2

Is Embodiment Necessary?

Mark H. Bickhard

Department of Philosophy, Lehigh University,
Bethlehem, PA, USA

Every known instance of a genuine cognitive agent is embodied, and it is clear that embodiment has a major influence on, and can be a major help for, artificial agents (Pfeifer & Scheier, 1999). It is also clear that the notion of embodiment has multiple interpretations (Ziemke, 2001; Wilson, 2002; Svensson & Ziemke, 2005). But these points leave open a basic question: Is there a sense of embodiment in which being embodied is *necessary* to cognition? Is embodiment in some way necessary to the nature of cognition, or is it (merely) an important but secondary consideration for (most) implementations of cognitive agents (whether biological or artificial)? I will argue that embodiment is in fact necessary—it is essential in the nature of representation, and, therefore, of cognition.

There are several parts to this argument. First, there are considerations of approaches to the modeling of representation that do not have as consequences any such necessity of embodiment. I will argue that these approaches do not and cannot succeed. Second, there is the development of an alternative model of the nature of representation and cognition. This will be outlined, and it is clear that it requires embodiment, because it requires genuine interaction between a cognitive agent and the world. Finally, a word or two will be in order to look at the *kind* of embodiment that is involved in this *interactive* approach to representation and cognition.

CRITIQUES

Critiques of standard models of representation can be partially compressed because they are all heirs to an underlying error. This error has many manifestations, some of which have been known for millennia and some of which have been discovered relatively recently. Because the central error is held in common among the various models of representation on offer in the literature, only minimal particularization of the central critiques is required in order to demonstrate the applicability of those critiques to specific models.

This central error (or family of errors) has to do with the normativity of representation: the sense in which representation can be true or false. One criterion for a model of representation is that the model be able to account for the simple possibility that the representation is in error. This can be difficult if the representational relationship is purportedly constituted in some sort of factual relation between representation and represented—for example, a causal or informational or nomological relation—because the proper factual correspondence to constitute a representation cannot exist unless the environmental end of the correspondence exists, so that the representation of that existence, in such a view, *must* be correct. There have been multiple attempts to avoid this problem,[1] but I will focus primarily on a strengthened variant of it that has *not* been addressed.

This variant is the criterion of being able to account for *organism* (or *system*) *detectable* error—that the system can itself detect its own errors. Such possibilities of detection may be restricted to certain kinds of organism complexity, and may be quite fallible, but we know that they occur, so any model of representation that cannot account for the possibility of such detections is at best incomplete, and any model that precludes such detections is refuted. This criterion is, in fact, of central importance to any complex cognitive system because system detectable error is necessary in order for error guided behavior and learning to occur, and error guided behavior and learning underpin major portions of most species' cognitive world.

Nevertheless, there is no attempt to address this criterion in the major approaches to representation in the literature. One reason why, so I argue, is that there is no major approach that can possibly account for system detectable error.

Some sense of the depth of the problem posed by this criterion can be found by realizing that it is equivalent to the radical skeptical argument: we cannot check our own representations for truth or falsity because, in order to do so, we would have to step outside of ourselves to obtain independent epistemic access to what we are attempting to represent and then compare what we are trying to represent with our attempted representation of it (Rescher, 1980). We cannot step outside of ourselves, so this is impossible—therefore, so this argument goes, checking our representations for error is impossible. Again, however, we know that error guided behavior and learning occur, so there must be something wrong with this argument. It is not "merely" an armchair philosophical argument: it is a long-standing manifestation of an error in fundamental assumptions about representation.

[1] Without success, I argue elsewhere (Bickhard, 1993, 2004a, in press, in preparation; Bickhard & Terveen, 1995).

The radical skeptical argument also illustrates a second, related, criterion: we must be able to compare represented with representation, so we must have access to our own representational contents—we (or the organism, or system, or agent, or central nervous system, etc.) must have access not only to the represented (which is what the skeptical argument focuses on) but also to the *content* that we are applying to the represented. We must have access to this content in order to make the comparison in order to determine whether or not the content truly applies to the represented. So, in order to engage in any such comparison, we must have access to both sides that are to be compared. Most models fail this second criterion as well as the first.

