interact with the mind-independent world. Here I will look at two arguments from SE cognition to irrealism: Varela et al's argument in *The Embodied Mind* and the one in Chemero 1998. Later, in section 4, I will argue that SE cognitive scientists can be realists. Before doing any of this, though, I should point out that the unmodified word 'realism' has so many uses as to be almost empty. By realism, I mean the coupled claims that (1) at least some of perceptions/thoughts/theories are accurate, and (2) that the objects of our accurate perceptions/thoughts/theories exist in a mind-independent world.

Varela et al 1991 take Brooks's robots to be a paradigmatic case of their enactive approach to cognition. Brooks's robots, Varela et al say, are *structurally coupled* to the environment, which is to say that robot and environment are not separate. Instead, they are mutually specifying in that the robot's behavioral abilities and sensory systems determine what its environment is, which in turn determine what the robot does. So, for example, the robot Allen has a ring of 12 ultrasonic sensors which it uses to determine the distance to the nearest object at each "hour" around it's body. With just these 12 sensors, Allen can wander around most cluttered environments successfully. The only things that can perturb Allen—that is, influence its behavior—are sufficiently large things that reflect ultraviolet light. According to Varela et al only these things are part of Allen's world. This makes Allen closed in an important sense: only very particular stimuli can elicit a reaction from Allen, and the way Allen reacts determines the significance of those stimuli. This is what it is for Allen to enact, or "bring forth" a world. Varela et al take Allen as a model for all animals. All animals are closed as Allen is, structurally coupled to a world composed of very specific stimuli; all animals enact or bring forth a world that is determined by the nature of their sensorimotor systems, which in turn determine the significance of the perturbations.

This sounds suspiciously like old-fashioned idealism. Animals have their own worlds, determined by what they sense. But Varela et al reject this interpretation.

> It is precisely this emphasis on mutual specification that enables us to negotiate a middle path between the Scylla of cognition as the recovery of a pregiven outer world (realism) and the Charybdis of cognition as the projection of a pregiven inner world (idealism). These two extremes both take representation as their central notion: in the first case representation is used to recover what is outer; in the second it is used to project what is inner. Our intention is to by pass entirely this logical geography of inner versus outer by studying cognition not as projection or recovery but as embodied action.

By rejecting representationalism, Varela et al claim to stake out a position that is neither realist nor idealist. On their view, animals and worlds are not separate, so there is no need for animals to represent the world. Without representations, there is nothing besides the world for the animal to interact with, but the worlds animals—including humans, including philosophers and physicists—interact with is strictly limited, and determined by sensorimotor capabilities.

Andy Clark, toward the end of a book that cheerfully outlines and defends SE cognitive science, dismisses this as a mere distraction.

> Varela et al. use their reflections as evidence against realist and objectivist views of the world. I deliberately avoid this extension, which runs the risk of obscuring the scientific value of an embodied, embedded approach by linking it to the problematic idea that objects are not independent of the mind. My claim, in contrast, is simply that the aspects of real-world structure which biological brains represent will often be tightly geared to specific needs and sensorimotor capacities. (Clark, 1997, 173)

This dismissal, which is Clark's only comment on the issue, is far too glib. We might try to get a handle on Clark's reasoning by noting that he also dismisses Varela et al's anti-representationalism (less glibly, with argument even). If animals are representing, we can ask what they are representing, whether what they are representing exists, and whether they are representing it accurately. We can also wonder whether there is anything beyond the representations. So, to repeat a mixed metaphor, Clark's representationalism puts us off Varela et al's middle path and back into Odysseus's boat. Once there, Clark chooses Scylla (realism) because, well, it's not a distraction. As much as one might argue with the philosophical probity of such a move, it is easy to see its point. Clark, more perhaps than anyone else, has worked to define SE cognitive science as a modest revolution, one that treats both babies (realism, representation) and bathwater (Cartesianism) appropriately.

It is, alas, not so easy. Chemero (1998) argued that Clark's representationalism doesn't save situated, embodied cognitive science from worries about realism. This papragraph and the next are a paraphrase of Chemero's argment. We can see this by considering that the representations

involved in SE cognition are, to use Clark's term, action-oriented. As Clark describes them, action-oriented representations (AORs) are representations that *both* describe a situation *and* suggest an appropriate reaction to it; they are essentially representations of affordances. AORs are doubly indexical, in that they are both local and personal: they are local in that they relate to the circumstances currently surrounding an animal; they are personal in that they are related to the animal's needs and the skills that it has. With AORs as the basis of human cognition and, quite probably, the only representations available to most non-humans, it is a few small steps to worries about realism.

