Genuine Representation in Artificial Systems

Mark H. Bickhard

Cognitive Science, LeHigh University
Bethlehem, PA 18015, USA
mhb0@lehigh.edu

Abstract. The greatest challenge to a model of the emergence of representation is that of the normativity of representations: the possibility of being true or false. The strongest version of that challenge is to be able to account for *system detectable*representational error, as is used in error guided behavior or error guided learning. No model in the standard literature, and, arguably, no spectator model of any kind, can account for it. Genuine representation, however, with content and truth value - system detectable truth value - emerges in the selection of actions and interactions in autonomous agents, whether natural or artificial, organisms or robots. Representation is most fundamentally of future potentialities for interaction, rather than of past encounters as standard approaches would have it. Representation is intrinsic to agents, not to passive spectators. The fundamental aspirations of Artificial Intelligence to create genuine artificial minds will be met in robotics.
Keywords. Representation, pragmatism, robots, agents, emergence.

Genuine representation and cognition are not possible in standard symbol manipulation or connectionist architectures (Campbell, this session). There is a very large family of multifarious reasons for the ultimate failures of such approaches, but those reasons for failure are united as manifestations of an underlying false assumption about the nature of representation. Do these foundational failures mean that representation in artificial systems is impossible? No.

But what is required is a new model of representation, one that avoids the aporia of standard approaches, and a different kind of system architecture that fits that model, a type of system that, in fact, is at the heart of one of the most active frontiers of research today. I will sketch such a model of representation—called *interactivism*—and will situate it in three perspectives: 1) a historical perspective, 2) a strengthened meta-theoretical requirement for a valid model of representation—one that is not addressed in contemporary literature, and 3) a diagnosis of one of the key ubiquitous conceptual conflations that has made it so difficult to resolve the nature of representation. The historical perspective has to do with the development of pragmatism; the meta-theoretical requirement is for models of representation to be able to account for system detectable representational error; and the diagnosis is to point out a ubiquitous conflation between epistemic contact and epistemic content.

1 Pragmatism

Studies of the mind and of mental phenomena have classically taken consciousness as the primary framework, and have accepted a passive model of vision as the primary analogy for understanding consciousness. 'Mind as spectator' is the guiding framework within which representation and cognition are approached in this tradition.

'Mind as spectator' has been the dominant framework for attempting to understand the mind since the ancient Greeks, and it is still dominant. It yields models of representation as produced by the processing of passively received inputs In such models, the informational or causal connection between some element in the mind—or bit pattern in a computer, or activation-vector in a connectionist vector space—and something in the world is taken to constitute a representational relationship. That element is supposed to represent that something in the world by virtue of its informational or causal relationship to that something: my mental representation of the desk in front of me is taken to represent that desk by virtue of the causal chain of light bouncing off the desk, entering my retina, stimulating certain nervous activity, etc., etc., ultimately yielding the generation of the CNS state that is supposed to constitute that mental representation.

But, as discussed by Robert Campbell (this session) a moment ago: 1) Informational causal relationships are ubiquitous in the universe, and certainly not all of them are representational. 2) Any element in such a relationship with one property or thing in the world will also be in such a relationship with an unbounded number of other properties or things in the world, and at most one of them can be represented. 3) No such element bears any intrinsic announcement about what causal or informational relationships it might be involved in, so no such element can be taken as a representation except by some agent that already knows about that element and its crucial representational relationship—such as in Morse code, or using neutrino flux to 'observe' nuclear processes in the sun—but requiring such knowledge about purported representational elements and their representational relationships in order for them to *be* representations at all introduces a fatal circularity. And so on, and so on. This approach simply does not work.

About a century ago, however, C. S. Peirce introduced a different framework for attempting to understand mind and mental phenomena—pragmatism (Joas, 1993; Rosenthal, 1983). Pragmatism takes action and interaction, rather than consciousness, as the crucial framework. It is inherently not a passive conception of mind. Minds cannot be primarily spectators. Representation cannot be simply passive reception, unconnected with action.