FODOR

With this pair of criteria in mind, then, I will take a look at some of the central contemporary models of representation, and I begin with Jerry Fodor. Fodor's model is a version of an information semantics, with the crucial representation constituting relationship, in this case, being a nomological relationship between the represented and the representing state in the organism (Fodor, 1987, 1990, 1991, 1998, 2003). As such, the model encounters difficulties accounting for the possibility of representational error, and Fodor has proposed an ingenious attempted solution.

Fodor's model turns on the intuition that false evocations of a representation are dependent on correct evocations, but that there is no reverse dependency—the dependency is asymmetric. Thus, a cow representation may be evoked by cows, but also perhaps by a horse on a dark night. But the possibility of evocation by the horse is dependent on the possibility of evocation by cows, and this dependency is not reciprocated: evocations by cows could continue even if there were never any possibility of evocations by horses.

One problem with this proposal is that such asymmetric dependencies, even among nomological relationships, do not suffice to pick out representation at all. For example, a neurotransmitter docking with a receptor triggers ensuing activity in a nomological manner, and a mimicking poison molecule docking with the same receptor also evokes nomologically related activity—and the poison molecule's possibility of such evocation is asymmetrically dependent on the possibility of the neurotransmitter evoking such activity. Yet there is at best a biologically functional relationship here, not a representational relationship (Levine & Bickhard, 1999).

Setting this concern aside, however, we still find that the alleged representational error is characterizable as error only for an external observer who could (1) (supposedly) determine the counterfactual asymmetries involved among various families of nomological relationships in order to characterize what a representation is supposed to represent—that is, to characterize the content, (2) epistemically access the represented, and (3) compare the two in order to determine whether the content holds of the represented.

Note that this external observer is in precisely the position that the radical skeptical argument points out that no actual epistemic agent can be in for itself. First, no epistemic agent can have access to its own relevant counterfactual nomological relations in order to determine content, and, second, to access the represented in order to make the comparison with the content is precisely the representational problem all over again. This is the circularity that is at the center of the skeptical argument.

For Fodor, the consequence is that content is not accessible, the represented is not independently accessible, and system detectable error is therefore impossible. Consequently, error-guided behavior and learning are not possible. But error-guided behavior and learning occur, therefore the model is refuted.

MILLIKAN

In Millikan's etiological model, representing X is a particular kind of function that some things or conditions might have, and having such a function is constituted in having (or being properly derived from something that has) the right kind of evolutionary selection history (Millikan, 1984, 1993). The crucial selection history is one of undergoing a sufficient number of generations of selection for having the (functional) consequence in question.[2] In this sense, having a function is constituted, roughly, in being designed to have that function by evolutionary selection as the designer.

I would first note that this approach does not have the problem of accounting for the possibility of representational error *per se*: what something is supposed to represent—content—is determined by evolutionary history, while what is being represented is in the present. The two are thus pulled apart in a way that permits the possibility that the content will be incorrectly applied to a present entity or state of affairs—in a way that is false. This is a distinct advantage over information semantic approaches. The etiological approach, however, ultimately does not succeed either.

One specific problem with this approach is that it renders function, thus representation, causally epiphenomenal: having a function is constituted in the past, not in the present, and present state of a system is not specific to having that requisite past evolutionary history. This is illustrated by Millikan's example of a lion that pops into existence that is molecule by molecule identical to a lion in the zoo: the lion in the zoo has all the right evolutionary histories, therefore its organs and processes have functions, including functions of representing, but the science fiction lion has no evolutionary history, therefore no functions at all,

[2] The number of generations required is a matter of discussion and dispute (Godfrey-Smith, 1994). This literature cannot use the locution above of having a functional consequence, and must make the point much more indirectly, because no such consequence can be a functional consequence until the requisite number of generations of selections have transpired.

whether representational or not. Yet the two lions, by assumption, are causally, dynamically, identical: etiological function makes no causal difference.

This is a thought experiment, but an equivalent actual case occurs every time in evolutionary history that something occurs for the first time and is selected for. If sufficient generations of such selections ensue, then whatever organ is involved will come to have that consequence as its function, and it will serve that function when it produces that consequence (in the right conditions). But this first time and all subsequent times till the magic of constituting a function occurs are instances in which identical (or extremely similar) consequences occur but are not (etiologically) functional. So, again, we have some systems that are causally identical to others, yet do not have functions, whereas the others do have functions. Again we have that etiological function is causally epiphenomenal.

