

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Whither Representation?

Permalink

<https://escholarship.org/uc/item/87v0x7c2>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 20(0)

Author

Bickhard, Mark H.

Publication Date

1998

Peer reviewed

Whither Representation?

Mark H. Bickhard* (mhb0@lehigh.edu)

Cognitive Science
17 Memorial Drive East
Lehigh University
Bethlehem, PA 18015

Abstract

Cognitive Science is founded on notions of representation, and shifts in models of representation have constituted the major internal revolutions in the field. Symbol System and related conceptions were long dominant, but the frontiers passed first to connectionism and more recently to autonomous agent orientations. In spite of its foundational role, representation has never received a consensual or adequate characterization within cognitive science. This is not surprising, given that millennia of effort in philosophy have also failed to achieve consensus or adequacy, but the situation nevertheless constitutes something of a scandal or impasse in a field in which representation is so central. More recently, workers in dynamicist and autonomous agent approaches have argued that representation is not even a useful notion. I argue that this confusion and impasse concerning representation is due to a fundamental misconception about the nature of representation, and offer an alternative model.

Cognitive Science is founded on notions of representation (Gardner, 1987), and shifts in models of representation have constituted the major internal revolutions in the field. Symbol System and related conceptions (Newell, 1980) were long dominant, but the frontiers passed first to connectionism and more recently to autonomous agent orientations (Bickhard & Terveen, 1995; Clark, 1997; Maes, 1990). In spite of its foundational role, representation has never received a consensual or adequate characterization within cognitive science. This is not surprising, given that millennia of effort in philosophy have also failed to achieve consensus or adequacy, but the situation nevertheless constitutes something of a scandal or impasse in a field in which representation is so central. Standard names for this impasse have even become common, such as "the empty symbol problem."

In recent work in robotics and autonomous agents, the very notion and usefulness of representation has come under attack, with some arguing that representation is an unnecessary and misleading conception that would be better jettisoned. Getting the dynamics right is the fundamental criterion, so the argument goes, and representation is just the wrong notion to be helpful in accomplishing that (Beer, 1990, 1995; Brooks, 1990, 1991a, 1991b, 1991c; Nehmzow & Smithers, 1991, 1992; Pfeifer & Verschure, 1992a, 1992b; Port & van Gelder, 1995). Such a position, of course, is not itself consensual, even among contemporary dynamicists and roboticists, but that such a rejection of representation is one of the major positions in the current frontiers of cognitive science is at least a bit ironic, and it

leads to the obvious question: where *should* models of representation go next?

I have an answer to that question, and an argument for that answer that I will adumbrate in this paper.

A Diagnosis

When proposing a resolution of a problem, it can be useful to have a diagnosis of that problem. I have argued that the impasse with respect to adequate *models* of representation is due to a venerable error in assumptions about the *nature* of representation (Bickhard, 1980, 1993, 1996; Bickhard & Terveen, 1995). There is, in fact, an entire family of related arguments, some of ancient provenance, some discovered just recently. For current purposes, I will focus on one slice through these arguments.

The empty symbol problem refers to the fact that standard symbols are empty of content—representational content. Content is what specifies, for a content understander, what the representation is supposed to represent—that *this* symbol represents dogs, say, and not cats. In that sense, standard symbols are symbols only in the formal sense that there is a way of causally differentiating tokens into types. This is obviously not problematic in the case, for example, of differentiating bit patterns into their respective pattern categories. In not having content, however, they are in fact not genuine representations at all, and, thus, not symbols in any ordinary sense of the term. The problem, then, is one of accounting for representational content (Coffa, 1991).

Furthermore, content is a system relative notion, and not *any* system relative content will due. In particular, human beings have no in-principle difficulty in understanding printed words and sentences as having content, and, for that matter, it is a practical headache, but not an in-principle problem, for humans to understand content for bit patterns. The impasse is with regard to modeling, and constructing, organisms and machines that have content for themselves. The impasse is in providing a naturalistic account of content, one that does not simply pass the problem off to some other content understanding system or content providing source, such as a human user or designer.

Representation is normally considered to be constituted as some sort of correspondence between the representation and whatever it is that is to be represented. These correspondences are variously considered to be causal in nature, or lawful, or informational, or isomorphic, but although such a correspondence approach is almost universal, it has encountered deep problems (Bickhard, 1993; Bickhard & Terveen, 1995; Fodor, 1987, 1990a, 1990b; Loewer & Rey, 1991; Shanon, 1993). One class of

problems is that such correspondences are ubiquitous, and most of them are clearly not representational. Every lawful connection in the universe creates an instance of some such correspondence. The standard assumption is that the class of correspondences has to be narrowed further somehow—that correspondence is necessary, but not sufficient to representation—and further criteria are sought that will refine the class of correspondences into those that are in fact representational.

