

Robots and Representations*

Mark H. Bickhard

Cognitive Science, Psychology, Philosophy

17 Memorial Drive East

Lehigh University

Bethlehem, PA 18015 USA

610-758-3633

610-758-6277 fax

mhb0@lehigh.edu

<http://www.lehigh.edu/~mhb0/mhb0.html>

Abstract

The design of complex interactive robots inherently yields a form of representation — an interactive form. Interactive representation is, arguably, the foundational form of representation from which all others are derived. It constitutes the emergence of representational truth value *for the system itself*, a criterion not addressed in current literature.

There is a form of representation that arises naturally in the design of complex interactive systems — robots. This form arguably constitutes an emergence of the fundamental form of representation, out of which increasingly complex forms are constructed and derived. Furthermore, this form of representation naturally satisfies an essential meta-epistemological criterion for original representation: system detectable truth value. No alternative approach to representation in the current literature addresses this criterion. Recognizing and exploiting the emergence of this form of representation in robotics and dynamic systems is a rich frontier for exploration.

In standard artificial intelligence and cognitive science models, inputs are received and processed, and, perhaps, outputs emitted. The critical consideration that arises in robot design is the possibility of a closure of this sequence of process such that robot outputs influence subsequent inputs via the environment, and, therefore, influence subsequent internal states and processes in the robot. That is, the critical consideration is the closure of input, processing, and output to include full *interaction*, not just action.

This simple closure introduces several important possibilities. In particular, possible internal states that might be consequent on some action or course of action can be functionally indicated in the robot. Because such possible consequent states in the robot will depend, in part, on the environment, those states, or some one of those states, may or may not actually occur — the environment may or may not yield the appropriate input(s) in response to the output(s) to induce those indicated states in the robot. If none of those indicated states are entered by the robot, then the indications are false, and are *falsified for the robot*. The error in such indications is detectable by and for the system itself.

In effect, to indicate such internal states as consequent on particular interactions on the part of the robot is to implicitly predicate of that environment whatever properties are sufficient to support those indications. It is to anticipate that the environment will in fact respond as indicated, if the interaction is engaged in. Some environments will possess a sufficiency of those properties, and will yield one of the indicated states, while other environments will not possess such properties, and will not yield any of the indicated states. For those environments that do *not* yield an indicated state, to set up such an indication is to set up an implicit predication, an anticipation, that is false, and potentially falsifiable by the system (Bickhard, 1993, in press; Bickhard & Terveen, 1995).

The possibility of error, and especially of system detectable error, is a fundamental meta-epistemological criterion for representation. Whatever representation is, it must be capable of some sort of truth value. Conversely, something *is* representation *for a particular system* only if it is capable of some sort of truth value *for that system*. This is critical because many states and conditions and phenomena are representational — can have truth value — but only for some user or designer or observer outside of the system itself, not for the system itself (Bickhard, 1993; Bickhard & Terveen, 1995).

Moderately complex robots, then, naturally involve a form of representation that is representational for the robot, not just for an observer or analyst or designer or user of the robot. This claim generates five questions: 1) How can notions such as ‘indication’ in the above discussion be made good in a functional manner in a robot, without committing a logical circularity by presupposing the very representationality that is allegedly being modeled? 2) Why would it be useful for a robot to have such representations of interactive potentialities? 3) How could such a notion of representation possibly be adequate to “normal” representational and cognitive phenomena such as representation of objects; representation of abstractions, such as numbers; language; perception; rationality; and so on? I will only outline an answer to the first of these questions, referring others to other sources. 4) On what basis would a robot set up such indications? And 5) How does this model of representation relate to contemporary research in artificial intelligence, cognitive science, connectionism, and robotics? My responses to this question too will, obviously, be abbreviated.

The Functional Story

First, I need to address the question of how interactive representation could be implemented without presupposing representation. All that is needed are some architectural principles adequate to the model that are themselves strictly functional — not representational. This is, in fact, rather simple. The indicated internal outcome states for an interaction function like final states in an automaton recognizer, but for an automaton that emits outputs to an interactive environment (Bickhard, 1980a). The indication of such states can be implemented with pointers — a pointer, say, to some location that will contain a “1” in the state being indicated. This is certainly not the only architecture that will implement the notions required, but it does suffice.