To return now to the general critique, no organism has access to the evolutionary histories of its parts, therefore no organism (or system) has access to its own representational contents. Therefore system detectable error is not possible.

Furthermore, the problem of access to what is being represented is identical in this case to the case for information semantics: that is the fundamental representational problem all over again. Again we encounter the fundamental circularity of the radical skeptical argument.[3]

THE SYMBOL SYSTEM HYPOTHESIS

Cognitive science was long dominated by computational approaches in which relevant processes were symbol manipulation processes (Franklin, 1995), and certainly such approaches are still very prevalent. There are, obviously, many powers and advantages afforded by such design approaches over, for example, simple associationistic approaches, but, regarding representation *per se*, they are hopeless. Basic representations in such models are taken to represent something in virtue of being in a correspondence with that something—a correspondence that somehow encodes its distal end—with the crucial correspondence variously taken to be one of a "stand-in" or perhaps a structural isomorphism (Newell, 1980; Vera & Simon, 1993). But such models cannot account for the bare possibility of representational error, and have no way to address the possibility of system detectable representational error (Bickhard, 2004a, in press, in preparation). If the crucial representational (encoding) relationship exists, then it is correct, and if it does not exist, then the representation does not exist, and there is no third possibility for modeling the representation existing but being incorrect.

[3]Dretske's approach to representation (Dretske, 1988) is also an etiological approach, though with a learning etiology instead of an evolutionary etiology, and, so suffers from similar problems. For more detailed analyses of Fodor's and Millikan's models, as well as those of Dretske (1988) and Cummins (1996), see Bickhard (2004a, in press, in preparation).

CONNECTIONIST REPRESENTATION

There was great hope for connectionist approaches in the 1980s because, among other reasons, they offered a seemingly natural way to realize learning processes. Certainly, connectionist models accommodate training in ways that symbol manipulation processes do not, but, again, with respect to representation *per se*, they do not offer any real advantages. In particular, a trained connectionist net establishes an informational correspondence with some class of input patterns, and, therefore, is a version of an information semantical approach.[4] Therefore, they suffer all of the same problems as do other versions, including the impossibility of system detectable error: no connectionist system has access to its own informational relationships, nor to what those relationships might be with, in order to be able to make any relevant comparisons regarding the possibility of error.[5]

AGENTIVE ANTI-REPRESENTATIONALISM

With the development of dynamic and agentive approaches in the late 1980s and 1990s, a strong anti-representationalism emerged. The very concept of representation was argued to be unnecessary and even misleading in designing and understanding complex organisms and systems (Brooks, 1991; van Gelder, 1995). One of the stronger *counter* arguments turned on what were called representation-hungry situations, in which representational tracking, such as of a hidden predator, was required (Clark & Toribio, 1995). These discussions were confused because different detailed conceptions of what constituted representations were often involved: for example does tracking *per se* constitute representation or is representation constituted only in manipulable symbols that track?

In either case, the underlying conceptions of representation were of the basic representational correspondence sort—something that encodes—and did nothing to avoid the fundamental problems of accounting for the possibility of error and of system detectable error. That is, my claim here is that both the anti-representationalists and the representationalists made equivalent fundamental assumptions about the nature of representation, and that both were wrong—neither could account for the normativities, the endogenous truth values, of representation.

If this is correct, then clearly the issues regarding the embodiment of cognitive systems cannot be properly addressed within such frameworks. In fact, issues of embodiment of cognitive systems cannot be properly addressed within

[4] Ironically, this model differs from Fodor's primarily in that Fodor's transductions are nomological, while connectionist nets are trained. In at least one of Fodor's incarnations, transduction is the evocation of representation via some process other than inference, and, by this definition, trained connectionist nets *are* Fodorian transducers (Bickhard, 1993; Bickhard & Terveen, 1995).

[5] For extensive discussion of connectionist systems, see Bickhard and Terveen (1995).

any of the frameworks that have dominated the history of cognitive science:[6] representation as correspondence or representation as encoding is a strictly input processing notion of representation. In any such view, there is no necessary involvement of any body, other than to house an input processing system. There is no necessity for action, thus none for a body that can act. In such a view, then, embodiment can at best constitute a source of constraint, quite possibly enabling constraint, on processing and on action, but can have no deep relationships with representation *per se*.