There is such a criterion, but, unfortunately, it does not solve the problem. The point is that there *are* some representations that are constituted as correspondences, and examining them yields an easy criterion, but that criterion is of no help. An example is Morse code: patterns of dots and dashes correspond to—encode—various characters and numerals, and thereby constitute representations of them. But such encoding systems depend on their users already knowing—representing—the dot and dash patterns, already knowing the characters and numerals, and already knowing the correspondences between them. The problem of representation, thus, is massively presupposed by such encodings. It is not solved. Naturally occurring “encodings” avoid the conventionality of Morse code, but do not add any additional modeling benefits: an instance of smoke “represents” a fire, a neutrino flux “represents” fusion processes in the sun, only to someone who already has all of the relevant representations.

To be sure, new knowledge can be gained via such genuine encodings—that is what telegraphs are for, and that is why people construct neutrino detectors. But what cannot be done is to generate new encodings per se, new representational content. The existence and the use of encodings presupposes all of the basic representational issues. Because standard correspondence approaches to representation assume that correspondences will create some kind of an encoding, sometimes explicitly using the term, sometimes not, I have called such approaches “encodingism”. It is the assumption or presupposition that *all* representation has the character of encodings. It cannot be correct because, although genuine encodings have representational content, they have content only by virtue of the encoding-user already knowing what that content is to be. Encodings cannot be the fundamental form of representation (Bickhard, 1980, 1993; Bickhard & Terveen, 1995).

In order for correspondences to constitute representations, they have to have content, and correspondences do not constitute nor provide content. Content has to come from somewhere else. Conventional encodings, such as Morse code, obtain content via definition. Dot and dash patterns are defined as encoding characters and numerals. But such definitions succeed only because the characters and numerals used to define the dot and dash patterns are themselves already representational. They already have content for the human learning the code, and they provide that content to the defined pattern. Again, content is not accounted for, but merely presumed.

If content has to come from somewhere, and if correspondences per se cannot provide it, then where could it come from? Fodor’s argument for radical innatism

(Fodor, 1981) turns on this issue (Bickhard, 1991) and concludes that all basic contents must be innate—there is no account of how they could be learned or developed. But Fodor’s argument turns on a logical point, that content must be provided from some source, and *that* constraint, if it holds at all, holds just as strongly for evolution as it does for learning and development. If content has to come from somewhere, from where did evolution get it?

The assumption that content must come from somewhere cannot be correct. There was no representational content at the moment of the Big Bang, and there is now. Content has to have emerged. We have no consensual account of content emergence, but that is what is required. The above arguments, plus the much larger family of which they are a part, urges that no such account is possible on standard approaches, and, therefore, that some alternative conception of representation is needed in order to be able to model the emergence of representational content. I will now turn to such an alternative.

Interactive Representation

There is no action in the typical encoding story of representation. If the relevant system engages in interaction, that is secondary to the nature and existence of the representations that might play a role in selecting or guiding those interactions. About a hundred years ago, however, and in contrast to the millennia of correspondence approaches, Peirce introduced the possibility of an action based framework for understanding the mind, including representation (Joas, 1993; Rosenthal, 1983). I will not argue the specifics of Peirce’s models, because I think he got it wrong in the details, but, nevertheless, I am proposing an action based, a pragmatic, model of representation.

Consider an interactive agent faced with the problem of choosing a next action. In simple cases, that “choice” may be a simple trigger based on current system state and current input, but, in more complex cases, mere triggering will not do (Bickhard & Terveen, 1995). The basic point is that triggering requires more reliability about the appropriateness of what is triggered than can be assured in all cases. When triggering is not sufficient, a richer resource for action selection is the indication within the system of the anticipated internal outcomes of various actions, should those actions in fact be executed (under these conditions). With such indications of future potentialities, the system can both select actions with anticipated desirable outcomes, and can check internally to find out if those anticipated outcomes did in fact occur once the interaction is over.

If they did not occur, then several kinds of options can be present: do the interaction again, do some different interaction, engage in problem solving, initiate learning processes. The critical point here is that the failure of the indicated outcome to obtain serves as error information, and error information can guide and initiate many other kinds of processes. Mere triggering does not generate (in itself) error feedback (see Bickhard & Terveen, 1995, for a discussion of built-in error feedback such as pain).