To indicate the interaction itself, upon which the indications of final states are based, requires only a pointer to the subsystem — perhaps the subroutine or interactive recognizer — that would engage in those interactions. So, a pointer to a subsystem together with a pointer or pointers to final states associated with that subsystem suffices for the implicit predication of interactive representation, but none of these pointers themselves are or require representation. Insofar as there is representation here, it is genuinely emergent in the architectural organization.

The Usefulness of Interactive Representations

Choice. Why would it be useful for a robot to have such indications? For two reasons: First, if there are multiple interactions possible in a particular environment, the indicated internal outcomes of those interactions can be used in selecting which interaction to actually engage in (Bickhard, 1997b). A frog seeing a fly might set up indications of the possibility of tongue-flicking-and-eating, while a frog seeing a shadow of a hawk might set up indications of the possibility of jumping in the water. A frog seeing *both* needs some way to decide, and internal outcome indications provide a basis for such decision — e.g., select the interaction with the indicated outcomes that have the highest priority relative to current set-points. (Note that if the relevant outcomes are presumed to be *represented*, rather than indicated — as must be the case if the outcomes are considered to be external outcomes in the environment — then there is a circularity involved in using such notions to model representation.)

Error. The second reason why such indications might be useful is that they create the possibility of error, and — most importantly — the possibility of the detection of error by the system. Detection of error, in turn, can be useful for guiding heuristics and strategies of interaction, and for evoking and guiding learning processes. Any general form of learning, in fact, *requires* such system detection of error (Bickhard & Terveen, 1995). In slogan form: Only anticipations can be falsified; therefore only anticipations can be learned.

On the Adequacy of Interactive Representation

Interactive, or robotic, representation might seem adequate for the kinds of interactive properties that interactive indications will implicitly predicate of the environment, but there are many other things to be represented that do not *prima facie* look like interactive properties. To make good on claims of the adequacy of interactive representation as a general form of representation would require a programmatic treatment of many or most of these representational phenomena. There isn't space to even begin that explication here (see, for example, Bickhard, 1980a, 1980b, 1992, 1993, in press, forthcoming; Bickhard & Campbell, 1992; Bickhard & Richie, 1983; Bickhard & Terveen, 1995; Campbell & Bickhard, 1986, 1992), but I will outline an approach to the interactive representation of physical objects in order to indicate that this is at least a plausible programme.

Complexities of Interactive Indications.

Before addressing objects *per se*, I need to outline some forms of complexity that can be involved in interactive indications. The first is that there may be multiple interactive possibilities indicated at a given time. The second is that interactive indications can be conditionalized on each other: interaction **A** with possible outcome **Q** might be indicated, and, if **A** is engaged in and **Q** is in fact obtained, then interaction **B** with possible outcome **R** becomes possible. There are other kinds of complications possible, but branchings and conditionalized iterations of interactive indications will suffice for briefly addressing the problem of object representation.

Webs. Branchings and conditional iterations yield the possibility of interactive indications forming potentially complex webs or nets of indications. In effect, the whole of such a web is indicated as currently possible, but actually reaching some parts of the web will be contingent on perhaps many intermediate interactions and outcomes.

Objects. Some sub-networks in such a complex web may have two critical properties: 1) A subnet may be closed in the sense that, if any part of it is reachable — possible — then all parts of it are. That is, all points (possible states) in such a web indicate the totality of all other points in the web, perhaps with necessary intermediate interactions. 2) Such a closed subnet may be invariant under some interesting class of possible interactions, in the sense that many possible interactions will leave the closed subnet invariant in internal structure and still reachable.