INTERACTIVE REPRESENTATION

This is in strong contrast to models of representation that derive representation from action and interaction, from general pragmatist principles (Rosenthal, 1983; Joas, 1993). If representation, thus cognition, is derived from (inter)action, then some sort of embodiment is required in order for such action and interaction to be possible: actions, thus embodiment, are not mere auxiliaries to representation, but, instead, are essential to it.

The task at this point, therefore, is to outline such an action-based model of representation, and show that it is a viable alternative to the passive, encoding correspondence, models addressed above. In particular, the task is to show that the model can account for the possibility of representational error and of system detectable error.

Consider a complex agent—animal or artificial. It faces an ongoing task of selecting what interactions to engage in and of guiding interactions underway. In order to select (or guide), there must be available in some functionally accessible way indications of what interactions are possible: it does no good to select opening the refrigerator to get something to eat if you are in the forest, miles from the nearest refrigerator. A frog, for another example, might have two possibilities for flicking its tongue and eating a fly, one possibility for eating a worm and one for jumping in the water. It does no good for the frog to select a tongue-flicking interaction for eating a fly when no detection of anything that might support such tongue-flicking-and-eating has occurred.

There are two aspects to this point: (1) some means for indicating interaction possibilities and (2) some way of selecting among those possibilities. Both aspects are of fundamental importance, but I will be focusing on the indications here.[7]

[6]There is a more general critique that underlies this one. If this more general critique is correct, then the fundamental problems extend much further back in time: to the pre-Socratics (Bickhard, 2006, in press, in preparation).

[7]The processes of selection underlie motivation (Bickhard, 2003), while, so I will argue, the indications underlie representation.

TRUTH VALUE

The crucial point about such indications of interaction possibilities for current purposes is that they constitute the emergence of the most fundamental aspect of representation: truth value. In particular, such indications, if selected, may turn out to be correct, in the sense that the interaction proceeds as indicated, or they may not, if the actual interaction violates the range of indicated possibilities.[8] The frog might flick its tongue, but not be able to proceed with eating because it was just a pebble, not a fly.

Note that the functional task of selecting among indicated interactive possibilities is inherent for any complex agent. The evolution of such indicative capabilities, therefore, is necessary for the evolution of complex agents. Consequently, the evolutionary emergence of representational truth value is inherent in the evolution of complex agents. And, therefore, the evolutionary emergence of representation is inherent in the evolution of complex agents.[9]

MORE COMPLEX REPRESENTATIONS

However, indications of interactive possibilities may possess truth values, but they do not look much like “standard” representations, such as of objects. Can this framework account for more complex representing?

At this point, I can borrow from one of the few extant models of representation in the literature that is based on action: that of Jean Piaget. In particular, Piaget has outlined a model of the representation of small manipulable objects in terms of organizations of potential actions that I can translate directly into the language of the interactive model (Piaget, 1954).[10]

Consider a child’s wooden toy block: it offers (or affords)[11] multiple possible interactions, ranging from visual scans to manipulations to throwing, chewing, and so on. A fundamental manifestation of the persistence of objects is that, if any of these potential interactions are available, then they all are with, perhaps, intervening interactions. A particular visual scan, for example, might require an intermediate turning of the block so that what was the reverse side is now visible. In this sense, all visual and manipulation interactions are reachable from each other. Furthermore, this organization of internally reachable interaction possibilities has the property of being invariant under an important class of additional

[8] This point turns on the normativity of such indications being “correct” or “incorrect,” and, therefore, on a more primitive, pragmatic, kind of normativity. Elsewhere, I develop a model of the emergence of normative function—a non-etiological model—that serves as the framework for this emergence of representational normativity (Bickhard, 1993, 2004a, in press, in preparation; Christensen & Bickhard, 2002).

[9] For discussion of how such anticipative indicating might occur in central nervous systems, see Bickhard and Terveen (1995), Bickhard (in preparation).

[10] For discussions of some of the differences between Piaget’s model and the interactive model, see (Campbell & Bickhard, 1986; Bickhard, 1988; Bickhard & Campbell, 1989).

[11] For discussions of Gibson’s model, see Bickhard and Richie (1983).

interactions: for example, if the block is left on the floor, it can (sometimes) be recovered by walking back into that room. The organization of internally reachable interaction potentialities, thus, is itself reachable given a reversal of locomotions and transportations—it is invariant under such locomotions and transportations. This is a further manifestation of the permanence of objects. Clearly, this organization is not invariant under transformations that destroy the block. For small children, these properties constitute what objects are.

