

The Biological Foundations of Cognitive Science

Mark H. Bickhard

Mark H. Bickhard
Cognitive Science
17 Memorial Drive East
Lehigh University
Bethlehem, PA 18015
610-758-3633
mhb0@lehigh.edu
mark.bickhard@lehigh.edu
<http://www.lehigh.edu/~mhb0/mhb0.html>

The Biological Foundations of Cognitive Science

Mark H. Bickhard
Abstract

Computationalism suggests that mental properties and processes can be modeled independently of most specifics about the brain — mind is like a computer program, and program properties do not depend on details of the computers that execute them. But computationalism suffers from fatal conceptual problems concerning the nature of representation. An alternative model of representation, called *interactivism*, is outlined and its implications for the biological foundations of Cognitive Science are explored. Interactivism profoundly reconfigures our conceptions of those biological foundations.

The Biological Foundations of Cognitive Science

Mark H. Bickhard

One of the founding assumptions that Cognitive Science extracted from the computer model was that mental phenomena can be modeled relatively independently of the brain. Mental phenomena are, in this view, assumed to be akin to the execution of a computer program, and program properties are independent of the specifics of the computer that is running that program. Similarly, mental properties are assumed to be relatively independent of the brain in which they are realized. Furthermore, the body is “just” a robot body housing that computer, and, perhaps, incorporating the sensory transducers that feed inputs into the computer. Recognition that the program/computer distinction might provide a model for the mind/brain relationship was, in fact, one of the early excitements of the field.

The re-emergence of connectionist models in the 1980s was accompanied by claims of greater faithfulness to actual brain organization and process compared to symbol manipulation computationalist models. But the architectural and vector process differences between computationalism and connectionism do not alter the common underlying assumption that, while the brain may realize the relevant processes to constitute mind, the critical properties of those processes may be still be modeled independently of neural level details in the brain.

More recently, dynamic systems and autonomous agent approaches have argued for the importance not only of specifics about the brain, but of details concerning the body as well. Cognition, in this view, is inherently embodied, and cannot be understood in the disembodied forms provided by either computationalism or connectionism (Beer, 1995; Maes, 1990).

These dynamic and agent approaches, however, have also reintroduced basic questions about what the aim of modeling in Cognitive Science should be. In particular, many have argued that mental phenomena such as representation are simply not relevant to either the design or the modeling of embodied autonomous agents. Representation is a vestigial notion from computationalism that should be rejected (Brooks, 1991). Dynamics is everything, so it is held, and getting the dynamics right, whether in design or in modeling, is the only relevant criterion.

It is clear that the relationships between cognition and the brain, and cognition and the body more broadly, are far from settled. I wish to argue that there are, in fact, close relationships between brain and bodily processes, on the one hand, and mental phenomena, on the other. The relationships involved, however, are neither the strong independence of computationalism and connectionism, nor the eliminative anti-emergentism of some dynamicists. Instead, they are relationships of emergence of mental processes in particular kinds of brain and bodily processes — and in biological processes more broadly.

The Emergence of Function and Representation

The biological foundations of cognition are those biological properties and processes that are essential for the emergence of cognition. I will explore the implications for such foundations of an approach to Cognitive Science called *interactivism*. Interactivism is, in the first instance, a model of the nature and emergence of representation, and, therefore, of cognition. The emergence of representation, in turn, depends on the emergence of normative function. So, the model to be outlined is that of the emergence of normative function, followed by the emergence of representation as a particular kind of function; the biological foundations are those that are crucial at each step of emergence. Interactivism is in important respects an instance of dynamic and autonomous agent models, but differs in basic ways from standard positions in both camps. Correspondingly, it offers a novel range of implications for the biological foundations of cognition.

My focus in this paper is on those biological foundations, and the interactive model itself has been presented elsewhere (Bickhard, 1993, 1996, 1997, 1998; Bickhard & Terveen, 1995; Bickhard & Richie, 1983; Campbell & Bickhard, 1986), so I will present an inspissated version of the model and the related arguments — hopefully, just enough to support the exploration of the implications for biological foundations.

Consider a far-from-equilibrium system, such as a candle flame or Benard cells in a pan of water. Benard convection cells form when there is sufficient heat differential between the top and the bottom of a layer of water. The convection cells will persist so long as the heat differential is maintained. Maintenance of the heat differential, in turn, is dependent on whatever external processes are creating it, such as a fire or electrical heat source underneath the pan holding the water. The crucial point for my current purposes is that the maintenance of the far-from-equilibrium system is dependent on processes external to the system itself.