A toy block, for example, will offer many possible interactions to a child. The block can be visually scanned, dropped, thrown, chewed, manipulated, and so on. Furthermore, any of such possibilities will indicate the availability of the rest of them: a particular visual scan becomes momentarily unavailable if the block is turned over, but is recoverable if the block is turned back over to its original position. Still further, this entire organization of interactive possibilities remains invariant for many interactions such as manipulations, throwings, hidings, storing in the toy box, locomotions on the part of the child,

and so on. It does not remain invariant, however, under burning or crushing.

The point is that a block, for an infant, is no more than such closed and invariant organizations of interactive possibilities. The infant knows nothing of substances or molecules — those all come later in theorizing about such things as blocks. This, I propose, is the general manner in which interactive representation can handle object representations. This is a Piagetian sort of account of object representations (Piaget, 1954), and I suggest similarly Piagetian flavored approaches to other representational issues (though I do not endorse all of Piaget's model: Bickhard, 1988; Bickhard & Campbell, 1989). My basic point in this paper, however, is simply that interactive representation is not immediately blocked by the multitudinous forms of representation that must ultimately be accounted for.

How to Set Up Interactive Indications

Interactive indications might be useful for a robot, but on what basis could they be set up? The general answer is simple: on the basis of previous outcomes of previous interactions. Previous outcomes differentiate environments into those that yield that outcome and those that don't. Conditionalized indications of interactive potentialities are, in effect, conditionalized on those prior differentiations. Any environment that will yield state **Q** from interaction **A** will also — according to the conditional — be an environment in which interaction **B** is possible, with possible consequent state **R**. So, if **Q** is obtained, then direct (not conditionalized) pointers to **B** and **R** should be set up. As with the rest of the model, many complexities can arise here, but this general point suffices to show that the problem of how to set up interactive representations is not itself an aporia.

Relationships to Contemporary Work

Symbol Systems. The symbol system hypothesis involves inputs and subsequent processing (Newell, 1980), but does not involve any necessary notion of closure from outputs back to inputs. It does not involve any recognition of the necessity for *interaction*. Correspondingly, it does not solve the problem of system detectable error. Error can be defined from a user or designer perspective, and it can be useful to do so, but this is still error, therefore representation, for the user or designer, not for the system itself.

In fact, the symbol system hypothesis is one of a large class of approaches to representation that attempt to construe representation in terms of the processing of inputs, with no recognition of the essential importance of full interaction (Bickhard, 1993; Bickhard & Terveen, 1995; Coffa, 1991; Dretske, 1981, 1988; Fodor, 1987, 1990; Hanson, 1990; Loewer & Rey, 1991; Millikan, 1984, 1993; Palmer, 1978; Pylyshyn, 1984; Smith, 1987). Inputs are, in these models, supposed to set up correspondences — usually some special kind of correspondences — between internal states of the system and particular properties or entities in the environment, and *are supposed to represent those properties or entities by virtue of those correspondences*. These

approaches have difficulty accounting for error at all, and do not even address the problem of system detectable error. The difficulty in accounting for error arises because, if the correspondence exists, then the representation exists and it is correct, while, if the correspondence does not exist, then the representation does not exist, and so cannot be *incorrect*.

There is a great deal of effort being currently devoted to attempting to account for error in some such model (Dretske, 1988; Fodor, 1990; Hanson, 1990; Loewer & Rey, 1991; Millikan, 1993). I will not address the vicissitudes of these attempts (though I do not think that any of them succeed) because none of them even address the stronger criterion of system detectable error. On their own terms, error, if definable at all, is definable from the perspective of some observer or user or designer or analyzer of the system in question, not for the system itself. There are, in addition, many other problematics of this general approach to representation, none of which arise as problematic for the interactive approach (Bickhard, 1980b, 1993; Bickhard & Terveen, 1995).

Connectionism. Interactive representation is a direct alternative to standard symbol manipulation approaches to representation. What about connectionism? The key to the answer to this question is to note that a connectionist net serves to differentiate instances of one input pattern class from instances of some other class of input patterns (Bickhard & Terveen, 1995; Churchland, 1989; McClelland & Rumelhart, 1986; Rumelhart, 1989; Rumelhart & McClelland, 1986; Smolensky, 1988; Waltz & Feldman, 1988). This is precisely what a transducer is supposed to do in standard information processing models (Bickhard & Richie, 1983; Fodor & Pylyshyn, 1981; Loewer & Rey, 1991). A connectionist net can be trained to differentiate new instances of new patterns, while a transducer has to be designed or evolved specifically for its differentiation task, but, beyond the difference between design versus training origins, the tasks that are performed are the same sorts of differentiation tasks.

