Mechanism is not enough

Mark H. Bickhard
Lehigh University

I will argue that mechanism is not sufficient to capture representation, thus cognition. More generally, mechanism is not sufficient to capture normativity of any sort. I will also outline a model of emergent normativity, representational normativity in particular, and show how it transcends these limitations of mechanism. To begin, I will address some illustrative attempts to model representation within mechanistically naturalistic frameworks, first rather generally, and then in the cases of the models of Fodor and Millikan.

Keywords: encodingism, Fodor, interactivism, mechanism, Millikan, pragmatism, representation, self maintenance, thermodynamics

1. Encoding models of representation

Attempts to model representation generally assume that representation is some form of encoding, in which a representation encodes that which it represents, and the focus of the models is on capturing the correct relationship that constitutes an encoding relationship. Deep issues are entwined in this endeavor: the constitutive relationship must, if the model is to be a naturalistic model, be within the realm of facts and causes, while the representational relationship that is to be modeled bears a form of normativity — the normativity of truth and falsity. Success in any such model would constitute a refutation of Hume's argument that norms cannot be derived from facts.[1]

Although I will offer a model of emergent representation, my first critical point is that no such model is possible within this general encoding framework. An alternative framework is required.

1.1 Real encodings

Encodings certainly do exist: Morse code is a good example. But Morse code also illustrates why encoding approaches to representation cannot capture foundational representation. "…" encodes "s" in Morse code insofar as relevant people already

Pragmatics & Cognition 15:3 (2007), 573–585.
ISSN 0929–0907 / E-ISSN 1569–9943

represent strings of dots and dashes, and the letters of the alphabet, and the Morse coding relationships among them. Morse code is useful because dots and dashes can be sent over telegraph wires, while alphabet characters cannot. The crucial point here is that codes already presuppose representation, and, thus, cannot provide an account of foundational representation. Encodings are stand-ins for already available representations, such as "…" standing-in for "s".

Lest it seem that these points turn on Morse code being a conventional code, consider a 'natural' code, such as the neutrino count in some deep gold mine encoding certain properties of fusion processes in the sun. This relationship is also an encoding relationship; it is not conventional; and it is useful because the neutrino count is accessible by scientists on earth, while the processes in the interior of the sun are not. Nevertheless, again, the encoding relationship exists only insofar as relevant scientists already represent neutrino counts, fusion processes, and the relationships between them.

Attempts to model representation as encodings, therefore, are inherently circular. They presuppose representation in order for the encodings to exist at all. These circularities are manifest most strongly with respect to the issues of representational normativity. Two such issues that I will consider are accounting for the possibility of representational error — a task to which much attention has been paid in the last decades — and a strengthened criterion of accounting for the possibility of *system detectable* representational error — a task that has received virtually no attention. Lest system detectable representational error seem unimportant or at least postponable, I point out that: (1) if system detectable error is not possible, then error guided behavior and learning are not possible, and (2) the possibility of system detectable representational error is precisely what the radical skeptical argument purports to refute.

That skeptical argument points out that, in order to check the correctness or error of our representations, we would have to somehow stand outside of ourselves so that we had an independent epistemic access to what we were attempting to represent, and yet also have access to the representational contents that we were attributing to what was being represented, so that we could compare the two and determine if the content appropriately applied to what was being represented. Because we cannot step outside of ourselves in this manner, so the argument goes, we cannot determine whether or not our representations are correct. We cannot determine their truth value: system detectable representational error is impossible. Both the venerability of this argument and the long standing lack of success in transcending it, as well as the simple fact that error guided behavior and learning *do* occur, highlight the importance of the criterion of system detectable representational error as applied to purported models of representation. Any model that cannot account for it is thereby refuted.

The critical criteria of circularity and inability to account for system detectable error apply, for example, to both Fodor and Millikan. I begin with Fodor.

1.2 Fodor

The center of Fodor's model is the claim that representation is constituted in nomological, lawlike relationships (Fodor 1987, 1990, 1998, 2003). Activities in the retina represent, encode properties of the light in virtue of the fact that there are lawlike relationships between light properties and retinal activities.