The candle flame is an example of a different kind of system property. The flame makes contributions to its own maintenance through time. It maintains above combustion threshold temperatures; in a standard gravitational field and atmosphere, it induces convection, which both brings in oxygen and eliminates combustion products; and it vaporizes wax into fuel. A candle flame, like the Benard cells, is a far-from-equilibrium system, and is dependent for its continued existence on the maintenance of particular properties and processes that support its far-from-equilibrium conditions. Unlike the Benard cells, however, the flame is *self-maintenant* — it contributes to its own far-from-equilibrium maintaining conditions (Bickhard, 1993).

The candle flame, however, can do nothing if it is running out of candle. The flame cannot adopt relevantly differing self-maintaining interactions with its environment when that environment itself changes. Some far-from-equilibrium systems, however can. A science fiction candle flame could do that if it could, perhaps, undertake to seek new fuel when it detected that the candle was almost gone. A paramecium, for a non-science fiction example, is capable of swimming, and continuing to swim, so long as it is swimming up a sugar gradient, but will tumble for a moment if it “finds itself” swimming down the sugar gradient (Campbell, D. T., 1974, 1990). Such systems can alter their methods of self-maintenance in ways appropriate to their current environments. They tend to maintain their own property of being self-maintenant: they are *recursively self-maintenant*.

Self-maintenant and recursively self-maintenant systems are, I argue, the key emergent forms of far-from-equilibrium systems in which normative function and representation emerge, respectively (Bickhard, 1993). I will outline that model and some of the arguments in its favor, and then turn to examining the implications for biological foundations.

Normative Function. Simply put, the candle flame’s heat serves a function for the flame insofar as it contributes to the flame’s maintenance. Function, in this view, is contribution, or tendency to contribute, to the maintenance of a far-from-equilibrium system, and is thereby always relative to some such system (Bickhard, 1993; Christensen & Bickhard, 2002). The heart of a parasite, for example, would likely serve a function for the parasite, but would be dysfunctional for the parasitized host.

This model of function is in strong contrast to the dominant approach in terms of evolutionary selection history (Millikan, 1984; Godfrey-Smith, 1994). The idea in these approaches is that a heart has the function of pumping blood, and not, say, that of making heart beat sounds, because it is in virtue of evolutionary predecessors having pumped

blood that this heart exists at all. The having of a function is constituted in having the right kind of etiology, the right kind of evolutionary history.

I will not focus on details of these models, but wish to simply point out that such a model entails that nothing can have a function if it does not have the right history, whether or not it contributes to self-maintenance. Millikan accepts the consequence that if a lion were to magically pop into existence that was molecule for molecule identical to a lion in the zoo, the heart of the science fiction lion would not have a function because it would not have the proper history (Millikan, 1984). Accepting such examples that will never occur may be a price well worth it if the etiological approach satisfied all other desiderata, but, unfortunately, what the science fiction example demonstrates is that the etiological approach cannot model function in terms of current state of a system. But only current state can be causally efficacious. The etiological approach, therefore, yields a model of function that is inherently epiphenomenal — nothing causal in the world depends on the presence or absence of having a function, that is, on the presence or absence of having the proper history, *per se*. What does make causal difference is current state — Does the heart pump blood, or not? — but current state is not sufficient for function in this view.

Clearly, the model of function as contribution to self-maintenance is definable in terms of current state — that is how it has *been* defined — and is therefore not epiphenomenal.

Representation. Function emerges in self-maintenant systems, and representation in recursively self-maintenant systems. The basic idea is that the *selection* of some process of interaction between the system and its environment has the *function* of contributing to the self-maintenance of the system, but *it could be in error*. That selection involves a kind of anticipation about the environment — it anticipates that the current environment is of the kind for which the selected processes are appropriate for helping to maintain relevant far-from-equilibrium conditions. The paramecium “anticipates” that swimming is the right thing to do, and it could be wrong — it could be swimming up a saccharin gradient, for example, not a sugar gradient.

This, I claim, is the most primitive level of emergence of representational truth value. Recursively self-maintenant selections involve anticipations that can be false; they involve implicit predications — e.g., predicating of this environment that it is of a type appropriate for swimming — that can be false. Representational truth value, in this manner, emerges quite naturally in the solution to the evolutionary problem of action selection.