For my current purposes, the important point is that an action- and interaction-based model of representation can address more complex forms of representation. There are many further kinds of representation (or purported representation) and cognition to address—for example abstract representation, perception, language, and so on[12]—but it is clear that the model is not restricted to just simple anticipations of single interaction possibilities. It is a candidate for capturing the basic ground of representation.

EMBODIMENT IS NECESSARY

There is a narrower point, however, that is sufficient for the point at issue here: Is embodiment necessary? The narrower point is that interactive representation clearly exists (whether or not it accounts for *all* representation), and it clearly requires interaction, and, therefore, requires some form of embodiment in order to be able to interact. It clearly exists because it is required for complex agents to function in the world, and it is representational because it manifests truth value. So embodiment is necessary for any form of cognition or representation that is of an anticipative interactive kind; such cognition and representation do exist, and therefore embodiment is necessary.

The case can be made stronger, however: if indicated interactions are engaged in and they fail to honor the indicated range of possibilities, then the indicated interactions are false, and they are falsified for the organism or system itself.[13] System detectable error can be accounted for within this model. Because no other model currently on offer can account for system detectable error, there is a strong claim that embodiment is required for *all* representation and cognition: absent an alternative model of representation that can satisfy the system detectable error criterion

[12] See, for example, Bickhard (2003, 2004b, 2005b, 2006, 2007a, in press, in preparation).

[13] It is important to note that detections of such failures of interaction anticipation do not require detecting any external, environmental, anticipated consequences. Such external anticipations would require being represented, and, thus, would not avoid the skeptical argument. Anticipated *internal* processes, however, do not require representation in order to be monitored: this can be done strictly internally and functionally. System detectable error, therefore, is possible. The skeptical argument is a valid argument, but it is based on an assumption that all representation has the form of encoding, and that assumption is false. The argument, then, is unsound, and constitutes a reductio of encoding models of representation.

without requiring interaction, thus embodiment, the hypothesis that all representation and cognition requires (interactive) embodiment is the only one left standing.[14]

WHAT KIND OF EMBODIMENT?

Embodiment is involved, and even required, for multiple phenomena. Any living being must be realized in some form that can maintain itself in its environment(s). Action and interaction are constrained and enabled not only by environments, but also by the specific forms of embodiment that must engage in those interactions—consider, for an extreme example, the differences between a caterpillar and a butterfly.

But perhaps these consequences and necessities of embodiment are different from those necessitated by representation and cognition. What can be said about the minimal form of embodiment that is necessitated for representation *per se*?

The central property involved in the interactive form of representation is the functional anticipation of future interactive processes (Bickhard, 2005a). I have been focusing on the necessity that there be involvement of the environment in determining the course of those interactive processes.[15] This involvement of the environment imposes a basic requirement on the nature of the embodiment of a representational system: the interactions must be of a sort that influences the environment in such ways as will, in turn, influence the internal interactive processes of the organism (or system). The interactions must be capable of being *full* interactions, not just inputs being processed into outputs—outputs that, in turn, have no influence on subsequent inputs.

A minimal embodiment requirement, therefore, is that the system or organism be capable of full interactions, with outputs having causal influence on inputs—with circular causal flow. This is actually a quite minimal requirement: many

[14] The claim here is not that all representation is directly interactive. Encodings do exist—Morse code is a perfectly good example—but genuine encodings are in fact representational stand-ins, and thus, they require something to be stood-in-for. Encodings cannot stand-in for other encodings unboundedly—that is the infamous regress of interpreters of symbolic encodings—so there must be some other form of representation that can ground any such encoding stand-in hierarchy. That other form of representation, thus the *basic* form of representation, is interactive representation. And that is what requires embodiment.