Such a model encounters an immediate challenge regarding the possibility of representational error: either the lawlike relationship exists, in which case the representation exists and it is correct, or else the relationship doesn't exist, in which case the representation doesn't exist. But there is a third possibility that must be addressed: the representation exists, but it is incorrect. Unfortunately, there are only two possibilities here: the relationship exists or it doesn't.[2]

Fodor's attempt to avoid this problem turns on a purported asymmetric counterfactual dependency between cases of correct representation and cases of incorrect representation. The idea is illustrated by the claim that a "cow" representation would incorrectly be used to represent a horse only if the "cow" representation already represented cows. So, there is a dependency of the false instance on the correct instances. But the dependency is not reciprocated: "cow" representations would represent cows even if they never falsely represented horses. So, the dependency is asymmetric.

Asymmetric dependency captures an intuition that false representation is somehow parasitic on true representation, but it is nevertheless subject to at least two fatal problems. First, there are counterexamples. Consider a poison molecule mimicking a neurotransmitter, docking on a receptor molecule in a neuron, and stimulating activity in that neuron. Here we have lawlike relationships, and an asymmetric dependency: the poison molecule would not initiate the neural activity if the neurotransmitter did not, but the neurotransmitter is not similarly dependent on the poison molecule. But there is no representation here. At best, there are functional and dysfunctional relationships, though even that depends on an absent model of normative function.

Second, there is no possibility of accounting for system detectable representational error. To detect error in its own representations on this model would require that the system, or organism, be able to access its own content — "cow" perhaps — and what it was trying to represent — a horse, perhaps — and compare them. But to 'know' what it was trying to represent is the representation problem all over again, as well as presupposing precisely the outside perspective that the skeptical argument points out we cannot have. So, such a check is circular and impossible.

Furthermore, the organism cannot have access to its own content. Content is constituted in lawlike relationships that participate in complex counterfactual relationships among what would be the case given various presences and absences, or possibilities and impossibilities, of the true and false instances of representation. No system, no organism, has access to any such structure of counterfactuals.

On Fodor's account, then, system detectable error, thus error guided behavior and learning are impossible. The account is therefore refuted.

1.3 Millikan

I turn next to a very different account of representation, that of Millikan (1984, 1993). Millikan's account is embedded in an underlying account of normative function: X represents Y insofar as it is the direct or derived function of X to do so. The overall model is sophisticated, complex, and ingenious, but my critical points depend only on the basic framework of the model.

In this model, something — a kidney, perhaps — has a function — e.g., filtering blood — insofar as the evolutionary ancestral kidneys have been selected for accomplishing this causal task. That is, filtering blood is one of many of the causal consequences of kidneys, and that particular consequence is constituted as the function of the kidney insofar as that particular causal consequence (unlike, say, contributing to the mass of the overall organism) has been selected for. Having a function, then, is constituted in having the right kind of evolutionary history.

This model has several counter-intuitive consequences, many of which are discussed by Millikan and like-minded theoreticians and philosophers. If a lion were to improbably pop into existence in the room and, by assumption, it were molecule by molecule identical to a lion in the zoo, the lion in the zoo would have proper evolutionary histories, and, therefore, its organs would have functions. But the lion in the room would have no evolutionary history, and, therefore, its organs would not have functions. Counter-intuitive this may be, but it is generally considered to be a small price to pay for the overall success of the model in accounting for normative function: the sense in which function and dysfunction make sense — in which a kidney can be functionally successful or unsuccessful.

But there is a further consequence much less often discussed. The two lions are, by assumption, causally, dynamically, identical. But one has functions and one does not. So, the having of functions — or not — makes no difference in the causal or dynamic properties of the world. Having functions is a matter of history, but identical current states will be dynamically the same, no matter how different the histories leading up to them. This model of function renders function causally epiphenomenal; and, therefore, any model of representation based on it, is also causally epiphenomenal.

Interestingly, this model seems to have a ready way in which to account for the possibility of representational error. If something has the function, due to its evolutionary history, of representing, say, cows, but is currently being applied to a horse, then it is in error. Such an account of error is possible because the content of the representation is determined independently (i.e., in terms of evolutionary history) of what is currently being represented, whereas in Fodor's model, for example, the content *depends* on what is being represented.