Start by realizing that humans and other animals' action-oriented representations will concern only the actions they undertake. Animals, that is, will represent only affordances for animals like themselves. And given the differences in the activities of animals of different species, we should expect the affordances perceived by animals of different species, or animals of the same species at different developmental stages, to be widely divergent, and even contradictory. (Indeed, this is a central point of the idea of affordances. See Michaels and Carello 1981.) This is to be expected if one assumes, as SE cognitive scientists do, that perceptual systems evolved to guide behavior. Neither humans nor beetles have AORs that represent the mind-independent world exactly correctly. Indeed, representing the mind-independent world is not what AORs are supposed to do; they are supposed to guide action. So the set of human affordances, i.e., action-oriented represent*ed*s, is just as tightly geared to specific human needs and sensorimotor capacities as those of any other type of animal. This leaves us with a multiplicity of conflicting sensorimotor systems, each of which is appropriate for guiding the adaptive behavior of animals whose systems they are.

This multiplicity of differing, conflicting, sets of action-oriented representeds is a problem for realism. Because there is no reason to assume that there is any criterion for appropriateness of AORs other than appropriately guiding behavior, there is no reason to think that humans, but not beetles, have AORs that reflect mind-independent reality. This becomes a problem when one adds the premise, common among proponents of SE cognitive science, that distinctively human "higher thought" (language use, theorizing, science, etc.) is based upon AORs. Indeed, many (Deacon 1997, Clark 1997, Millikan 1998) have suggested that public language evolved to fit facts about our pre-linguistic brains. So, because AORs are tightly geared to needs and sensorimotor capacities and hence not accurate reflections of the mind-independent reality, so too will be language-based higher thought. If both our everyday perceptual categories and the categories of our sciences are built upon a skewed foundation, they too will be skewed. Essentially, our AORs don't map the mind-independent world because they aren't supposed to;

our theories don't map the mind-independent world because they are built from our AORs. We must, then, reject one part of realism as defined above. Although, we are justified in believing that there is a mind-independent external world, we have no justification to believe that our perceptions, thoughts and theories are accurate reflections of it.

This leaves us able to believe in what Putnam (1985) calls *internal realism* at best.[6] In internal realism, one believes the entities of a theory to be real, as long as one understands them as theory bound, and not elements of an independent reality. And though not everyone agrees as to whether internal realism is a kind of transcendental idealism (Putnam 1985, Millikan 1993), or a kind of nominalism (Hacking 1983), everyone does agree that it is not a kind of realism.

So… Varela et al argue that non-representational, situated, embodied cognitive science is inconsistent with realism. Chemero argued that representational, situated, embodied cognitive science is inconsistent with realism. It would seem then that controversies Clark wanted to sweep under the rug are leaving an unseemly lump. Must situated, embodied cognitive scientists worry give up on realism?

4. The Joys of Plasticity. In this section, I will argue that, for reasons specific to SE cognitive science, the argument[7] just presented, despite their initial plausibility, are not to be believed. (See Mandik and Clark 2002 for a different response to these arguments.) To make this case, I will *re*-present an argument made by Paul Churchland in a somewhat different context. The context in question is the debate between Churchland (1988) and Jerry Fodor (1984, 1988) over the theory-ladenness of perception (Hanson 1958; Kuhn 1962; Churchland 1979) that took place in this journal. Some background is appropriate. Though the debate over theory-ladenness is most directly about objectivity, whether perception is theory-laden is directly relevant to realism. If perception is theory-laden, the theories we believe affect what we perceive. How, then, can theory-laden perception can be of theory-independent reality? Kuhn (1962), for example, has argued that holders of different theories perceive different (at least partly theory-determined) worlds. Attempting to head these worries off, Fodor (1984) argues that, because perception is modular, perception is not theory-laden. Perceptual mechanisms, Fodor argues

[6] This is unsurprising considering that Chemero's argument is parallel to Putnam's. Putnam argued that because ideal scientific theories have many models that conflict with one another, there is no reason to trust that any one gets the theory-independent world just right. Chemero argued that because we actually do have many equally good, conflicting sensory systems, there is no reason to trust that one maps the animal-independent world just right.

[7] I am putting Varela et al's non-representationalist argument against realism aside for the time being. I will bring it back in section 5.