The interactive model of representation that I will outline belongs to this pragmatist tradition. Pragmatism has a much shorter history than classical approaches, and much less has been done within its framework. For much of this century, in fact, pragmatism suffered the fate of being mis-interpreted as nothing more than a simple minded precursor to verificationism, and, of course, everyone knows that *that* approach failed (e.g., Fodor, 1998). Nevertheless, there are a number of important pragmatist models of representation, starting with Peirce himself and including, for example, Jean Piaget. I do not think that these models got it right, though there is much of importance to learn from them. Here I will focus on the interactive model itself, since I have explored the details of the similarities and differences with other pragmatist oriented models elsewhere (e.g., Bickhard, 1988; Bickhard & Campbell, 1989).

Pragmatism, then, forms the outer framework for the interactive model. Already at this level we can see that classical architectures will not suffice: Neither symbol manipulationist nor connectionist models of representation grant any essential role to action or interaction in their conceptions of representation. They are both committed to the passive spectator model regarding representation and cognition. They can be wedded to action systems, but, though those action systems may *make use* of the purported passively-generated representations, action has nothing to do with the purported representational nature of elements in those systems at all. Interactivism requires genuine agents in the world in order for genuine representation to emerge. For representation in artificial systems, it requires artificial interactive systems—robots, or autonomous agents (Bickhard, 1982, 1993; Bickhard & Terveen, 1995).

2 Representational Error

One important perspective for understanding the problems of standard approaches to representation, and, therefore, the advantages of a pragmatist—action based—approach, derives from considerations of representational error. In fact, one of the many problems vexing standard conceptions of representation is that of accounting for the very *possibility* of representational error. If, according an informational model of representation, the required informational relationship exists, then the representation exists, and that representation is correct. If, on the other hand, that crucial informational relationship does not exist, then the representation does not exist, and, therefore, *it cannot be incorrect.* The possibility of representational error proves to be difficult to account for in these models.

There are attempts to do so. Jerry Fodor's attempt turns on the notion of asymmetric dependency: a COW element in the mind represents cows, and not horses on dark nights, because, even though COW can be evoked by either cows or by horses on dark nights, the possibility of it being evoked by horses on dark nights is dependent on the possibility of COW being evoked by cows, and that dependency is asymmetric. It is not reciprocated. Horses on dark nights would not evoke COW unless cows did, but cows would evoke COW even if horses on dark nights never did.

There are rather simple counterexamples to this model (Campbell, this session; Bickhard, 1993), and much controversy surrounding it (along with all of its alternatives in the contemporary literature) (e.g., Loewer & Rey, 1991). What I wish to focus on is that, even if Fodor's model were accepted, it would at best characterize representational error in a way that is determinable, if at all, only from the perspective of some *external observer* of the epistemic situation—some observer who could determine the relevant evocation relationships and their dependencies, asymmetric or otherwise, on each other. In particular, neither Fodor's model, nor any other in the "spectator" literature, can account for *system detectable* representational error—error detectable by the system itself. Fodor attempts, unsuccessfully, to account for the mere possibility of error, but does not even address *system detectable* error.

Since at least some minds—e.g., human minds—at least some of the time can detect such representational error for themselves, this criterion of system detectable error would seem to be critical to any acceptable model of representation. Neither error guided behavior nor error guided learning would be possible without it. Furthermore, since the difference between error and no error is the fundamental normative issue regarding representation, any model that can account for error only from some *external perspective* on the agent in question cannot account for that error in any naturalistic way. If error can be defined only for an observer, then that observer's error cannot be naturalistically modeled. To attempt to do so simply requires some second observer of the first one, and then a third for the second one, and so on—a classic infinite regress.

I propose, then, that system detectable error be added to the criteria that an acceptable model of representation must satisfy. No 'spectator' model in the literature even attempts it. Nevertheless, it provides an introduction to a pragmatist model that easily accounts for the possibility of system detectable representational error—to the interactive model.