Nevertheless, the same problems of circularity and the impossibility of *system detectable* error emerge again, in related though different forms. As before, for an organism to detect error in its own representation, comparison must be made between what is being represented and the representational content attributed to it. Again, having any independent epistemic access to what is being represented is the representational problem itself — circularity — and presumes the external autonomous perspective that the skeptical argument holds impossible.

Further, as for Fodor's model, there cannot be any access to the content that is being attributed to what is being represented, this time because that content is constituted in the organism's evolutionary past. Even human evolutionary biologists could at best have theoretical hypotheses about what the relevant selection history might have been, and nothing akin to that is available to any other species. System detectable error, on this model too, is impossible.

2. Interaction and representation

There are several further notable models of representation in the literature, such as those of Dretske (1988) or Cummins (1996), but they too fail the circularity and system detectable error criteria (Bickhard 2003b, in preparation), and I will turn at this point to an outline of a model of representation that satisfies such criteria. This model argues that representation is an evolutionary emergent response to the problem of the guidance and selection of interaction. In this sense, it participates in a general pragmatist orientation, and, therefore, has kinships to Piaget's model (though also differences). It is also relevant that, in modeling representation as emergent in interaction systems, action and interaction are made central to representation and cognition, unlike the encoding models that would have action and interaction *make use* of representation, while representation per se has no need, in such models, of action or interaction at all. Action based models of representation and cognition, thus, make much stronger contact with both evolutionary biology and with robotics — with both natural and artificial agents (Buisson 2004; Cherian and Troxell 1995; Stojanov 2001; Stojanov and Kulakov 2003).

2.1 Forms of stable process

Consider two fundamental forms of stability in the world: energy well stability and the stability of far from thermodynamic equilibrium systems. Energy well stability is constituted in patterns of process that remain stable unless energy that is above some crucial threshold impinges on them. If no such above threshold energy happens to strike them, they can remain stable for cosmological time periods. Atoms are a canonical example.

For my purposes, the most important aspect of energy well stabilities is that they remain stable if isolated and left to go to equilibrium. This is in strong contrast with the other form of stability, that of far from equilibrium systems. In contrast to energy well stabilities, far-from-equilibrium stabilities cannot be isolated: if they are, they go to equilibrium and cease to exist. Far from equilibrium stabilities must be maintained. They must be in interaction with an environment that supports and maintains the far from equilibrium conditions. There is a fundamental asymmetry here between energy well stabilities and far from equilibrium stabilities with respect to their interrelations with their environments: energy well stabilities are maximally stable in isolation, while far from equilibrium stabilities require maintenance and cannot continue to exist in isolation.

A canonical example of a simple far from equilibrium system would be a pan of water that is being heated and self organizes into Benard cells of boiling water. Self organization in far from equilibrium systems is of central importance, but my current focus is more on two special sub-categories of far from equilibrium systems: self-maintenant systems, and recursively self-maintenant systems.[3]

2.2 Self-maintenant systems and function

Self-maintenant systems are far from equilibrium systems that make contributions to their own far from equilibriums conditions — hence, to their own stability. An example is that of a candle flame. The candle flame maintains above combustion threshold temperature, it melts wax so that it will percolate up the wick, it vaporizes wax in the wick so that it is available for combustion, and it induces convection which brings in fresh oxygen and gets rid of waste products. It is self-maintenant in several senses.

The success of a self-maintenant system in being able to maintain essential conditions is vital to the continued existence of the system. If the candle flame runs out of wax, it cannot continue its self maintaining processes, and it ceases to exist. The fact of continued existence of such a system is dependent on the successful maintenance of the essential conditions. That is, maintaining those conditions *serves the function* of supporting that existence.

This is a normative sense of function in that self maintenance can be successful or unsuccessful, better or worse, independent of any outside criteria — the criterion of continued existence of the system is internal to the system, and intrinsic to its nature as far from equilibrium. Such a notion of serving a function is also inherent in the current dynamic conditions of the system, independent of its history, and thus, does not suffer from the causal epiphenomenality of the evolutionary history approach to function.