Connectionism and information processing approaches agree on the next step in the modeling: the differentiations accomplished by transducers or by connectionist nets are taken to be representations of that which has been differentiated. It is this step that immediately encounters a host of problems, not the least of which is how such "representations" could possibly be in error, and how the system itself could detect such error.

In the interactive model, in contrast, such differentiations are useful in order to serve as the basis for setting up interactive indications, and there is no requirement that they be alleged to be representational in order that they serve that function. Sensory input processing, in other words — or connectionist net processing — does not directly generate representations of what is on the environmental end of those inputs, but, instead, sets up indications of what is interactively possible in the environments differentiated by that input processing. Input processing occurs — and serves a necessary function — in both cases, but the interactive model does not require or assume that the system has any representation of what produces the inputs. (Humans do have such representations

of the sources of [some of] our inputs, and that must be accounted for, but flatworms don't, and it's not clear that frogs do — primitive forms of representation do not necessarily involve such representations.)

Robotics. Interactive representation cannot exist in standard symbol manipulation or connectionist architectures — there are no closed interactions involved in their purported representations — but it arises naturally in robot design. How does this compare with contemporary robotics literature? The problematics of standard conceptions of representation are clearly recognized (Beer, 1990, 1995; Bickhard, 1996; Brooks, 1991a, 1991b, 1991c; Prem, 1995), and robotic and autonomous agent design problems naturally lead to the sorts of indications that constitute interactive representation — whether or not they are recognized as representational (Kuipers, 1988; Kuipers & Byun, 1991; Maes, 1990; Nehmzow & Smithers, 1991, 1992; Stein, 1994). Recognizing the emergence of representation and representational truth value in those designs, with all the concomitant implications for related domains of cognition and robotic functioning, is a rich frontier for robotics and dynamic systems (Bickhard, 1997a, 1997b, forthcoming; Bickhard & Campbell, 1996a, 1996b; Bickhard & Terveen, 1995; Brooks, 1994; Cherian & Troxell, 1995a, 1995b; Hooker, 1995; Hooker, Penfold, & Evans, 1992; Hooker & Christensen, in preparation).

In particular, the problem of action selection in robots and artificial agents can be solved in the general case only via some sort of selection of actions on the basis of their anticipated consequences. That is, the problem of action selection forces the emergence of interactive representation (Bickhard, 1997b). Once recognized, interactive representation is the framework within which complex representation can be designed and learned, and upon which higher order cognition can be constructed, including rationality and language (Bickhard, 1992, forthcoming, in press; Bickhard & Campbell, 1996a, 1996b; Bickhard & Terveen, 1995). Interactive representation opens the door to genuine higher cognition for artificial agents — just as it did for natural agents (Bickhard, 1992, in press, forthcoming; Bickhard & Campbell, 1992; Bickhard & Terveen, 1995).

Pragmatism. Since the ancient Greeks, mind has been studied primarily as a passive consciousness, processing inputs. The shift to conceiving of mind as an action system is fundamentally due to the philosopher C. S. Peirce (Joas, 1993). Peirce's pragmatism has influenced many people, in spite of its relatively short history, and this literature contains multiple parallels and partial convergences with the interactive model.

Most important are Piaget (e.g., 1954), the later Wittgenstein (1958), Heidegger (1962), and Merleau-Ponty (1962). An action framework, however, does not guarantee viability of the model in detail. The most common subsequent error is a vestigial, sometimes subtle and only implicit, commitment to classical correspondence notions of representation. This can be manifest in arguments *for* such classical notions, even within a dynamic or robotic framework (e.g., Clark, 1997; Clark & Toribio, 1995), or in arguing *against* representation where representation is construed *only* in this classical form (e.g., Brooks, 1991a;

Port & van Gelder, 1995), or in arguing for a version of the *idealism* that correspondence conceptions of representation have so often generated (e.g., Maturana & Varela, 1980, 1987; Varela, Thompson, & Rosch, 1991). These and many other positions are discussed in Bickhard & Terveen (1995).