I claim that this much already provides a model of the emergence of primitive representation. Further analysis is

required to show in what sense that is true, and it is certainly a primitive form, so further modeling is required to account for less primitive representation. Nevertheless, such indications of further interactive potentialities and their anticipated outcomes constitutes emergent representation, with emergent content, that is natural in the sense that it does not depend on any outside user or designer or interpreter.

The first point is that the indication of the potentiality for a particular interaction and outcome can be false. The outcome may not occur. So we have the emergence of something with truth value.

Furthermore, that truth value is (fallibilistically) detectable by the system itself by checking on whether or not the indicated outcome(s) did or did not occur (since they are *internal* outcomes, this checking can be a strictly functional matter). That is, the emergent truth value is of and by and for the system (there are additional issues here regarding the nature of function that I will not address in this paper—see Bickhard, 1993; Christensen, 1996; Godfrey-Smith, 1994; Millikan, 1984, 1993). It is not an observer defined truth value.

Still further, such an indication of interactive potentiality is a predication about the current environment. It predicates that this environment is of an appropriate sort that it will yield the indicated outcome(s) if the indicated interaction is engaged in. And, to reiterate, it is a predication that can be false.

Further once again, such a predication has content, emergent—and implicit—content. The predication predicates of the environment that this environment has those properties, *whatever* they are, that will support the obtaining of the indicated outcomes. It is the predication of this content that can be false, and can be falsified.

Here we encounter something novel. Content is standardly assumed to be explicit, not implicit. “Dog” represents dogs because the content of that symbol as an encoding is known to specify dogs. An interactive indication, in contrast, predicates an implicit content, not an explicit content. It predicates that this is a “yields outcome X when interaction Y is engaged in” type of environment. The content is in terms of yielding particular internal outcomes, and the external properties that would make that true are, again, implicit not explicit.

Nevertheless, I claim that this rather primitive, interactive, implicit content, form of representation is foundational to all others. Certainly it is emergent, it yields truth value, and it does have content—all aporia on standard encoding accounts. The promissory notes that must be addressed (though most cannot be here) include: What about explicit content? What about more familiar representations, such as of objects? What about representations of abstractions, such as of numbers? What about other representational phenomena, such as memory, categorization, rationality, episodic memory, and so on? What about language? Most of these questions must be referred elsewhere (Bickhard, 1980, 1992, 1993, forthcoming, in press-a, in press-b; Bickhard & Campbell, 1996; Bickhard & Richie, 1983; Bickhard & Terveen, 1995; Campbell & Bickhard, 1986).

Explicit content, however, is one that can be addressed fairly quickly. Consider the possibility that the system has engaged in an interaction and has obtained the indicated internal outcome. The predication is true—this is in fact an environment of, say, type A. What else might the system know about type A environments? Presumably one thing that it might know is that type A environments are also, say, type B environments. That is, detecting a type A environment indicates that the system is also in a type B environment, so that, if it were to engage in the interactions appropriate to type B, it would obtain the internal outcomes appropriate to type B. There can be predictive—indicative—relationships among the primitive indicative predications.

But note that the general indicative relationship that “type A environments are type B environments” is removed from specific environmental encounters. Such an indication could be present in and for a system even if that system is not and has not for a long time encountered a type A environment. Such indicative relationships, then, are a more general form of knowledge—knowledge about environmental possibilities, not just about (the possibilities of) this current environment.

Of particular relevance here is that the content of such indicative relationships is (more) explicit. The contents of what constitutes a type A environment and what constitutes a type B environment are implicit, but the indication that type A environments (whatever they are) *are* type B environments is explicit. Explicit, but still potentially false and falsifiable, so we still have emergent truth value, and we now have explicitness.

Such indicative relationships can iterate: type A indicates type B, which, in turn, indicates type C, and so on. And they can branch: type A might indicate all of B and C and D as possibilities. As such, they can link into potentially vast webs of representations about what sorts of dependencies obtain in possible environments, and it is in the possibilities of such webs that the answers lie to such questions as “what about representations of objects?”

A simple manipulable object, such as a toy block, will offer many interactive potentialities, such as visual scans from multiple perspectives, manipulations of various kinds, dropping, throwing, chewing, and so on. Furthermore, any one of these interactions suffices to indicate the potentialities of all the others, perhaps with intermediate interactions, such as rotating a toy block to obtain a particular visual scan. Still further, the entire pattern of interactive potentialities will remain invariant under a large class of possible interactions, such as transporting, putting in the toy box, walking away, and so on—again perhaps with intermediate interactions such as returning to the toy box—but it will not remain invariant under others, such as burning or crushing. Such internally mutually reachable, and invariant-under-classes-of-basic-physical-actions, webs of interactive indicative relationships is, I claim, the basic form of simple object representations. It is a Piagetian model of object representation (Piaget, 1954), and I would suggest similarly Piagetian representational models for many other kinds of representation, such as of numbers (Piaget’s model is also a pragmatist, action based, model—

thus the ability to borrow such Piagetian submodels—but the interactive model also diverges from Piaget in many crucial ways; see Bickhard, 1988; Bickhard & Campbell, 1989; Campbell & Bickhard, 1986).