2.3 Recursively self-maintenant systems and representation

There is much more to be developed concerning this model of function,[4] but here I will turn to the still more special subcategory of recursively self-maintenant systems. The candle flame has no choices available if it runs out of wax, but some systems do. My canonical example here is the bacterium that can swim, and keep swimming if it finds itself swimming up a sugar gradient, but will switch to tumbling if it finds itself going down a sugar gradient (Campbell 1974, 1990). Swimming is self maintaining, is functional, insofar as it takes the bacterium into higher sugar concentrations, but that same activity is dysfunctional if it takes the bacterium into lower sugar concentrations. If it is headed in the wrong direction, tumbling switches in, and it tries swimming again.

Such a system is *recursively* self-maintenant in the sense that it maintains the condition of being self-maintenant in the face of variations in the relevant environmental conditions. It can detect and appropriately select what to do in various conditions so that what it does succeeds in being self-maintenant across those various conditions.

Central for my purposes, however, is that such success is not guaranteed. The selection of a particular interaction presupposes that the conditions that make that interaction appropriate obtain in the current situation. Interaction selection implicitly predicates those appropriateness conditions to the environment, and such a predication may be *true or false*. Being oriented up the sugar gradient supports the functionality of swimming, while being oriented up a saccharin gradient falsifies that presupposition. Here we have representational normativity — truth and falsity — emergent in the selection of interaction: the stability, or continued existence, of the system constitutes an ultimate normative criterion with respect to which interactions succeed or fail, and the presuppositions under which interactions would succeed can be true or false in particular conditions in which those interactions might be engaged in (or perhaps simply indicated as possible, as for more complex organisms, such a frog, which can in general have multiple possible interactions available at a particular time and must select among them).

Representation as emergent in interaction selection and indication does not look much like canonical representation, such as of a child's toy block. How such a primitive origin of representation could account for something like object representation is a crucial question. The answer is roughly Piagetian: In more complex organisms, potential interactions are indicated, not triggered as for the bacterium. Such indications are conditional: under some circumstances, such and such a range of interactions would be possible. And some such interactions may serve to create the conditions for other interactions: If X is successfully engaged in, then Y, Z, and a few others may then be possible. Indications of potential interactions, then, can branch and iterate, and, therefore, can form potentially quite complex webs of interaction potentialities. Object representations emerge as various kinds of properties and conditions within subwebs of such an overall interactive web (Bickhard 1993, 2003b).

2.4 Challenges

There are many challenges for such a model of representation, of how it can account for multifarious representational phenomena such as objects, abstractions such as numbers, concepts, and so on (Bickhard 1993, 2000, 2003a, 2003b, 2003c, 2005, in preparation). And there are a number of properties of this model of representation that are not common, such as that representing is future oriented, toward interaction potentialities, not past oriented, toward a locus of past causal or informational or nomological contact. Interactive representation is also inherently modal: it is with respect to possibilities, potentialities, not (past) actualities. And the content of representation in this model is intrinsically implicit, not explicit as in an encoding model. Such properties and more enable the interactive model to account for multiple phenomena that are at best difficult on standard models, and avoid various fatal problems that standard encoding models cannot avoid.

I will mention here just the two that were discussed above for the Fodor and Millikan models: circularity and system detectable error. In order to detect error in the interactive model, the organism need only to engage in the indicated interaction and detect if the flow of interaction remains within the indicated bounds. If it does not, then the indication, the representation, is false and is falsified by and for the system itself. Furthermore, no external epistemic perspective is required because the detection is internal and functional, not external and (therefore) epistemic. No separate representation of what is being represented is required in order to compare to the content of the representation because the indications are internal and are internally functionally accessible; therefore, there is no circularity. And because representational error can in this way be detected functionally by the system itself, error guided behavior and learning are possible. The radical skeptical argument is a valid argument — but it is unsound: it presupposes the encoding framework for representation, and is transcended by a pragmatist, future oriented, anticipatory, model of representation.

3. Why mechanism is not enough

Much more can and ultimately needs to be said concerning further challenges to and properties of the interactive model of representation, but I will here take it as established that the model is at least a viable candidate. If the critical points are sound, it is the only viable candidate because all alternatives in the literature fail those criteria. The question that began this discussion can then be addressed: is mechanism sufficient for representation and cognition? The negative answer is clear: it is not. But more can be said about why that is the case.