(1983, 1984), are modular, i.e., innately structured and informationally encapsulated. Because perception is informationally encapsulated, theories (not handled by perceptual modules) do not change perceptual mechanisms or the output of perceptual modules. Therefore, Fodor argues, perception is not theory-laden. Furthermore, because perceptual modules are innately structured, they're the same in all (normal) humans. There is, then, no sense in which humans who believe different theories perceive a different world. Their perceptual mechanisms produce the same output given the same input. Although they may hold different theories about what they perceive, they perceive the same thing. So, Fodor holds, the modularity of mind makes perception a neutral basis for theoretical disputes, and this can form the basis for objectivity and scientific realism.[8]

In a reply to Fodor, Churchland (1988) argues that the modularity of perceptual mechanisms does not get Fodor what he wants, which is *theory-neutral* perception. A perception is theory-neutral "just in case its truth is not contingent upon the truth of any general empirical assumptions, just in case it is free of potentially problematic assumptions." (1988, 170). Innate perceptual mechanisms are not free of empirical assumptions; their empirical assumptions are simply hard-wired by evolution. Churchland points out that having innate, informationally-encapsulated perceptual models "merely dooms us to a single point of view, a point of view that is epistemologically just as problematic as the infinity of other sets of empirical assumptions that might have been hard-wired into us instead." (1988, 1970) Furthermore, Churchland continues, because these mechanisms are hard-wired and encapsulated, we are stuck with them, no matter how faulty we might learn that they are. Far from getting us theory-neutral perception, that is, modularity gets us instead "universal dogmatism" in which all humans have perceptual systems that share the same, probably false, but unshakeable assumptions. Swinging this around to meet up more directly with the arguments of section 3 above, what Churchland shows is that Fodor's argument is not an argument for theory-neutrality at all. Rather, we end up exactly where section 3's arguments left us, with evolutionarily-hard-wired, hence action-oriented, hence anthropomorphically-biased, empirical assumptions built into our perceptual systems. As argued above. this leads to grave doubts that our perceptions are of a human-independent world, to internal realism at-best. Indeed, some evolutionary psychologists have reached the same conclusions from their devotion to innately structured, evolved, perceptual modules: both Plotkin (1993) and Boyer (2001) embrace non-realist conclusions.

[8] Churchland (1979) embraces both scientific realism and theory-laden perception. But Churchland's scientific realism is highly non-standard.

Although Churchland's argument—not surprisingly—rolls right off Fodor's back (see Fodor 1988), it does lead to worries about the argument from situated, embodied cognitive science. First, Fodor, Plotkin and Boyer are not good company for situated, embodied cognitive scientists. If SE cognitive scientists are agreeing with these guys, something must be terribly wrong. And something has gone terribly wrong. We can see what it is by reminding ourselves of something that Fodor, Plotkin, and Boyer agree upon, in opposition to Churchland and—as far as I know—all SE cognitive scientists. Fodor and evolutionary psychologists are *nativists*. In contrast, situated, embodied cognitive scientists follow Churchland in fully embracing neural, perceptual, behavioral, and conceptual *plasticity*. And embracing plasticity is fatal to the argument outlined in section 3. To see this, we can continue re-playing Churchland's arguments (1988). Churchland points out that the uniformity of innate perceptual mechanisms (and universal dogmatism) only lasts until a mutant shows up, with a different but equally good (or better) perceptual mechanism.

> In fact, we begin to become such mutants or aliens ourselves, when we change our sensory modalities by augmenting them with unusual instruments such as phase–contrast microscopes, deep-sky telescopes, long-baseline stereoscopes, infrared scopes, and so forth. And the metamorphosis is completed when, after years of professional or amateur experience, we begin to see the world appropriately and efficiently with these new senses. (1988, 171)

The point of all these examples of perceptual augmentation is not merely that we can augment our perception. While interesting, mere augmentation doesn't get us out of the argument described in section 3. If the action-oriented foundations for our perceptual systems are crooked, making the building taller won't help. More important for our purposes is the metamorphosis of perceptual systems.[9] Churchland argues that plasticity in our perceptual systems allows them re-organize to take full advantage of the technological extensions he lists.[10]

There is every reason to think that our perceptual systems can be re-organized. First, no matter what you think about SE cognitive science, pick up any neuroscience textbook to see that our brains are exceptionally plastic. Second, for SE cognitive scientists, consider the centrality of cognitive artifacts in cognition. Our brains, behavior and perception must be capable of transformation if we are to take advantage of such artifacts. Indeed, Clark 2003 argues that the

[9] In an extended critique of arguments by Varela et al and Chemero, Mandik and Clark (2002) point to perceptual extensibility, but not to the re-organization that might come with it. Clark 2003, however, is a big fan of cognitive re-organization.

[10] See also Churchland 1979, which is book-length paean to plasticity.

ability to transform to take advantage of perception-transforming technology is the most distinctively human feature of our brains. If this is correct, then there is every reason to think that we can re-align our perceptual foundations, adjust our perceptions to deal with new situations, despite our biased action-oriented representations. Indeed, the AORs themselves can be linguistically, theoretically, and technologically altered, potentially moving closer to alignment with a mind-independent world. So if we assume plasticity in our perceptual systems, we can avoid the non-realist conclusions of section 3.