The interactive model of representation, then, arguably accounts for emergent content and truth value, and provides resources for modeling other representational phenomena such as objects and numbers. It is a plausible candidate for a foundational form of representation.

Consequences

The interactive model of representation does not look much like standard correspondence models. Representations are commonly supposed to be produced by the processing of (transduced) inputs, while interactive representation is emergent in certain kinds of interactive system organization. That contrast does raise the question, however, of what role, if any, input processing might play in the interactive approach.

Recall that actually interacting with an environment and arriving at some internal outcome serves to differentiate or classify that environment as being of, say, type A. On the basis of such a differentiation, further indications, such as of the possibilities of interactions of type Q, R, and S, might be evoked, and the full interactive model come into play. My focus here, however, is the initial differentiation.

A simple form of interaction, and, therefore, a simple form of differentiation, is a form in which there is no output—there is only input and the processing of inputs. That is, it might suffice to arrive at a differentiation that the current environment is of type A to only process inputs, with no full *interaction*. But it is precisely, in standard models, the results of such input processing that are assumed to constitute encoded representations by virtue of the correspondences that are established (factually) in any such differentiation (there will always be a factual correspondence with whatever has been factually differentiated). The difference in the two sorts of models, at least to this point, is that encoding models assume that somehow content has magically come into being in such correspondence, while the interactive model claims that, however necessary and useful differentiations are, more is required in order to get content. In particular, indications of further potentialities, perhaps based on such differentiations, are required in order to constitute content.

So, the facts of sensory input processing pose no problems for the interactive model. Those are the basic form of simple differentiations of environmental types. But such processing is understood quite differently—as part of an entire interactive process, with representation emergent in the future looking anticipations, not in the backward looking correspondences.

On the interactive view, then, standard approaches have it half right. They focus on differentiations and presume that that is everything. They miss the full interactive perspective, with its future oriented anticipations. Such a point holds for all correspondence models of representation, including symbol system models, connectionist models, and many others (Bickhard & Terveen, 1995).

On the interactive view, genuine representation can be emergent only in interactive systems, not in passive input processing only systems. Standard computer models and standard connectionist models, then, cannot have genuine representations. On the other hand, interactive representation may well be emergent in interactive systems even if it is not recognized as such by the designers of the systems. Interactive anticipatory representations are a natural solution to a natural problem of action selection that cannot be avoided in complex interactive systems.

The controversy about whether or not representation is a useful notion for understanding and modeling autonomous agents, then, looks quite different from the perspective of interactivism. First, interactivism agrees that standard encodingist conceptions of representation are inappropriate, unnecessary, and can be badly misleading if taken seriously, in the design and modeling of interactive agents. On the other hand, a genuine form of representation—interactive representation—is unavoidable in most interactive systems, and they can be better understood and better designed when that is recognized, and its potential power exploited (Bickhard, 1997a, 1997b; Cherian & Troxell, 1995a, 1995b; Christensen, Collier, Hooker, in preparation; Hooker, 1996). So, representation *is* an essential notion in robotics and autonomous agents, just not standard encodingist notions of representation.

More deeply, representation *cannot* be accounted for in input processing terms, but it emerges naturally, and in initially quite simple forms, in the interactive view. Simple interactive representations emerge in quite primitive organisms, and became progressively more complex and differentiated into multiple kinds throughout evolutionary history, but the basic interactive character remains through all levels of evolution, and into the design of artificial agents as well.