The shift to an action framework is of fundamental importance, and interactivism shares this move with a growing minority of positions in the literature. The interactive model, however, contributes some essential additions and corrections of its own. System detectable error, for example, is not addressed elsewhere, but is, arguably, essential for an adequate naturalization of representation, whether in organisms or machines.

Conclusion

Interactive representation arises naturally in interactive systems. This is evident for robots, in particular, but it is also apparent in simple interactive organisms. Insofar as interactive representation serves as the fundamental form of representation out of which more complex forms are constructed and derived, this serves to connect the most complex human representational abilities with primitive, emergent, evolutionary precursors.

Interactive representation naturally satisfies a critical meta-epistemological criterion — system detectable error — that must be satisfied by any model of original representation, and that is not satisfied, or even addressed, by standard conceptions of representation. As such, interactive representation offers a promising approach to representation in particular and cognitive problems in general.

References

- Beer, R. D. (1990). *Intelligence as Adaptive Behavior*. Academic.
- Beer, R. D. (1995). Computational and Dynamical Languages for Autonomous Agents. In R. Port, T. J. van Gelder (Eds.) *Mind as Motion: Dynamics, Behavior, and Cognition*. MIT.
- Bickhard, M. H. (1980a). A Model of Developmental and Psychological Processes. *Genetic Psychology Monographs*, 102, 61-116.
- Bickhard, M. H. (1980b). *Cognition, Convention, and Communication*. New York: Praeger.
- Bickhard, M. H. (1988). Piaget on Variation and Selection Models: Structuralism, Logical Necessity, and Interactivism *Human Development*, 31, 274-312.
- Bickhard, M. H. (1992). How Does the Environment Affect the Person? In L. T. Winegar, J. Valsiner (Eds.) *Children's Development within Social Contexts: Metatheory and Theory*. (63-92). Erlbaum.
- Bickhard, M. H. (1993). Representational Content in Humans and Machines. *Journal of Experimental and Theoretical Artificial Intelligence*, 5, 285-333.
- Bickhard, M. H. (1996). Troubles with Computationalism. In W. O'Donohue, R. F. Kitchener (Eds.) *The Philosophy of Psychology*. (173-183). London: Sage.