These are quite primitive forms of representation, and do not look much like paradigm cases, such as those of objects or numbers, and that poses a challenge to this model: how can the interactive model account for representations beyond the simple cases of predicating types of interactive appropriateness to current environments? That question, and others, are addressed elsewhere (Bickhard, 1993, 1999, 2000; Bickhard & Terveen, 1995; Campbell & Bickhard, 1986), so I will only indicate a little of the direction of response here.

The selections of a recursively self-maintenant system necessarily involve some sort of appropriate sensitivity to the environment. The paramecium must detect that the current environment is of a type that is usually appropriate for swimming. The selections, then, are conditional on prior detections, prior outcomes of earlier engagements with that environment. Such engagements may be relatively simple and passive, but they may also be potentially complex and fully interactive themselves: the internal outcome of an *interaction* could serve to detect a relevant environmental type as well or better than a passive input processing detection.

The relevant implication here is that interactions can not only be selected and be indicated as potential selections, they can also function to differentiate environmental types for subsequent indications of appropriateness of interactions. Selections, then, can be based on indications of potentiality, and indications of potentiality can be *conditional*: if *this* interactive outcome has occurred, then *that* interaction is now indicated as appropriate.

Conditional indications of interactive potentiality are the key to more complex forms of representation. Such conditional relationships can branch — a single interaction outcome can function to indicate multiple further interactive potentialities — and they can iterate — completion of interaction A may indicate the potentiality of B, which, if completed, would indicate the potentiality of C, and so on. That is, they can form potentially quite complex webs of conditional interactive potentialities, and particular kinds of such webs, I argue, yield the emergence of higher kinds of representation, such as of manipulable objects (Bickhard, 1980, 1993; Piaget, 1954). Representation of abstractions, such as of numbers, requires still further development (Campbell & Bickhard, 1986).

Correspondence Models of Representation. For current purposes, I will take the interactive model to be a viable contender as a model of the nature of representation. Before turning to the examination of the biological foundations, note that this model of representation, like the model of function, is in strong contrast to dominant approaches. In particular, standard information semantic approaches to representation will construe the differentiations or detections involved in the interactive model as themselves

being representations of the properties that they differentiate or detect. In accomplishing such differentiations, which are necessary for the subsequent selections of a recursively self-maintaining system to be appropriately sensitive to the environment, the system does thereby create informational relationships, informational and perhaps causal correspondences, with whatever properties it has detected. But, whereas the interactive model needs only that those detections be *in fact* of properties that are relevant to further indications of interactive potentiality, and does *not* require that those properties be *represented*, it is precisely such detections or differentiations that are standardly construed as being representations of whatever properties they have differentiated — especially when those differentiations have occurred via the passive processing of inputs such as allegedly occurs in perceptual systems (Bickhard & Richie, 1983). Correspondence models construe representation as looking backward in time, down the input stream, to some privileged locus of that input stream, such as the surface of a table reflecting light into the eye. The interactive model construes representation as future oriented, to the potentialities for further interaction that are afforded by the current environment, and to the conditional webs of potential interactions that can extend indefinitely in space and time.

Correspondence models of representation suffer from a host of fatal problems (Bickhard, 1993, 1996; Bickhard & Terveen, 1995). I will outline only a few of them here. The first is that such models cannot account for their own foundations. It is clear that not all informational or causal correspondences are representations. In fact, hardly any are: the universe is full of such correspondences, since, for example, every single instance of every causal law yields such a correspondence. So one of the tasks facing a correspondence model is to specify what special *kind* of correspondence is representational — is to model what makes alleged *representational* correspondences so different from all the others. There are multiple candidates for these extra criteria to be imposed on the class of correspondences per se (Fodor, 1990, 1998; Stich & Warfield, 1994), but the point that I wish to focus on here is that none of them provide a naturalized model of the representational *content*, the mental content, for these alleged representations.

Representational content is that which specifies for the system what a representation is *supposed* to represent. It is the normative aspect of representation. It is what makes it possible for a representation to be in error: if a content of “cow” is predicated of what is in fact a horse, then that predication is false. Correspondence models are in a serious bind: if the special correspondence exists, then the representation exists, and that representation is correct, but if the special correspondence does not exist, then the representation does not exist, and it cannot be false. Information semantics, and other correspondence models,

have grave difficulties accounting for the possibility of representational error (Fodor, 1990; Levine & Bickhard, 1999; Loewer & Rey, 1991). And, therefore, accounting for representational content per se.