3 Action Selection

I will approach this issue of error and error detection by way of a broader problem that faces all agents of more than minimal complexity—the problem of selecting actions. In the simplest cases, action can be selected by simple triggering: current input plus current internal state triggers some appropriate—or hopefully appropriate—action. This may suffice, perhaps, for paramecia, but more complex agents will, in general, face situations in which more than one action is possible and a selection needs to be made (a before-the-action selection problem), or in which the appropriateness of an action and its outcomes is not so 'obvious' and reliable so that evolution can simply make the selection ahead of time and install a triggering device in the organism itself (an after-the-action evaluation problem). Instead, some means needs to be available for determining what actions or interactions might be possible now, some means for selecting among them, and some means for evaluating that selection, and how desirable that interaction turned out to be.

3.1 How to Select

I will skip over, for the time being, the issue of how a system can determine what actions are possible now, and focus on the later two issues: how to select actions, and how to evaluate the selection once they are complete. A *prima facie* answer to 'how to select actions' is in terms of their consequences: pick those available actions with the consequences best suited to current goals. This answer, I suggest, contains the right intuitions, but, if those consequences are external, then they need to be represented, and that introduces a circularity if we are trying to model representation.

If those consequences are strictly *internal*, however, then they do not need to be represented. If a system can indicate possible actions, perhaps by pointers to the subroutines that would engage in those actions, and can indicate their anticipated *internal* outcomes, perhaps by pointers to the internal states that would be expected upon completion, then those indicated internal outcomes can be used to select actions, and there is no circularity if this somehow provides the framework for understanding representation.

Such selections, of course, will be with respect to current goals, and here is another potential circularity. If goals need to be representational, then the circularity of defining representation in terms of representation bites again. But goals too can be internal, not external, and can be functionally assessed, not representationally assessed. Goals need only be switches that switch one way—out of the current subsystem—if some internal set-point has been attained, e.g. blood sugar level, and switch another way—back into the subsystem—if the set-point has not been attained. Continue eating if blood sugar level is not yet above threshold, and do something else if it is above threshold. That is, goals need only be, e.g. set-points for functional servomechanisms, or TOTEs (Miller, Galanter & Pribram, 1960). (Strictly, explicit goals are not required at all: this architecture of pointers and goals and subroutines illustrates a possibility, it is not a necessity; Bickhard & Terveen, 1995.) Functional set-points of this sort do not need to be represented in order to be functionally available. Nothing has to *represent* blood sugar level in order to differentially switch control depending on blood sugar level.

3.2 How to Evaluate

The same means that provides for action selection—indication of anticipated outcomes—also provides the possibility of evaluation of those anticipations. In particular, a system can functionally determine whether or not an indicated internal outcome state actually obtains after an interaction is complete. If it does not, then the system can switch control appropriately, perhaps engaging in the interaction again, perhaps switching to an alternative interaction, perhaps invoking learning processes if such are available.

It is possible that the anticipated outcomes do occur, but the goal is not attained or furthered. In this case, the fault lies with the selection process, not with the anticipations *per se*. That can be a critical consideration, of course, but my current concerns focus on the anticipations.

4 Interactive Representation

Such interactive anticipations, I claim, constitute representation. Simple versions constitute a very primitive form of representation, but representation it is nevertheless, and it is arguable that all other forms of representation are derived from this one (Bickhard, 1980, 1993). Note that, if such interactive representations do constitute representation, then I have shown why representation emerges naturally both in the evolution of biological agents, and in the development of artificial agents: interactive representation, in the form of interaction anticipations, is required for any even minimally sophisticated action selection. Action selection is the reason why representation emerges naturally—in the pragmatic view.

4.1 Predications

The first perspective on the sense in which such anticipations constitute representation is that they constitute implicit predications about the current environment. They predicate that the current environment is one in which the indicated interaction will in fact yield the indicated outcomes. If the interaction is an X type of interaction, and the indicated outcomes are Y, then the indication constitutes a predication that this environment is appropriate to X interactions with Y outcomes. The frog, for example, upon seeing a fly, may have an indication that its current environment is an environment appropriate for tongue flicking and eating. Primitive interactive representation, then, emerges as predications about the interactive nature of the current environment.