Before concluding this section, a few caveats are in order. First, my calling on plasticity above is no guarantee of realism. It just shows that Chemero's argument from SE cognitive science to irrealism need not be believed. Second, this is a *local* issue in two senses. It is local in that it only concerns the particular purported connection between embodied cognitive science and the untenability of realism. If you worry about defending scientific realism from general, empiricist arguments, I haven't helped your case here. The point here has been to show that there is no *special* reason to be anti-realist relating to SE cognitive science. The issue is also local in that the help I've offered against the realist only applies in the SE neighborhood in cognitive science, where we believe that cognition is situated and embodied and that perceptual systems are plastic. The points raised in this section concerning plasticity do not affect those who live across town and believe in evolved, innately-structured perceptual modules. If you deny that perceptual systems are highly plastic, as evolutionary psychologists and other nativists do, you probably ought to believe in universal dogmatism and/or some kind of non-realism (internal realism, Varela et al's enaction, etc.). Plotkin and Boyer agree.

5. But what sort of realism? OK, so situated, embodied cognitive scientists can be realists. That is, they can believe that there is a mind independent world, and that some of our perceptions and thoughts get it right. But to say that is to say very little, considering the wide variety of realisms on the market. In this section, I'll outline one variety of realism that is particularly appropriate for SE cognitive science; there are no doubt others that will do the trick. Before saying what kind of realism is appropriate, it's worth commenting on one that isn't appropriate. The argument that allows SE cognitive scientists to be realists is one that Churchland offers in a different context, and depends upon the plasticity of our perceptual systems. Since SE cognitive scientists are also committed to plasticity, it might seem that Churchland's own realism, described in *Scientific Realism and the Plasticity of Mind* (1979), would be just the right kind of realism. Why not, after all, let Churchland solve all our problems for us? Alas, it's not that easy. Churchland's own single-sentence description of his scientific realism makes this plain. "Excellence of theory as the fundamental measure of all ontology." (1979, 2) Because

Churchland thinks *all* knowledge is theoretical, this serves as a means to secure realism about both the unobservable entities that natural scientists talk about and, in those cases where our theories are appropriate, the middle-sized objects that populate our everyday environment. But this is plainly inappropriate for SE cognitive science, which is founded on the claim that most cognition is interaction with the environment, and has nothing to do with theory. Indeed, SE cognitive science is wholeheartedly opposed to the idea that minds are, or are analogous to, theories.[11,12] So Churchland's scientific realism is not the right realism.

SE cognitive science, with its focus on action, needs a realism that moves away from theory to focus on *practice*. Ian Hacking's entity realism (1982, 1983), suitably modified, will do the trick. Hacking's realism moves debates over scientific realism out of the realm of scientific theory, into scientific practice. The existence of theoretical entities, Hacking argues, is secured by our ability to manipulate them during experiment. His primary example is electrons. Experimentalists are justified in believing in electrons because they can build equipment that exploits the properties of electrons to investigate something else (e.g., neutral bosons). Essentially, when electrons become tools in the experimenter's kit, their reality is on a par with all the rest of the experimenter's tools (glassware, computers, etc.). Hacking writes:

> The more we come to understand some of the causal powers of electrons, the more we can build devices that achieve well-understood effects in other parts of nature. By the time that we can use the electron to manipulate other parts of nature in a systematic way, the electron has ceased to be something hypothetical, something inferred. It has ceased to be theoretical and has become experimental. (1983, 262)

Theoretical entities are real when they become part of the practice of the experimentalist. This, Hacking insists, is independent of whether the experimentalist has a true theory of electrons. Two important caveats. First, although Hacking's theory has been the subject of considerable debate (Shapere 1993; Resnik 1994; Reiner and Pierson 1995; Clarke 2001; Massimi 2004), for now I will simply assume that entity realism is tenable and explain how it can be fit to situated, embodied cognitive science. Second, some readers will react to my call upon entity realism here thusly: "That's not what *I* mean by 'realism'". And it is surely true that Hacking's realism is not standard scientific realism. Entity realism, I am about to argue, seems especially

---

[11] Incidentally, this is something Churchland and Fodor agree upon. For Churchland, a mind is one, big theory; for Fodor, the mind is a bunch of little theories. Taking minds to be (analogous to) theories is also a key feature of some theories of development (Gopnik 1996; Gopnik and Meltzoff 1997).

[12] Note that there is no contradiction between claiming that perceptual or cognitive systems are plastic and claiming that they are not theories. For the situated, embodied cognitive scientist, modifications of perceptual systems come from learning new skills, not alterations of theories.

appropriate for SE cognitive science. But, note that the previous section argued that there are no special reasons that SE cognitive science leads to irrealism. So, if you want to be a realist, but entity realism is not for you, there is no in-principle reason that some other form of realism might not also work. That said, onward.