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*. (121-147). Cambridge, MA: MIT Press.
- Bickhard, M. H. (1980). *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. (1991). The Import of Fodor's Anticonstructivist Arguments. In L. Steffe (Ed.) *Epistemological Foundations of Mathematical Experience*. (14-25). New York: Springer-Verlag.
- Bickhard, M. H. (1992). How Does the Environment Affect the Person? In L. T. Winegar, J. Valsiner (Eds.) *Children's Development within Social Context: Metatheory and Theory*. (63-92). Hillsdale, NJ: 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). 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. (1997b). Is Cognition an Autonomous Subsystem? In S. Ó Nualláin, P. McKeivitt, E. MacAogáin *Two Sciences of Mind: Readings in Cognitive Science and Consciousness*. (115-131). Amsterdam: John Benjamins.
- 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-a). Levels of Representationality. *Journal of Experimental and Theoretical Artificial Intelligence*.
- Bickhard, M. H. (in press-b). Dynamic Representing and Representational Dynamics. In E. Dietrich, A. Markman (Eds.) *Cognitive Dynamics: Conceptual Change in Humans and Machines*. MIT.
- Bickhard, M. H., Campbell, R. L. (1989). Interactivism and Genetic Epistemology. *Archives de Psychologie*, 57(221), 99-121.
- Bickhard, M. H., Campbell, R. L. (1996). Topologies of Learning and Development. *New Ideas in Psychology*, 14(2), 111-156.
- 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. (1990). Elephants don't Play Chess. In P. Maes (Ed.) *Designing Autonomous Agents*. (3-15). Cambridge, MA: MIT Press.
- Brooks, R. A. (1991a). Challenges for Complete Creature Architectures. In J.-A. Meyer, S. W. Wilson (Eds.) *From Animals to Animats*. (434-443). Cambridge, MA: MIT Press.
- Brooks, R. A. (1991b). How to Build Complete Creatures Rather than Isolated Cognitive Simulators. In K. VanLehn (Ed.) *Architectures for Intelligence*. (225-239). Hillsdale, NJ: Erlbaum.
- Brooks, R. A. (1991c). Intelligence without Representation. *Artificial Intelligence*, 47(1-3), 139-159.
- Brooks, R. A. (1991d). New Approaches to Robotics. *Science*, 253(5025), 1227-1232.
- Campbell, R. L., Bickhard, M. H. (1986). *Knowing Levels and Developmental Stages*. Basel: Karger.
- 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.
- Christensen, W. D. (1996). A complex systems theory of teleology. *Biology and Philosophy*, 11, 301-320.
- Christensen, W. D., Collier, J. D., Hooker, C. A. (in preparation). Autonomy, Adaptiveness, Anticipation: Towards autonomy-theoretic foundations for life and intelligence in complex adaptive self-organising systems.
- Clark, A. (1997). *Being There*. MIT/Bradford.
- Coffa, J. A. (1991). *The Semantic Tradition from Kant to Carnap*. Cambridge.
- Fodor, J. A. (1981). The present status of the innateness controversy. In J. Fodor *RePresentations* (257-316). Cambridge: MIT Press.
- Fodor, J. A. (1987). *Psychosemantics*. Cambridge, MA: MIT Press.
- Fodor, J. A. (1990a). *A Theory of Content*. Cambridge, MA: MIT Press.
- Fodor, J. A. (1990b). Information and Representation. In P. P. Hanson (Ed.) *Information, Language, and Cognition*. (175-190). Vancouver: University of British Columbia Press.
- Gardner, H. (1987). *The Mind's New Science*. New York: Basic.
- Godfrey-Smith, P. (1994). A Modern History Theory of Functions. *Nous*, 28(3), 344-362.
- Hooker, C. A. (1996). Toward a naturalised cognitive science. In R. Kitchener and W. O'Donohue (Eds.) *Psychology and Philosophy*. (184-206). London: Sage.
- 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.
- Loewer, B., Rey, G. (1991). *Meaning in Mind: Fodor and his critics*. Oxford: Blackwell.
- Maes, P. (1990). *Designing Autonomous Agents*. Cambridge, MA: MIT Press.
- 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.
- 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). Cambridge, MA: MIT Press.
- 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). Cambridge, MA: MIT Press.
- Newell, A. (1980). Physical Symbol Systems. *Cognitive Science*, 4, 135-183.
- Pfeifer, R., Verschure, P. (1992a). Beyond Rationalism: Symbols, Patterns and Behavior. *Connection Science*, 4(3/4), 313-325.
- Pfeifer, R., Verschure, P. (1992b). Distributed Adaptive Control: A Paradigm for Designing Autonomous Agents. In F. J. Varela, P. Bourguine (Eds.) *Toward A Practice of Autonomous Systems*. (21-30). Cambridge, MA: MIT Press.

- 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.
- Rosenthal, S. B. (1983). Meaning as Habit: Some Systematic Implications of Peirce's Pragmatism. In E. Freeman (Ed.) *The Relevance of Charles Peirce*. (312-327). La Salle, IL: Monist.
- Shanon, B. (1993). *The Representational and the Presentational*. Hertfordshire, England: Harvester Wheatsheaf.

*Thanks are due to the Henry R. Luce Foundation