Such models can address complex representations as being constructed as combinations of more basic representations, and the more basic representations might be themselves constructed out of still more basic representations, but this composition and decomposition process must have a base, a ground or foundation, of representations that are not constructed out of nor defined in terms of any other representations. It is this foundation that cannot be accounted for. The models only permit the creation of new representations in terms of already available representations. There is no account of newly emergent representation (even if a new *correspondence* — a new differentiation — were created, perhaps in a connectionist net, there is no account of how the system can get any content about what those correspondences are *supposed* to represent).

One move at this point is to posit that all basic representations must be innate, and all representations that, say, adult human beings are capable of are created as complexes out of those innate atoms (Fodor, 1981). But the problem here is a logical one: there is no model of how correspondence representational content can come into existence, and evolution is just as helpless in the face of that problem as is individual level learning and development. Correspondence models, in other words, presuppose a foundational level of representation that they are incapable of modeling, that they themselves render impossible. They involve an internal contradiction in their assumptions, an incoherence (Bickhard, 1993).

Correspondence models render the emergence of new representation impossible. Neither evolution nor development nor learning can generate emergent representation, if this approach is correct. But there were no representations at the moment of the Big Bang and there are representations now. Therefore, representation has emerged. Therefore, representational emergence is possible. Therefore, any model that renders representational emergence impossible is false. Therefore, correspondence models are false.

Note that the interactive model renders representational emergence almost trivial: any new construction of system process organization that happens to involve the right kinds of indications of interactive potentiality will constitute an emergent representation.

Biological Foundations

Far-from-equilibrium systems. The clearest implication of the interactive model for biological foundations is that genuine representation, thus genuine cognition, can emerge only in far-from-equilibrium systems — more precisely, recursively self-maintenant systems. It is only in far-from-equilibrium systems that function can emerge, and only a certain kind of interactive indicational function that constitutes emergent representation.

This is already a major shift from computer model or connectionist approaches. If the interactive model is correct, various properties of cognition can be simulated in computers and in connectionist nets, but they cannot be more than simulations. There are actually several deficiencies of computer models. The first is that they are not self-maintenant systems — by most standards, they are not even far-from-equilibrium systems (their dependence on externally supplied power does keep them at some distance from thermodynamic equilibrium).

Open systems. A second is that computers are not in any relevant sense open systems. They process inputs or data, but do not interact with their environments. Recursive self-maintenance requires environmental interactions that achieve a closure in the sense of circling back to support the far-from-equilibrium conditions that made those interactions possible in the first place (Christensen & Hooker, 2000; Ruiz-Mirazo & Moreno, 1998).

Interactive timing. A third involves issues of interactive timing. Successful interactions require appropriate timing. Mere speed is not sufficient; an interaction can fail from being too fast just as easily as from being too slow. Computers do have timing information in their internal clocks, but there is no timing in the Turing machine formalisms (and equivalents) for computers, and the architecture of the timing in a computer is not an evolutionary possibility. If the brain involved clock driven processes akin to those in a computer, every evolutionary change in the brain would have to have involved simultaneous adaptive changes in the interactive circuitry *and* precisely coordinated changes in the timing circuitry. That is vanishingly improbable for any instance, and simply impossible for any extended evolution (Bickhard & Terveen, 1995).

The brain solves the timing problem in a very different way: Put clocks everywhere, and construct all functional relationships out of relationships among those clocks. Since clocks are “just” oscillators, this translates into: build the central nervous system out of oscillators, and construct all functional relationships out of modulatory relationships among those oscillatory subsystems. This is, in fact, the basic architecture of

neural functioning. Neurons and neural circuits are oscillatory, involving baseline levels of oscillation which are modulated by influences from other neurons and neural circuits. Some kinds of neurons never fire at all, but do modulate the activities of others. And evolution has created a virtual tool kit of temporal and spatial ranges of modulatory influences, from tiny and very fast gap junctions, to classical synapses, to volume transmitters that diffuse throughout a local population of neurons, to graded release of transmitters that are not all or nothing, and so on (Bickhard & Terveen, 1995).

Such oscillatory and modulatory architectural principles are at least as powerful as classical conceptions: a limit case of one system modulating another is for the first system to switch the second on and off, and switches are sufficient for the construction of computers. They are more powerful in that they inherently provide timing, while Turing machines do not, and computers involve timing in a biologically impossible form.

The idealization of neurons into threshold switches that so often occurs in computer perspectives, or of simple activation level transformers as in connectionist models, is seriously unfaithful to what actually occurs in the brain (Bickhard & Terveen, 1995). The contrasting oscillatory and modulatory architectural consequence of the interactive model is not logically forced, but it is forced by evolutionary considerations, and it is consonant with actual brain processes.