[15] There is also the requirement that those functional anticipations be normative, in the sense that they can be true or false. I have set aside the normativity requirement here, but it too is crucial. For an account of the emergence of basic normativity, see Bickhard (2004a, in press, in preparation). This emergent normativity, interestingly, imposes its own requirements on the nature of normative systems. In fact, they impose constraints on forms of embodiment—and thereby serve as the foundation for interactive representation. In particular, they constrain normativity to certain kinds of far from thermodynamic equilibrium systems, of which all living systems are examples. The nature of such systems, in turn, requires forms of interactivity, which constitute the ground for interactive representation—representation builds on normativity in a very natural way. Normativity, thus, imposes additional constraints on what kinds of systems can manifest representation and, thus, be cognitive systems, and these impose additional constraints on forms of embodiment. But those constraints on *forms* of embodiment do not add to the basic *necessity* that cognitive systems be embodied.

kinds of systems that have unusual embodiments—perhaps, for example, being constituted in some distributed network fashion over large, and perhaps even ongoingly changing, spatial regions—could satisfy it. What cannot satisfy it are passive input processors or throughput processors for which the outputs have no particular influence on the inputs. That is, passive computer systems, including passive connectionist systems, cannot be interactive in the required sense.

Animals and robots, clearly, *can* be interactive in the required sense. And, as mentioned at the beginning of this discussion, they are embodied. Although embodied and capable of being interactive, it is nevertheless an interesting question whether mechanical robots can capture full representation and cognition, whether they can manifest the normativities involved in interactive anticipations being true or false.[16]

CONCLUSION

Embodiment is necessary for representation and cognition. The minimal requirements on embodiment are, on the one hand, sufficient to eliminate most artificial systems in the world today as candidates for being genuinely cognitive, but, on the other hand, they are quite general in what they do impose on sufficient forms of embodiment. But it is not just interaction that is required for representation, it is also normative anticipation of interaction, and that involves its own further requirements.

REFERENCES

Bickhard, M. H. (1988). Piaget on variation and selection models: Structuralism, logical necessity, and interactivism. *Human Development*, *31,* 274–312.

Bickhard, M. H. (1993). Representational content in humans and machines. *Journal of Experimental and Theoretical Artificial Intelligence*, *5,* 285–333.

Bickhard, M. H. (2003). An Integration of motivation and cognition. In L. Smith, C. Rogers, & P. Tomlinson (Eds.), *Development and Motivation: Joint Perspectives* (pp. 41–56). Leicester: British Psychological Society, Monograph Series II.

Bickhard, M. H. (2004a). Process and emergence: Normative function and representation. *Axiomathes—An International Journal in Ontology and Cognitive Systems, 14*, 135–169. Reprinted from: Bickhard, M. H. (2003). Process and emergence: Normative function and representation. In J. Seibt (Ed.), *Process Theories: Crossdisciplinary Studies in Dynamic Categories* (pp. 121–155). Dordrecht: Kluwer Academic.

Bickhard, M. H. (2004b). The social ontology of persons. In J. I. M. Carpendale & U. Muller (Eds.), *Social Interaction and the Development of Knowledge* (pp. 111–132). Mahwah, NJ: Erlbaum.

Bickhard, M. H. (2005a). Anticipation and representation. In C. Castelfranchi (Ed.), *From Reactive to Anticipatory Cognitive Embodied Systems* (pp. 1–7). *AAAI Fall Symposium, Technical Report FS-05.05*. Menlo Park, CA: AAAI Press.

Bickhard, M. H. (2005b). Consciousness and reflective consciousness. *Philosophical Psychology, 18*(2), 205–218.

Bickhard, M. H. (2006). Developmental normativity and normative development. In L. Smith & J. Voneche (Eds.), *Norms in Human Development* (pp. 57–76). Cambridge: Cambridge University Press.

[16] Elsewhere I argue that ‘Mechanism is not enough’ (Bickhard, 2007b).

Bickhard, M. H. (2007a). Language as an interaction system. *New Ideas in Psychology, 25*(2), 171–187.

Bickhard, M. H. (2007b). Mechanism is not enough. In Q. Gonzalez, M. Eunice, W. F. G. Haselager, & I. E. Dror (Eds.), *Mechanicism and Autonomy: What Can Robotics Teach Us About Human Cognition and Action?* Special issue of *Pragmatics and Cognition, 15*(3), 573–585.

Bickhard, M. H. (in press). The interactivist model. *Synthese*.

Bickhard, M. H. (in preparation). The Whole Person: Toward a Naturalism of Persons—Contributions to an Ontological Psychology.

Bickhard, M. H. & Campbell, R. L. (1989). Interactivism and genetic epistemology. *Archives de Psychologie, 57*(221), 99–121.