- Bickhard, M. H. (1997a). Is Cognition an Autonomous Subsystem? In S. O'Neill, P. McKeivitt, E. MacAogain (Eds.). *Two Sciences of Mind*. (115-131). Amsterdam: John Benjamins.
- Bickhard, M. H. (1997b). Emergence of Representation in Autonomous Agents. *Cybernetics and Systems: Special Issue on Epistemological Aspects of Embodied Artificial Intelligence*, 28(6), 489-498.
- Bickhard, M. H. (forthcoming). Critical Principles: On the Negative Side of Rationality. In Herfel, W., Hooker, C. A. (Eds.) *Beyond Ruling Reason: Non-formal Approaches to Rationality*.
- Bickhard, M. H. (in press). Levels of Representationality. *Journal of Experimental and Theoretical Artificial Intelligence*.
- Bickhard, M. H., Campbell, R. L. (1989). Interactivism and Genetic Epistemology. *Archives de Psychologie*, 57(221), 99-121.
- Bickhard, M. H., Campbell, R. L. (1992). Some Foundational Questions Concerning Language Studies: With a Focus on Categorical Grammars and Model Theoretic Possible Worlds Semantics. *Journal of Pragmatics*, 17(5/6), 401-433.
- Bickhard, M. H., Campbell, R. L. (1996a). Topologies of Learning and Development. *New Ideas in Psychology*, 14(2), 111-156.
- Bickhard, M. H., Campbell, R. L. (1996b). Developmental Aspects of Expertise: Rationality and Generalization. *Journal of Experimental and Theoretical Artificial Intelligence*, 8(3/4), 399-417.
- Bickhard, M. H., Richie, D. M. (1983). *On the Nature of Representation: A Case Study of James J. Gibson's Theory of Perception*. New York: Praeger.
- Bickhard, M. H., Terveen, L. (1995). *Foundational Issues in Artificial Intelligence and Cognitive Science — Impasse and Solution*. Amsterdam: Elsevier Scientific.
- Brooks, R. A. (1991a). Intelligence without Representation. *Artificial Intelligence*, 47(1-3), 139-159.
- Brooks, R. A. (1991b). New Approaches to Robotics. *Science*, 253(5025), 1227-1232.
- Brooks, R. A. (1991c). How to Build Complete Creatures Rather than Isolated Cognitive Simulators. In K. VanLehn (Ed.), *Architectures for Intelligence*. (225-239) Erlbaum.
- Brooks, R. A. (1994). Session on Building Cognition. Conference on The Role of Dynamics and Representation in Adaptive Behaviour and Cognition. University of the Basque Country, San Sebastian, Spain, December 9, 1994.
- Campbell, R. L., Bickhard, M. H. (1986). *Knowing Levels and Developmental Stages*. Basel: Karger.
- Campbell, R. L., Bickhard, M. H. (1992). Clearing the Ground: Foundational Questions Once Again. *Journal of Pragmatics*, 17(5/6), 557-602.
- Cherian, S., Troxell, W. O. (1995a). Intelligent behavior in machines emerging from a collection of interactive control structures. *Computational Intelligence*, 11(4), 565-592. Blackwell Publishers. Cambridge, Mass. and Oxford, UK.
- Cherian, S., Troxell, W. O. (1995b). Interactivism: A Functional Model of Representation for Behavior-Based Systems. In Morán, F., Moreno, A., Merelo, J. J., Chacón, P. *Advances in Artificial Life: Proceedings of the Third European Conference on Artificial Life*, Granada, Spain. (691-703). Berlin: Springer.
- Churchland, P. M. (1989). *A Neurocomputational Perspective*. MIT.
- Clark, A. (1997). *Being There*. MIT/Bradford.
- Clark, A., Toribio, J. (1995). Doing without Representing? *Synthese*, 101, 401-431.
- Coffa, J. A. (1991). *The Semantic Tradition from Kant to Carnap*. Cambridge.
- Dretske, F. I. (1981). *Knowledge and the Flow of Information*. Cambridge, MA: MIT.
- Dretske, F. I. (1988). *Explaining Behavior*. MIT.
- Fodor, J. A. (1987). *Psychosemantics*. Cambridge, MA: MIT Press.
- Fodor, J. A. (1990). *A Theory of Content and Other Essays*. MIT.
- Fodor, J. A., Pylyshyn, Z. (1981). How direct is visual perception?: Some reflections on Gibson's ecological approach. *Cognition*, 9, 139-196.
- Hanson, P. P. (1990). *Information, Language, and Cognition*. University of British Columbia Press.
- Heidegger, M. (1962). *Being and Time*. New York: Harper & Row.
- Hooker, C. A. (1995). *Reason, Regulation, and Realism: Towards a Regulatory Systems Theory of Reason and Evolutionary Epistemology*. SUNY.
- Hooker, C. A., Christensen, W. (in preparation). Very Simple Minds.
- Hooker, C. A., Penfold, H. B., Evans, R. J. (1992). Towards a Theory of Cognition Under a New Control Paradigm. *Topoi*, 11, 71-88.
- Joas, H. (1993). American Pragmatism and German Thought: A History of Misunderstandings. In H. Joas *Pragmatism and Social Theory*. (94-121). University of Chicago Press.
- Kuipers, B. J. (1988). The TOUR Model: A Theoretical Definition. From Kuipers, B. J., Levitt, T. Navigation and Mapping in Large-Scale Space. *AI Magazine*, 9(2), 25-43.
- Kuipers, B. J., Byun, Y. (1991). A Robot Exploration and Mapping Strategy Based on a Semantic Hierarchy of Spatial Representations. *Robotics and Autonomous Systems*, 9, 47-63.