Irreversibility and normative scaling. One further biological foundation that I would like to consider arises from looking more carefully at the case of a computer controlling a robot. A robot can interact with the world, and so would seem to satisfy the interactive condition. To be successful in its interactions, the computer would have to appropriately handle timing in some way, even if not in a biologically plausible way. If we suppose, for example, that among the capabilities of the robot is the ability to detect when its batteries are running low, and to seek out power sources to replenish its batteries when that occurs, then we would seem to have, in some minimal sense, a far-from-equilibrium system that is also minimally self-maintaining, and even recursively self-maintaining (because it can switch into and out of “power source seeking” similarly to the paramecia’s ability to switch into and out of “swimming”).

Could such a robot have emergent cognition? The contrast with the biological case arises in the fact that most of the robot’s body is *not* far-from-equilibrium, cannot be self-maintained, and certainly not recursively self-maintained. Conversely, the only part of the robot that *is* far from equilibrium, the battery, is *not* self-maintaining.

Biological organisms involve elaborate infrastructure — organs, bones, and so on — that are themselves (mostly) far-from-equilibrium and self-maintained in constant overturning of molecular constituents. What renders some parts of the overall far-from-equilibrium system as *infrastructure* is that infrastructure involves slower time scales of interaction and often higher energy levels than other interactions of the systems. Such infrastructure is required in order to control and to enable the multiple processes — metabolic processes — necessary to self-maintenance and to the creation, selection, and execution of recursively self-maintenant interactions (Ruiz-Mirazo & Moreno, 1998).

The robot is far-from-equilibrium only in the sense that its operations require power; the basic existence of the robot is not far-from-equilibrium, and does not require self-maintenance or recursive self-maintenance. In that sense, the robot does not have much at stake in its self-maintenant activities, at least relative to the biological case. At issue is the metaphysical significance of those differences for the emergence of function and representation.

Any types of emergent phenomena are likely to have marginal cases, and the robot example is clearly marginal. But the nature of its marginality is important, I contend, because it illustrates how profoundly far-from-equilibrium conditions, thus self-maintenant and recursively self-maintenant processes, permeate the emergence of function and representation in paradigm biological cases. Robots can be marginally far-from-equilibrium by design, and, therefore, perhaps marginally self-maintenant and recursively self-maintenant. They do not, however, form an evolutionarily marginal case. Neither function nor representation could have evolved via such a kind of marginality.

There is a fundamental difference between the far-from-equilibrium character of living beings and that of robots that organizes this marginality that robots illustrate into a kind of continuum. Living beings are *irreversible* far-from-equilibrium systems, while robots, in general, exhibit a reversibility. A robot can be turned off, then on again; if its battery runs down, nothing in particular is necessarily lost if that battery is later recharged. The thermodynamics of living systems are not reversible. If the far-from-equilibrium processes that constitute them are halted, they cannot be restarted. This is the deeper perspective on the significance of far-from-equilibrium infrastructure and metabolism in living systems: it is not just the *activities* of living systems that require far-from-equilibrium conditions, but their very ontological constitution is far-from-equilibrium, and it begins to irreversibly move toward equilibrium once that far-from-equilibrium process is sufficiently interrupted (including by the accumulation of errors during normal functioning — aging).

The significance of far-from-equilibrium infrastructure and metabolism is that the normative aspect of representation emerges in (functional) contributions to the maintenance of far-from-equilibrium conditions. If the far-from-equilibrium conditions are “minimal”, or of minimal importance, then so also are the normative aspects of function and representation. It is this normativity that is indirectly scaled by any scaling of the far-from-equilibrium status of the ‘body’ of the system involved. Function and representation are normative insofar as, and to the extent that, there is something that is ontologically at stake in their ‘successful’ functioning (Christensen & Bickhard, 2002).

Self-reproducing and living systems? This brings up the last point I wish to make about biological foundations. Cognition, if this model is correct, requires far-from-equilibrium systems that are recursively self-maintenant, in perhaps very complex ways. There is nothing in the model, however, that requires that these systems be *self-reproducing*. It is a matter of definition, then, whether it is required that they be *living*. (Not all animals are capable of self-reproduction, not even all animals that are capable of clear cognition: e.g., any sterile organism, be it insect, such as a bee, or mammal, or human.) The interactive model, then, suggests that designed systems could in principle be fully capable of function, representation, and cognition — but not with the architectural and process resources available in computer or connectionist models.