Bickhard, M. H. & Richie, D. M. (1983). *On the Nature of Representation: A Case Study of James Gibson's Theory of Perception*. New York: Praeger Publishers.

Bickhard, M. H. & Terveen, L. (1995). *Foundational Issues in Artificial Intelligence and Cognitive Science: Impasse and Solution*. Amsterdam: Elsevier Scientific.

Brooks, R. A. (1991). Intelligence without representation. *Artificial Intelligence, 47*(1–3), 139–159.

Campbell, R. L. & Bickhard, M. H. (1986). *Knowing Levels and Developmental Stages*. Basel, Switzerland: Karger, Contributions to Human Development.

Christensen, W. D. & Bickhard, M. H. (2002). The Process Dynamics of Normative Function. *Monist, 85*(1), 3–28.

Clark, A. & Toribio, J. (1995). Doing without representing? *Synthese, 101*, 401–431.

Cummins, R. (1996). *Representations, Targets, and Attitudes*. Cambridge, MA: MIT Press.

Dretske, F. I. (1988). *Explaining Behavior*. Cambridge, MA: MIT Press.

Fodor, J. A. (1987). *Psychosemantics*. Cambridge, MA: MIT Press.

Fodor, J. A. (1990). *A Theory of Content and Other Essays*. Cambridge, MA: MIT Press.

Fodor, J. A. (1991). Replies. In B. Loewer & G. Rey (Eds.), *Meaning in Mind: Fodor and His Critics* (pp. 255–319). Oxford: Blackwell.

Fodor, J. A. (1998). *Concepts: Where Cognitive Science Went Wrong*. Oxford: Oxford University Press.

Fodor, J. A. (2003). *Hume Variations*. Oxford: Oxford University Press.

Franklin, S. (1995). *Artificial Minds*. Cambridge, MA: MIT Press.

Godfrey-Smith, P. (1994). A modern history theory of functions. *Nous, 28*(3), 344–362.

Joas, H. (1993). American pragmatism and German thought: A history of misunderstandings. In H. Joas (Ed.), *Pragmatism and Social Theory* (pp. 94–121). Chicago, IL: University of Chicago Press.

Levine, A. & Bickhard, M. H. (1999). Concepts: Where Fodor went wrong. *Philosophical Psychology, 12*(1), 5–23.

Millikan, R. G. (1984). *Language, Thought, and Other Biological Categories*. Cambridge, MA: MIT Press.

Millikan, R. G. (1993). *White Queen Psychology and Other Essays for Alice*. Cambridge, MA: MIT Press.

Newell, A. (1980). Physical symbol systems. *Cognitive Science, 4*, 135–183.

Pfeifer, R. & Scheier, C. (1999). *Understanding Intelligence.* Cambridge, MA: MIT Press.

Piaget, J. (1954). *The Construction of Reality in the Child*. New York: Basic Books Inc.

Rescher, N. (1980). *Scepticism*. Totowa, NJ: Rowman and Littlefield.

Rosenthal, S. B. (1983). Meaning as habit: Some systematic implications of Peirce's pragmatism. In E. Freeman (Ed.), *The Relevance of Charles Peirce* (pp. 312–327). LaSalle, IL: Monist.

Svensson, H. & Ziemke, T. (2005). Embodied representation: What are the issues. In B. Bara, L. Barsalou, & M. Buccarelli (Eds.), *Proceedings of the 27th Annual Meeting of the Cognitive Science Society* (pp. 2116–2121). Mahwah, NJ: Lawrence Erlbaum.

van Gelder, T. J. (1995). What might cognition be, if not computation? *The Journal of Philosophy, XCII*(7), 345–381.

Vera, A. H. & Simon, H. A. (1993). Situated action: A symbolic interpretation. *Cognitive Science, 17*(1), 7–48.

Wilson, M. (2002). Six views of embodied cognition. *Psychonomic Bulletin & Review, 9*(4), 625–636.

Ziemke, T. (2001). Are robots embodied? In C. Balkenius, J. Zlatev, C. Brezeal, K. Dautenhahn, & H. Kozima (Eds.), *Proceedings of the First International Workshop on Epigenetic Robotics: Modeling Cognitive Development in Robotic Systems* (pp. 75–93). Lund, Sweden: Lund University Cognitive Studies.