- Loewer, B., Rey, G. (1991). *Meaning in Mind: Fodor and his critics*. Blackwell.
- Maes, P. (1990). *Designing Autonomous Agents*. MIT.
- Maturana, H. R., Varela, F. J. (1980). *Autopoiesis and Cognition*. Dordrecht, Holland: Reidel.
- Maturana, H. R., Varela, F. J. (1987). *The Tree of Knowledge*. Boston: New Science Library.
- McClelland, J. L., Rumelhart, D. E. (1986). *Parallel Distributed Processing: Explorations in the Microstructure of Cognition: Vol. 2 Psychological and Biological Models*. Cambridge: MIT.
- Merleau-Ponty, M. (1962). *Phenomenology of perception*. New York: Humanities Press.
- Millikan, R. (1984). *Language, Thought, and Other Biological Categories*. MIT.
- Millikan, R. G. (1993). *White Queen Psychology and Other Essays for Alice*. MIT.
- Nehmzow, U., Smithers, T. (1991). Mapbuilding Using Self-Organizing Networks in "Really Useful Robots." In J.-A. Meyer, S. W. Wilson (Eds.) *From Animals to Animats*. (152-159). MIT.
- Nehmzow, U., Smithers, T. (1992). Using Motor Actions for Location Recognition. In F. J. Varela, P. Bourguine (Eds.) *Toward A Practice of Autonomous Systems*. (96-104). MIT.
- Newell, A. (1980). Physical Symbol Systems. *Cognitive Science*, 4, 135-183.
- Palmer, S. E. (1978). Fundamental aspects of cognitive representation. In E. Rosch & B. B. Lloyd (Eds.), *Cognition and categorization*. Hillsdale, NJ: Erlbaum.
- Piaget, J. (1954). *The Construction of Reality in the Child*. New York: Basic.
- Port, R., van Gelder, T. J. (1995). *Mind as Motion: Dynamics, Behavior, and Cognition*. Cambridge, MA: MIT Press.
- Prem, E. (1995). Grounding and the Entailment Structure in Robots and Artificial Life. In Morán, F., Moreno, A., Merelo, J. J., Chacón, P. *Advances in Artificial Life: Proceedings of the Third European Conference on Artificial Life*, Granada, Spain. (39-51). Berlin: Springer.
- Pylyshyn, Z. (1984). *Computation and Cognition*. MIT.
- Rumelhart, D. E. (1989). The Architecture of Mind: A Connectionist Approach. In M. I. Posner (Ed.) *Foundations of Cognitive Science*. MIT, 133-160.
- Rumelhart, D. E., McClelland, J. L. (1986). *Parallel Distributed Processing. Vol. 1: Foundations*. MIT.
- Smith, B. C. (1987). *The Correspondence Continuum*. Stanford, CA: Center for the Study of Language and Information, CSLI-87-71.
- Smolensky, P. (1988). On the Proper Treatment of Connectionism. *Behavioral and Brain Sciences*, 11, 1-74.
- Stein, L. A. (1994). Imagination and Situated Cognition. *Journal of Experimental and Theoretical Artificial Intelligence*, 6, 393-407.
- Varela, F. J., Thompson, E., Rosch, E. (1991). *The Embodied Mind*. MIT.
- Waltz, D., Feldman, J. A. (1988). *Connectionist Models and Their Implications*. Norwood, NJ: Ablex.
- Wittgenstein, L. (1958). *Philosophical Investigations*. New York: Macmillan.

* Thanks are due to the Henry R. Luce Foundation for support during the preparation of this paper.

Published as: Bickhard, M. H. (1998). Robots and Representations. In R. Pfeifer, B. Blumberg, J.-A. Meyer, S. W. Wilson (Eds.) *From Animals to Animats 5*. (58-63). *Proceedings of the Fifth International Conference on Simulation of Adaptive Behavior*. Zurich, Switzerland, August 17 - 21, 1998. MIT.