4.2 System Detectable Error

Furthermore, that predication can be false, *and its falsity can be detected by the system itself.* We have inherent system detectable error. This point is so simple and so obvious that it deserves to be repeated that the property of system detectable error is not addressed by, and is not attainable within, any of the "spectator" models of representation that have been offered over the last two or three thousand years. As simple as it is, a model of system detectable error is completely novel.

4.3 Representational Content

If interactive anticipations—predications of interactive properties—are the emergent form of representation, what is their representational content? Content is what determines what the representation is supposed to be a representation *of.* Content is what determines that X is a representation of a ball, rather than a cat. Content is what determines that COW is supposed to represent a cow rather than a horse on a dark night, even if it is currently being mistakenly evoked by a causal interaction with a horse on a dark night.

On one level of analysis, the content is simply the predication: this is an X-interaction kind of environment. But this form of content is highly system dependent (in fact, the interactive model entails that all representation is inherently deictic and indexical, and that more system invariant forms of representation are attainments of some organisms by virtue of certain kinds of constructions of more complex forms of interactive representation—Bickhard, 1993, 1996, in press). The environmental properties that would support the indicated interactions and interaction outcomes are not *explicit* in the predication itself. They are *implicit*: they are whatever the environmental properties are that would support the indications.

This primitive implicitness of content is one of the unusual and powerful properties of interactive representation. Elsewhere, I argue, for example, that it provides a solution to the frame problems (Bickhard & Terveen, 1995). A form of explicit content can be constructed, however: if an X interaction is engaged in and the appropriate outcomes do occur, then perhaps a Y interaction would then be possible. Such conditionalized interaction possibilities can also be indicated by, for example, the appropriate iterations of pointer relationships. A given interaction outcome might also indicate the potentiality of *many* further interactions, not just one, so interactive indications can both branch and iterate. Interactive indications, therefore, can form potentially complex webs of conditional anticipations. I argue that such complex webs, of certain special kinds, underlie our more familiar higher-primate kinds of interactions, such as of objects (Bickhard, 1993, in press).

For my current purposes, the important point is that such conditionalizations constitute *explicit* content: it is explicit in such a conditionalization that an X environment is also (conditional on X) a Y environment. What properties constitute X and Y environments remain implicit, but the *relationship* between X and Y environments is explicit. The frog's implicitly defined F type environment may in fact be encountered mostly only when there is a fly about, and the implicitly defined E environment only when eating is possible, but it will be *explicit* for the frog that F environments are also E environments. Interactive representation, then, accounts for representational content—implicit at its base (with important consequences), but explicit in certain forms of constructed representations.

4.4 Evoking Interactive Representations

At this point I can address the issue that was postponed above: how to evoke interactive representations, how to determine what interactions might be possible now. In fact, I have already done so: A Y interaction, for example, is possible if an X interaction has just successfully completed. In general, the actual outcomes of an interaction contain differentiating information about the current environment, and the differentiation of an outcome-A type of environment might be just what is needed to indicate that Y and Z interactions are now possible.

4.5 Spectator Models

There is an important point to be noted here. The differentiation of, say, an A type of environment might be obtained by engaging in some full and perhaps extended and complex interaction. But it might also, in some cases, be obtained by a much simpler interaction, one with no outputs, and, therefore, not a full interaction at all. That is, it might be obtained by the passive processing of inputs.

If such a passive input processing differentiation occurs within what is commonly known as a sensory system, such as the visual system, the overwhelming likelihood is that the correspondence that is thereby created between the internal differentiating element and the external property or thing that is in fact involved in the differentiation will be construed as a representational correspondence—a sensory encoding (Carlson, 1986; cf. Bickhard & Richie, 1983). Differentiations are standardly construed as representations of that which has been differentiated. But that way lies the infinitude of problematics of all spectator approaches.