Conclusion

The interactive model is a version of the dynamic systems approach, and of autonomous agent approaches (Bickhard & Terveen, 1995). Unlike other versions, however, interactivism neither argues *against* the usefulness of correspondence representations (Port & van Gelder, 1995), nor for the *necessity* of correspondence, informational, representations (Clark & Toribio, 1995). Instead, it argues against the common underlying assumption that representation is correspondence — and it argues that function and representation emerge naturally in the evolution of dynamic autonomous biological agents, but that they emerge in a way that correspondence approaches cannot account for. Interactivism reconfigures our assumptions about the biological foundations for cognition, and, therefore, for Cognitive Science.

References

- Beer, R. D. (1995). 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 Publishers.

- 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. O'Nuallain, P. McKeivitt, E. MacAogain (Eds.). *Two Sciences of Mind*. (115-131). Amsterdam: John Benjamins.
- Bickhard, M. H. (1998). A Process Model of the Emergence of Representation. In G. L. Farre, T. Oksala (Eds.) *Emergence, Complexity, Hierarchy, Organization, Selected and Edited Papers from the ECHO III Conference. Acta Polytechnica Scandinavica, Mathematics, Computing and Management in Engineering Series No. 91, Espoo, Finland, August 3 - 7, 1998*, 263-270.
- Bickhard, M. H. (1999). Interaction and Representation. *Theory & Psychology*, 9(4), 435-458.
- Bickhard, M. H. (2000). Motivation and Emotion: An Interactive Process Model. In R. D. Ellis, N. Newton (Eds.) *The Caldron of Consciousness*. (161-178). J. Benjamins.
- Bickhard, M. H., Richie, D. M. (1983). *On the Nature of Representation: A Case Study of James Gibson's Theory of Perception*. New York: Praeger Publishers.
- Bickhard, M. H., Terveen, L. (1995). *Foundational Issues in Artificial Intelligence and Cognitive Science: Impasse and Solution*. Elsevier Scientific.
- Brooks, R. A. (1991). Intelligence without Representation. *Artificial Intelligence*, 47(1-3), 139-159.
- Campbell, D. T. (1974). 'Downward Causation' in Hierarchically Organized Biological Systems. In F. J. Ayala, T. Dobzhansky (Eds.) *Studies in the Philosophy of Biology*. (179-186). Berkeley, CA: University of California Press.
- Campbell, D. T. (1990). Levels of Organization, Downward Causation, and the Selection-Theory Approach to Evolutionary Epistemology. In Greenberg, G., & Tobach, E. (Eds.) *Theories of the Evolution of Knowing*. (1-17). Hillsdale, NJ: Erlbaum.
- Campbell, R. L., Bickhard, M. H. (1986). *Knowing Levels and Developmental Stages. Contributions to Human Development*. Basel, Switzerland: Karger.

- Christensen, W. D., Bickhard, M. H. (2002). The Process Dynamics of Normative Function. *Monist*, 85(1), 3-28.
- Christensen, W. D., Hooker, C. A. (2000). An Interactivist-Constructivist Approach to Intelligence: Self-Directed Anticipative Learning. *Philosophical Psychology*, 13(1), 5-45.
- Clark, A., Toribio, J. (1995). Doing without Representing? *Synthese*, 101, 401-431.
- Fodor, J. A. (1981). The present status of the innateness controversy. In J. Fodor *RePresentations* (257-316). Cambridge: MIT Press.
- Fodor, J. A. (1990). *A Theory of Content*. Cambridge, MA: MIT Press.
- Fodor, J. A. (1998). *Concepts: Where Cognitive Science went wrong*. Oxford.
- Godfrey-Smith, P. (1994). A Modern History Theory of Functions. *Nous*, 28(3), 344-362.
- Levine, A., Bickhard, M. H. (1999). Concepts: Where Fodor Went Wrong. *Philosophical Psychology*, 12(1), 5-23.
- 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.
- 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.
- Ruiz-Mirazo, K., Moreno, A. (1998). Autonomy and Emergence: How systems become agents through the generation of functional constraints. In G. L. Farre, T. Oksala (Eds.) *Emergence, Complexity, Hierarchy, Organization*, Selected and Edited Papers from the *ECHO III Conference*. *Acta Polytechnica Scandinavica*, Mathematics, Computing and Management in Engineering Series No. 91, Espoo, Finland, August 3 - 7, 1998, 273-282.
- Stich, S., Warfield, T. A. (1994). *Mental representation : a reader*. Oxford, UK: Blackwell.