The source of the problem with realism for situated, embodied cognitive science is that at root animals perceive affordances, and affordances are animal-dependent. Given this animal-dependence, in what sense are we justified in taking affordances to be part of the basic furniture of the universe. There are two ways to apply Hacking's entity realism to these particular entities. The first way is to expand Hacking's position from scientific to a more general realism. The second is to stay within Hacking's area of concern, and apply entity realism to affordances as theoretical entities. We will look at these in turn.

Expanding entity realism. Hacking makes very clear that his entity realism is about scientific entities, a position in the debate over the existence of theoretical entities which are observable only with specialized equipment or not at all. (See Hacking 1983, Chapter 7.) He takes for granted the more general realism, about the things we perceive, that was called into question by the arguments outlined in section 3 above. It is easy, though, to see how to expand entity realism to fit affordances and the middle-sized objects of the environment. The key is to realize that experimental practice is just one kind of practice. Just as experimentalists are justified in believing in electrons because they can use them in their investigations, the rest of us are justified in believing in entities that we use in our successful practices. I see no reason to think that this extension is untenable. Such an extension applies immediately and obviously to affordances. According to situated, embodied cognitive scientists, we perceive affordances and guide our actions based upon them. This behavior is successful even though humans are rarely able to report accurately, or at all, on what they perceive that allows them to behave successfully. (What are you perceiving, right now, that allows you to maintain upright posture in your chair as you read? What are you perceiving while driving that allows you to brake successfully, avoiding collision, while coming up to a stop light? Hint: It's not your car's speed. See Owens, Chiang and Muller 1996.) Thus humans most assuredly don't have true theories of the affordances they are using, yet they use them successfully in their practices. According to entity realism, then, affordances are real and we are justified in believing in them. Our perceptions are of affordances, which really exist, even though we can't say anything true about them.

This focus on practice also makes entity realism an appropriate response to Varela et al's non-representationalist irrealism (1991). Varela et al claim that the question of realism versus idealsm only comes up if one is a representationalist. Entity realism, though, does not depend on representations of any kind. It depends not upon what one thinks or theorizes but on what one does. Indeed, like Varela et al, Hacking (1983) is skeptical of the idea that thoughts and sentences are represention. Entity realism holds even for Varela et al's enactive theory of cognition.

Affordances as theoretical entities. Discussing affordances as theoretical entities is, it might seem, a bit odd. Affordances, after all, are what animals are supposed to perceive. If animals perceive affordances, in what sense are they unobservable, theoretical entities. There are two reasons to think of affordances as unobservable. First, the sense of 'observation' relevant to the philosophy of science involves describing, recording and reporting what is perceived. That is, most perception is not observation. (See, e.g., Torretti 1986.) But as noted above, humans are rarely aware that they perceive affordances. Furthermore, even when we are aware that we perceive affordances, we can almost never describe them. (Again: try to describe what you're perceiving that allows you to maintain your posture as you're reading this.) If humans are unable to describe or report on affordances they perceive, affordances are not observable in the sense relevant to the philosophy of science. Feyerabend puts the point this way: "There is no use appealing to observation if one does not know how to describe what one sees, or if one can offer one's description with hesitation only, as if one had just learned the language in which it is formulated." (1975, 59) Second, affordances *per se* do not reflect light and cannot be detected by laboratory instruments. Experimentalist in the lab physical measure properties of animals and environments and record the responses of animals. They then use this data, along with theoretical premises, to infer the presence and qualities of affordances. For example, Warren 1984 measured the leg length of his subjects and their energy consumption while climbing stairs of different heights, and recorded their judgments about which steps were climbable. This information was combined with premises about the nature and perceivability of affordances to determine characteristics of affordances for stair-climbing. Given this, it is appropriate to think of affordances as unobservable, theoretical entities.

In entity realism, one is justified in being a realist about theoretical entities when one can use them as tools in experimental investigations of other entities. To see an example of the this, consider the experiment described in Chemero, Klein and Cordeiro 2003. In that experiment, affordances related to gap-crossing were used to study the perception of events related to gap-crossing. Humans and other animals often have to step, hop, or jump discontinuities or breaks

in the ground. In recent years, a good deal of research has explored the combination of environmental and animal properties that determine whether gap-crossing affordances are present. In particular, experiments have shown that stable, repeatable relationships among anthropomorphic measures (eye height, leg length, flexibility, gait), postures (sitting, standing, walking, running), and environmental properties (gap size, gap depth, ground stability) determine the presence and perceivability of affordances. (See Mark 1987; Burton 1992, 1993, 1994; Jiang and Mark 1994; Mark, Jiang, King and Paasche 1999; Cornus, Montagne and Laurent 1999.) These findings allowed the manipulation of gap-crossing affordances in order to determine whether humans can perceive behaviorally-salient gap-crossing events. In particular, based on what is known about gap-crossing affordances, Chemero et al were able to make gap-crossing affordances appear and disappear in real-time, and measure the way in which subjects responded. They predicted, and found, that subjects would be able to accurately perceive behaviorally-salient events. (See Chemero et al 2003 for details.)