Those are avoided simply by noting that: 1) differentiations do not inherently represent that which they have differentiated, 2) differentiations can nevertheless be useful for invoking indications of further interactive possibilities, and 3) it is in such future oriented indications that representational content and truth value emerges, not in the past oriented differentiations. The frog does not have to represent *either* flies *or* BBs, but can indicate, can represent, the potentiality for tongue flicking andeating on the basis of differentiating flies or BBs (Bickhard & Campbell, 1996b). Differentiations play an essential role in the interactive model, without requiring that they be mistakenly construed as being representational *per se*.

5 The Spectator Conflation

I can now return to the third situating perspective mentioned at the beginning of this paper: a standard conflation that has contributed to the difficulty of resolving the problems of the nature of representation. In the interactive model, past oriented differentiations can ground the setting up of future oriented indications of interactive potentialities. The differentiations provide an epistemic *contact* with the current environment; the interactive indications provide a *content* about that current environment. Contact and content are two of the fundamental aspects that a model of representation must account for (Campbell, this session). In the interactive model, they are accounted for by distinct, though deeply related, parts of the model. *In spectator models, they are equated.*

That is the basic conflation: If the visual system has differentiating *contact with* a fly, then it is presumed to have representational content of that fly. In the spectator view, there really is no other possibility: what you passively observe *must* be both your contact with the world, and your content about that world. That way, the standard way for millennia, lies conceptual doom.

By distinguishing between epistemic contact and representational content, with one being past oriented and the other future oriented, the interactive model avoids this conflation, and thereby avoids the unbounded aporia which it generates. Differentiations don't need to represent in order to suffice as differentiations, and interactions don't need to have already occurred in order to be indicated as possible. Indicated interactive contents don't need to be contents about what has been differentiated, they only need to be appropriately *based* on such differentiations. Further, determining what has been differentiated is a different task from differentiating it in the first place, and taking into account an interactive possibility is a different process from engaging in that interaction.

Still further, determining whether some representation of what has been differentiated is *correct* directly yields the error problem. The best that the system can do is to engage in the differentiating process again, but that yields a circular check on correctness. (This argument that the check on representational accuracy is circular—all you can do is to check again—has been a core argument of radical skepticism for a *very* long time, and has never been refuted from within the spectator view; Hookway,1992; Rescher, 1980; Sanches, 1988.) If your intuitions say that you could check *consequences*, rather than the differentiated property or object itself, then—so long as you recognize that those anticipated consequences cannot themselves require representation (on pain of a different circularity) and, therefore, must be internal—then you are correct, and we have once again arrived at the interactive model of representation.

6 Conclusion

Genuine representation is possible in artificial systems. But representation must be recognized as fundamentally interactive in nature, not passive, not merely the result of spectatoring, and the artificial systems must be interactive in nature, not passive. They must be agents, robots. (There is an additional requirement that those systems be, in a certain critical way, far from thermodynamic equilibrium, but I will not address that point here; Bickhard, 1993.)

The shift in the model of representation that is involved overturns a millennia old tradition in favor of one a mere century old. The shift in architecture involves recognizing the limitations of passive systems, and the potentialities of what was once considered as 'merely' an engineering adjunct to Artificial Intelligence (Bickhard, 1982, 1997). Once, but no more: some of the most exciting work in Artificial Intelligence today is exploring autonomous agent and robotic design. The issue of representation is not resolved in this literature (Beer, 1990, 1995a, 1995b; Bickhard & Terveen, 1995; Clark, 1997; Clark & Toribio, 1995; Port & van Gelder, 1995), but, if the interactive model is correct, genuine representation emerges naturally and necessarily in complex interactive agents—they cannot avoid the action selection problem—whether or not the designers recognize that fact.