Representation is fundamentally normative. It involves the normativities of truth and falsity. Hume's argument that norms cannot be derived from facts, if it were sound, would preclude *any* account of representation within the framework of facts and causes. The argument is in fact unsound, but, nevertheless, it is very powerful and applies to encoding models of representation. It does not apply to pragmatist, anticipatory models of the emergence of representation.

The key to a naturalistic model of representation, then, is to provide a naturalistic account of normativity. Normativity involves distinctions between normatively positive and normatively negative — this is not just a differentiation, but an asymmetric differentiation in which some differentiated possibilities or categories are better than others with respect to some criterion. It is easy to get differentiations in physics, and it is easy to get preferences for some things over others in terms of human or other animal preferences or some other form of evaluation. But such preferences and evaluations are already normative and, therefore, cannot account for the independent emergence of normative phenomena. But physics provides little in the way of the kind of asymmetric differentiations that normativity requires. The two realms of the factual and the normative, then, would appear to remain split apart.

There is an exception, however. Thermodynamics involves the kinds of nonlinearities that can underlie asymmetric differentiations. This is crucial to the asymmetry between energy well stabilities and far from equilibrium stabilities that the above model is based on. Mechanics, on the other hand, involves linearities that do not support an asymmetry upon which something like normativity might be emergent. Mechanism is not enough.

4. Robots and living systems

The examples of representational systems mentioned above are living — bacteria and frogs. What about robots and other autonomous agents? Artificial agents such as these, interestingly, force an additional distinction within the general thermodynamic framework.

A robot with a full battery charge is not at equilibrium. The battery, at least, is at some distance from equilibrium, headed toward equilibrium as it discharges. Does this suffice for the emergence of normativity, and, therefore, perhaps for representation and cognition? This might appear to be a boundary issue, one that has no clear answer because it falls on the boundary between crucial categories. In part, I suggest that this is the case, but I think that a deeper distinction is relevant.

The robot with the charged battery requires that non-equilibrium state of the battery in order for the robot to operate. At full equilibrium, however, with discharged battery, the robot can no longer operate, *but it still exists*. Once the battery is charged again, the robot can operate again. So, the status of being distant from equilibrium is in this case crucial not to the continued existence of the system, as for, say, a bacterium or frog, but rather to the ongoing operation of the system, and that operationality can be restored by restoring the charge in the battery. Most of the robot is in fact in forms of energy well stability, and the non-equilibrium state of the battery has no relevance to the continued existence of those energy well processes.

What is most crucial to the case of living systems is that they are far from equilibrium not only with respect to their "operation", but with respect to their existence. The thermodynamics involved in the ontology of the kinds of system that they exemplify is an irreversible thermodynamics — they cannot be "restored", because of that irreversibility. They die if they go to equilibrium, and that is final. Furthermore, it makes a difference in the rest of the world whether or not they continue in their far from equilibrium conditions: the actions of a recursively self maintenant system have potentially very different consequences from those of the system at or on its way to equilibrium. So, the continued existence of such a system constitutes an asymmetric wedge that pries apart contributions to the continued existence of the system from damages to the continued existence of the system — just the kind of asymmetry needed to constitute the emergence of normativity, serving a function in this case.

5. Conclusion

In sum, the conclusion that mechanism is not sufficient for representation and cognition follows from the fact that mechanism does not provide the kind of nonlinear asymmetries needed to account for normativity, representational normativity in particular. Thermodynamics, on the other hand, does, and it is on this basis that the interactive model above is developed. Finally, robotics illuminates a boundary realm of examples that highlight the importance of the far from equilibrium thermodynamics being an irreversible thermodynamics crucial to the existence of the system, not just the operation of the system.

Notes

1. Hume's argument is an unsound argument (Bickhard 2003b, in preparation). But part of the burden of the argument in the text is that, although it is possible to model emergent representation, thus normativity, it cannot be done within the framework of mechanism.