This experiment used affordances as a tool to study events and their perception. The very possibility of such an experiment is justification for realism about affordances. This is true even though there are several competing theories of just what affordances are. Some (e.g., Reed 1996) believe that affordances are manifest properties of objects in the environment; others (e.g., Turvey 1992; Scarantino 2003) believe that affordances are dispositional properties of objects in the environment; others still (Chemero 2003; Stoffregen 2003) believe that affordances are relations between animals and environments. Yet despite this widespread theoretical disagreement, all parties agree about the basic experimentally-determined properties of affordances. Whatever affordances are, they are real and have well-known properties that can be used in experiments. Although, with just one published experiment, affordances are not entrenched to the extent that electrons are, we are justified in believing in affordances.

These two applications of entity realism to affordances show that entity realism works for objects of interest to situated, embodied cognitive scientists, even contested, animal-dependent ones.

6. What SE Realism is like. Refer to BCS; Rosch and Skarda in N&F book

6. Conclusion. There are many reasons to worry about SE cognitive science. Will it be able to account for distinctively human intelligence, like planning a menu or doing calculus? Indeed, will SE cognitive science be able to account for more complicated embodied behavior, such as picking just the right time to cross a busy street? Is SE cognitive science just a passing fad?

These are all legitimate concerns. All I hope to have shown is that, despite what you may have heard, metaphysics is not a reason to worry. That is, we can be both situated, embodied cognitive scientists and realists.

References

Agre, P. (1997). *Computation and Human Experience*. New York: Oxford.

Anderson, M. (2003) Embodied cognition: A field guide. *Artificial Intelligence* 149: 91-130.

Barwise, J. and Perry, J. (1981). Situations and attitudes. *Journal of Philosophy* 77: 668-91.

Barwise, J. and Perry, J. (1983). *Situations and Attitudes*. Cambridge: MIT Press.

Beer, R. (1995) "Computational and Dynamical Languages for Autonomous Agents." In Port and van Gelder 1995.

Beer 2003

Bingham, G. P., Bradley, A., Bailey, M., & Vinner, R. (2001). "Accommodation, occlusion and disparity matching are used to guide reaching: A comparison of actual versus virtual environments." *Journal of Experimental Psychology: Human Perception and Performance* 27: 1314-1334.

Boucekkine R., del Rio F., and Licandro O. (2003). "Embodied Technological Change, Learning-By-Doing And The Productivity Slowdown." *Scandinavian Journal Of Economics* 105: 87-97.

Beer, R. (1995). Computational and dynamical languages for autonomous agents. In *Mind as Motion* (Port, R.F. and van Gelder, T., eds.), pp. 121-147, MIT Press.

Beer, R.D. (2000). "Dynamical approaches to cognitive science." *Trends in Cognitive Sciences* 4, 91-99.

Beer, R.D. (2003). "The Dynamics of active categorical perception in an evolved model agent." *Adaptive Behavior, 11*, 209-243.

Boyer, P. (2001). "Natural epistemology or evolved metaphysics? Developmental evidence for early-developed, intuitive, category-specific, incomplete, and stubborn metaphysical presumptions." *Philosophical Psychology* 13:

Brooks, R. (1991). "Intelligence Without Representation." *Artificial Intelligence*, 47:139-159.

Brooks, R. (1999). *Cambrian Intelligence*. Cambridge: MIT Press.

Burton, G. (1992). Nonvisual judgement of the crossability of path gaps. *Journal of Experimental Psychology: Human Perception and Performance, 18*, 698-713.

Burton, G. (1993). Non-neural extensions of haptic sensitivity. *Ecological Psychology, 5*, 105-124.

Burton, G. (1994). Crossing without vision of path gaps. *Journal of Motor Behavior, 26*, 2, 147-161.

Carello, C. and Turvey, M. (2004) "Physics and Psychology of the Muscle Sense." *Current Directions in Psychological Science* 13: 25-28.

Chemero, A. (1998). "A Stroll through the Worlds of Animats and Persons." *Psyche* 14: 4.

Chemero, A. (2000). "Anti-representationalism and the Dynamical Stance." *Philosophy of Science* 67: 625-647.

Chemero, A. (2003.) "An Outline of a Theory of Affordances." *Ecological Psychology* 15: 181-195

Chemero, A. Klein, C., and Cordeiro, W. (2003). "Events as Changes in the Layout of Affordances." *Ecological Psychology* 14: 19-28.