Brooks , one of the leaders of contemporary robotics (1991a, 1991b, 1991c), is a good example. He argues, correctly in my judgment, that his robots do not contain any representations—*at least none that look like standard models of representation*. That is, they do not contain any standard manipulable elements that are supposed to be representational by virtue of correspondence, though he acknowledges that there might be some non-standard sense in which representation is involved (Brooks, 1991a). On the other hand, robots from Brooks' lab that can, for example, travel around rooms with obstacles and find best paths and shortcuts do so precisely with indications of potential actions of the robot and of their outcomes (Stein, 1994). Brooks even uses the vocabulary of anticipations in discussing them (Brooks, 1994). That is, Brooks' robots are realizations of interactive representational agents, even though he does not in general recognize that fact (Bickhard & Terveen, 1995).

Recognizing that action selection yields emergent representation provides a powerful new perspective on system design and modeling, including of biological organisms, including human beings (e.g., Bickhard, 1992, in press; Bickhard & Campbell, 1996a, 1996b; Cherian & Troxell, 1995a, 1995b; Christensen, Collier, Hooker, in preparation; Hooker, 1995, 1996). The conceptual shifts are required in order to account for, and to design, genuine representation.

7 References

Beer, R. D. (1990). *Intelligence as Adaptive Behavior.* Academic.

Beer, R. D. (1995a). 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.

Beer, R. D. (1995b). A Dynamical Systems Perspective on Agent-Environment Interaction. *Artificial Intelligence*, 73(1/2), 173.

Bickhard, M. H. (1980). *Cognition, Convention, and Communication.* New York: Praeger.

Bickhard, M. H. (1982). Automata Theory, Artificial Intelligence, and Genetic Epistemology. *Revue Internationale de Philosophie*, 36(142-143),549-566.

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 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.(1997). Is Cognition an Autonomous Subsystem? In S. Ó Nuall´in, P. Mc Kevitt, E. MacAog´in, *Two Sciences of Mind: Readings in Cognitive Science and Consciousness.* (115-131). Amsterdam: John Benjamins.

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. (1996a). Developmental Aspects of Expertise: Rationality and Generalization. *Journal of Experimental and Theoretical Artificial Intelligence*, 8(3/4), 399-417.

Bickhard, M. H., Campbell, R. L. (1996b). 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. (1991a). Intelligence without Representation. *Artificial Intelligence*, 47(1-3), 139-159.

Brooks, R. A. (1991b). Challenges for Complete Creature Architectures. In J.-A. Meyer, S. W. Wilson (Eds.) *From Animals to Animats.* (434-443). MIT Press.

Brooks, R. A. (1991c). New Approaches to Robotics. *Science*, 253(5025), 1227-1232.

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. (this session).

Carlson, N. R. (1986). *Physiology of Behavior.* Boston: Allyn and Bacon.

Cherian, S., Troxell, W. O. (1995). 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. (1995). Interactivism: A Functional Model of Representation for Behavior-Based Systems. In Moran, F., Moreno, A., Merelo, J. J., Chacon, P. *Advances in Artificial Life: Proceedings of the Third European Conference on Artificial Life*, Granada, Spain. (691-703). Berlin: Springer.

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.

Clark, A., Toribio, J. (1995). Doing without Representing? *Synthese*, 101, 401-431.

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

Hooker, C.A. (1995). *Reason, Regulation and Realism: Toward a Naturalistic, Regulatory Systems Theory of Reason.* Albany, N.Y.: State University of New York Press.

Hooker, C. A. (1996). Toward a naturalised cognitive science. In R. Kitchener and W O' Donohue (Eds.) *Psychology and Philosophy.* (184-206). London: Sage.

Hookway, C. (1992). *Scepticism.* London: Routledge.

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.
Miller, G. A., Galanter, E., & Pribram, K. H. (1960). *Plans and the Structure of Behavior.* New York: Holt, Reinhart, and Winston.
Port, R., van Gelder, T. J. (1995). *Mind as Motion: Dynamics, Behavior, and Cognition.* Cambridge, MA: MIT Press.
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.* (312-327). La Salle, IL: Monist.
Sanches, F. (1988/1581). *That Nothing is Known.* Cambridge.
Stein, L. A. (1994). Imagination and Situated Cognition. *Journal of Experimental and Theoretical Artificial Intelligence*, 6, 393-407.