2. Note that the general form of this argument applies to any purported 'natural' encoding relationship, whether it be taken as informational, causal, lawlike, structural, or whatever. Various models in the literature differ with respect to how they attempt to avoid this problem. Many, such as the information "semantics" that dominates cognitive psychology and cognitive science, or the related symbol system hypothesis, don't even address the issue (Bickhard 1993, 2003b, in preparation; Bickhard and Terveen 1995).

3. Self-maintenance and recursive self-maintenance constitute two simple forms of autonomy (Bickhard 2003b, in preparation; Christensen and Bickhard 2002).

4. One important difference is that this model takes "serving a function" as its primary locus, while the evolutionary models take "having a function" as their primary locus. Each must ultimately account for both, but the difference in initial locus is important: among other consequences, it means that the self-maintenance model of function is naturally much broader than the evolutionary model (or etiological models more generally) because, in such models, a function can be served *only* if something actually *has* that function, while in the self maintenance model, functions can be served without anything *having* that function. See Bickhard (2003b, in preparation; Christensen and Bickhard 2002) for further discussion.

References

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

Bickhard, M.H. 2000. "Motivation and emotion: An interactive process model". In R.D. Ellis and N. Newton (eds), *The Caldron of Consciousness*, Amsterdam: John Benjamins, 161–178.

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

Bickhard, M.H. 2003b. “Process and emergence: Normative function and representation”. In J. Seibt (ed), *Process Theories: Crossdisciplinary Studies in Dynamic Categories*. Dordrecht: Kluwer Academic, 121–155.

Bickhard, M.H. 2003c. “The biological emergence of representation”. In T. Brown and L. Smith (eds), *Emergence and Reduction: Proceedings of the 29th Annual Symposium of the Jean Piaget Society*. Mahwah, NJ: Erlbaum, 105–131.

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

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

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

Buisson J.-C. 2004. “A rhythm recognition computer program to advocate interactivist perception”. *Cognitive Science* 28: 75–87.

Campbell, D.T. 1974. “Evolutionary epistemology”. In P.A. Schilpp (ed), *The Philosophy of Karl Popper*. LaSalle, IL: Open Court, 413–463.

Campbell, D.T. 1990. “Levels of organization, downward causation, and the selection-theory approach to evolutionary epistemology”. In G. Greenberg, and E. Tobach (eds), *Theories of the Evolution of Knowing*. Hillsdale, NJ: Erlbaum, 1–17.

Cherian, S., and Troxell, W.O. 1995. “Interactivism: A functional model of representation for behavior-based systems”. In F. Morán, A. Moreno, J. Merelo, and P. Chacón, *Advances in Artificial Life*: Proceedings of the Third European Conference on Artificial Life, Granada, Spain. Berlin: Springer, 691–703.

Christensen, W.D. and Bickhard, M.H. 2002. “The process dynamics of normative function”. *The Monist,* 85(1): 3–28.

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

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

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

Fodor, J.A. 1990. *A Theory of Content*. Cambridge, MA: The MIT Press.

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

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

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

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

Stojanov, G. 2001. “Petitagé: A case study in developmental robotics”. *Proceedings of the 1st International Workshop on Epigenetic Robotics, Lund University Cognitive Studies.* Lund, Sweden, 85: 39–46.

Stojanov, G., Kulakov, A. 2003. “Interactivist approach to representation in epigenetic agents”. *Proceedings of the 3rd International Workshop on Epigenetic Robotics, Lund University Cognitive Studies.* Lund, Sweden, 106: 123–130.

Author's address

Mark H. Bickhard
Lehigh University
17 Memorial Drive East
Bethlehem, PA 18015

mark@bickhard.name
http://bickhard.ws/

About the author

Mark H. Bickhard joined Lehigh University in 1990 as Henry R. Luce Professor in Cognitive Robotics and the Philosophy of Knowledge. He is affiliated with the Departments of Psychology and Philosophy, and is Director of the Institute for Interactivist Studies. He is Editor of *New Ideas in Psychology*. He was Director of Cognitive Science from 1992 thru 2003. His work focuses on the nature and development of persons, as biological, psychological, and social beings. This work has generated an integrated organization of models encompassing *The Whole Person*, which is the tentative title of a book in preparation.