Churchland, P. (1979). *Scientific Realism and the Plasticity of Mind*. Cambridge: Cambridge University Press.

Churchland, P. (1988). "Perceptual Plasticity and Theoretical Neutrality". *Philosophy of Science* 55: 167-187.

Churchland, P. (1989). *A Neurocomputational Perspective*. Cambridge: MIT Press.

Clancey, W. (1997). *Situated Cognition: On Human Knowledge and Computer Representations.* Cambridge: Cambridge University Press.

Clark, A. (1997). *Being There*. Cambridge: MIT Press.

Clark, A. (1999). "An Embodied Cognitive Science?" *Trends In Cognitive Sciences* 3: 345-350.

Clark, A. (2003). *Natural Born Cyborgs*. New York: Oxford University Press.

Clarke, S. (2001). "Defensible Territory for Entity Realism." *British Journal for the Philosophy of Science* 52: 701-722.

Cornus, S., Montagne, G., & Laurent, M. (1999) Perception of a stepping-across affordance. *Ecological Psychology*, *11*, 249-267.

Cussins, A. (1992). "Content, Embodiment and Objectivity: The Theory of Cognitive Trails". *Mind* 101: 651-688

Deacon, T. (1997). *The Symbolic Species*. New York: WW Norton.

Dourish, P. (2001). *Where the Action Is*. Cambridge: MIT Press.

Feyerabend, P. (1975). *Against Method*. London: New Left Books.

Fodor, J. (1983). *The Modularity of Mind*. Cambridge: MIT Press.

Fodor, J. (1984). "Observation Reconsidered". *Philosophy of Science* 51: 23-43.

Fodor, J. (1988). "A Reply to Churchland's 'Perceptual Plasticity and Theoretical Neutrality'". *Philosophy of Science* 55: 188-198.

Fodor, J. and Pylyshyn, Z. (1981). "How direct is visual perception? Some reflection on Gibson's 'ecological approach'" *Cognition*: 9, 139-196.

Gibson, J. (1979). *The Ecological Approach to Visual Perception*. New Jersey: Houghton Mifflin.

Gopnik, A. (1996). "The Scientist as Child." *Philosophy of Science* 63: 485-514.

Gopnik, A. and Meltzoff, A. (1997). *Words, Thoughts and Theories*. Cambridge: MIT Press.

Grush, R. (1997). "The Architecture of Representation" *Philosophical Psychology*, 10: 5-24.

Hacking, I. (1982). "Experimentation and Scientific Realism." *Philosophical Topics* 13: 71-87.

Hacking, I. (1983). *Representing and Intervening*. Cambridge: Cambridge University Press.

Hanson, N. (1958). *Patterns of Discovery*. Cambridge: Cambridge University Press.

Harvey, I. (1992). "Untimed and Misrepresented: Connectionism and the Computer Metaphor. U. of Sussex Tech Report.

Harvey, I., Husbands, P., and Cliff, D. (1994). "Seeing the Light: Artificial Evolution, Real Vision." In D. Cliff (ed.) *From Animals to Animats 3*.

Hendriks-Jansen, H. (1996). *Catching Ourselves in the Act*. Cambridge: MIT Press.

Husbands, P., Harvey, I., and Cliff, D. (1995). "Circle in the Round: State Space Attractor for Evolved Sight Robots. *Journal of Robotics and Autonomous Systems*, 15: 83-106.

Hutchins, E. (1995). *Cognition in the Wild*. Cambridge: MIT Press.

Keijzer, F. (2001). *Representation and Behavior*. Cambridge: MIT Press.

Kelso, J. (1995). *Dynamic Patterns*. Cambridge: MIT Press.

Kirsh, D. (1991). "Today the earwig, tomorrow man." *Artificial Intelligence* 47: 161-184.

Kirsh, D. (1995). "The Intelligent Use of Space". *Artificial Intelligence* 72: 31-68.

Kirsh, D. and Maglio, P. (1994). "On Distinguishing Epistemic from Pragmatic Action". *Cognitive Science* 18: 513-549.

Kuhn, T. (1962). *The Structure of Scientific Revolutions*. Chicago: University of Chicago Press.

Jiang, Y and Mark, L. (1994). The effect of gap depth on the perception of whether a gap is crossable. *Perception and Psychophysics*, *56*, 6, 691-700.

Lakoff, G. and M. Johnson. (1999). *Philosophy In The Flesh*. New York: Basic Books.
Mandik, P. and Clark, A. (2002). "Selective Representing and World Making." *Minds and Machines* 12: 383-395.

Mark, L. (1987). Eye height-scale information about affordances: A study of sitting and stair climbing. *Journal of Experimental Psychology: Human Perception and Performance, 13*, 360-370.

Mark, L., Jiang, Y., King, S. S., and Paasche, J. (1999). The impact of visual exploration on judgements of whether a gap is crossable. *Journal of Experimental Psychology: Human Perception and Performance, 25*, 287-295.

Michaels, C. and Carello, C. (1981). *Direct Perception*. Englewood Cliffs, NJ: Prentice-Hall.

Millikan, R. (1993). *White Queen Psychology*. Cambridge: MIT Press.

Millikan, R. (1998). "A common structure for concepts of individuals, stuffs, and real kinds." *Behavioral and Brain Sciences* 21:55-65.

Norman, D. (1986). *The Psychology of Everyday Things*. New York: Basic Books.

Owens, D. A., Chiang, L., and Muller, C. (1996). "Visual control of vehicle stopping distance: Do we see speed or kinetic energy?" *Investigative Ophthalmology and Visual Science (Suppl.)* 37: S525.

Paley J. (2004). "Clinical cognition and embodiment." *International Journal of Nursing Studies* 41: 1-13.

Petitot, J. (1995). "Morphodynamics and Attractor Syntax." In Port and van Gelder 1995.

Pfeiffer, R. and Scheier, C. (1999). *Understanding intelligence*. Cambridge: MIT Press.

Plotkin, H. (1993). *Darwin Machines and the Nature of Knowledge*. New York: Harvard University Press.

Port, R. and van Gelder, T. (1995). *Mind as Motion*. Cambridge: MIT Press.

Putnam, H. (1985). *Realism and Reason: Philosophical Papers, Volume 3*. Cambridge: Cambridge University Press.

Reed, E.S. (1996a). *Encountering the world*. New York: Oxford University Press.

Reiner, R. and Pierson, R. (1995). "Hacking's Experimental Realism: An Untenable Middle Ground." *Philosophy of Science* 62: 60-69.

Resnik, D. (1994), "Hacking's Experimental Realism", *Canadian Journal of Philosophy* 24: 395- 412.

Rosa J. and Malter, A. (2003). "E-(embodied) Knowledge and E-commerce: How physiological factors affect online sales of experiential products." *Journal of Consumer Psychology* 13: 63-73.

Scarantino, A. (2003). "Affordances Explained." *Philosophy of Science* 70: 949-961.

Shapere, D. (1993). "Astronomy and Anti-Realism." *Philosophy of Science* 60: 134- 150.

Slocum, A., Downey, D., and Beer, R. (2000). "Further experiments on the evolution of minimally cognitive behavior: From perceiving affordances to selective attention." In Meyer, J. et al (eds.) *From Animals to Animats 6*. Cambridge: MIT Press.

Smith. Brian C. (1991). "The Owl and the Electric Encyclopedia." *Artificial Intelligence* 47: 251-288.

Smith, Brian C. (1996). *On the Origin of Objects*. Cambridge: MIT Press.

Smith, B. C. (1999). "Situatedness/Embeddedness." In Wilson and Keil (eds.) *The MIT Enclclopedia of Cognitive Science*, Cambridge: MIT Press.

Stoffregen, TA (2003). "Affordances as properties of the animal-environment system." *Ecological Psychology* 15: 115-134

Thelen, E. (1995). "Time -Scale Dynamics and the Embodiment of Embodied Cognition." In Port and van Gelder 1995.

Thelen, E. and Smith, L. (1993). *A Dynamic Systems Approach to the Development of Cognition and Action*. Cambridge: MIT Press.

Torretti, R. (1986). "Observation." *British Journal for the Philosophy of Science* 37: 1-23.

Turvey, M. (1990). "Coordination." *American Psychologist*, 45: 938-953.

Turvey, M. (1992). "Affordances and Prospective Control: An outline of the ontology." *Ecological Psychology*, *4*, 173-187.

Turvey, M. (1996). "Dynamic touch." *American Psychologist* 51: 1134-1152.

Turvey, M, Shaw, R, Reed, E, and Mace, W. (1981). Ecological laws of perceiving and acting: In reply to Fodor and Pylyshyn (198 1). *Cognition* 9: 237-304.

Ullman, S. (1980). "Against direct perception." *Behavioural and Brain Sciences* 3: 373-415.

van Gelder, T. (1995). "What might Cognition be if not Computation?" *Journal of Philosophy*, 91, 345-381.

Varela, F., Thompson, E. and Rosch, E. (1991). *The Embodied Mind*. Cambridge: MIT Press.

Wheeler, M. (1996). "From Robots to Rothko." In Boden, M. (ed), *The Philosophy of Artificial Life*. New York: Oxford University Press.

Wheeler, M. (to appear). *Reconstructing the Cognitive World*. Cambridge: MIT Press.