Journal of Theoretical and Philosophical Psychology
2015, Vol. 35, No. 1, 53–66
© 2014 American Psychological Association
1068-8471/15/$12.00 http://dx.doi.org/10.1037/a0038059

What Could Cognition Be If Not Computation . . . or Connectionism, or Dynamic Systems?[1]

Mark H. Bickhard
Lehigh University

Despite the fact that representation is at the center of artificial intelligence and cognitive science, there is still no consensual model of what representation, thus cognition, is. There are good reasons for this impasse: none of the approaches to modeling representation that are currently on offer is ultimately viable. I outline several of the deepest reasons for this conclusion, and apply 1 of these desiderata in particular to multiple models of representation on offer in the contemporary literature, showing that *none* of these models can satisfy this criterion. A brief overview of an alternative, pragmatist-based, approach to modeling representation and cognition is then presented, and it is shown that this model *does* satisfy *all* of the desiderata outlined earlier. Finally, I address 2 even more recent approaches to cognition—enactivism and action-oriented predictive brain models—in a little more detail, and show that they too fail to satisfy the desiderata.

Keywords: connectionism, dynamic systems, enactivism, information processing, interactivism, predictive encoding, representation

This document is copyrighted by the American Psychological Association or one of its allied publishers.
This article is intended solely for the personal use of the individual user and is not to be disseminated broadly.

Some Desiderata for Models of Representation

What characteristics should an acceptable model of representation have? What characteristics should fail a model on offer if that model cannot satisfy them? I will outline three such desiderata, and make use of primarily the third in surveying some representational models in the current literature. As will be evident, this third criterion—accounting for the possibility of system or organism detectable representational error—rather trivially refutes every model of representation on offer in the contemporary literature. It is striking, and perhaps even somewhat scandalous, that these refutations are so simple, and yet so overlooked or ignored.

These three desiderata, in fact, are part of a large family of criteria, critiques, and refutations of models of representation, some of which are of ancient provenance (including the criterion of organism-detectable error) and some noted more recently (Bickhard, 1993, 2009a). Nevertheless, the most common stance seems to be that there *must* be some way around the flaw(s)—after all, what else could representation be?—or else that the critical problems *cannot* be solved, so there is no point in addressing them—for example, the common avoidance of skeptical arguments. This is at least troublesome: we seem to have no viable model of the central notion of much (or most) of cognitive science, psychology, and philosophy. One likely reason is that there has until historically recently been no real alternative on offer to the standard symbolic, encodingist, information semantic, correspondence notions of representation.

I offer such an alternative; one that is historically related to pragmatism, especially that of Peirce. Pragmatism is only a little more than a century old, compared with millennia for standard frameworks for understanding mental phenomena, and has not yet been developed nearly as extensively. I will outline one exploration of the resources of pragmatism to explain representing below.

This article was published Online First September 22, 2014.

Correspondence concerning this article should be addressed to Mark H. Bickhard, Departments of Philosophy and Psychology, Lehigh University, Bethlehem, PA 18015. E-mail: mhb0@lehigh.edu

[1] With apologies to Tim van Gelder.

This document is copyrighted by the American Psychological Association or one of its allied publishers.
This article is intended solely for the personal use of the individual user and is not to be disseminated broadly.

Emergence

A first criterion for a model of representation is that it must account for the *emergence* of representation. Representation has to have emerged: it didn't exist 14 billion years ago and it does now. Any model of representation that makes this impossible is thereby refuted; any model that fails to account for this is thereby at least incomplete.[2]

It is also worth noting that pushing the problem of emergence onto evolution, such as in contemporary nativisms, does not resolve the issue (Allen & Bickhard, 2013a, 2013b; Fodor, 1975, 1981). If basic representational atoms must be innate because the only way in which new representations can be constructed is as combinations of already available representations, then any such in-principle restrictions apply to evolution just as much as to learning and development, and emergent representation is not possible. Such models are thereby refuted.

On the other hand, if evolution *can* generate emergent representation, then there is no model or explanation on offer of how evolution can accomplish this, but learning and development cannot. That is, if emergent representation is possible, such as in evolution, then why not in individual learning and development?

In either case, representation emerged, and any viable model of representation must be capable of accounting for that emergence.

Error

A second desideratum is to be able to account for the possibility of representational error: Being *about* something involves the possibilities of being true or false about that something.

Attempting to account for the possibility of representational error became a small in-house industry when it was realized that information semantics models seemed to make error not possible: If the crucial informational (causal, nomological) relationship that constitutes representation exists, then the representation exists and is correct; if the crucial relationship does not exist then the representation does not exist; and there is no third possibility for accounting for the representation existing but being false. There have been several attempts to provide such an account in the last decades, but, arguably, none of them succeed (Bickhard, 2009a).

Furthermore, even if they did succeed on their own terms, they address the problem of error only from the perspective of an external observer on the organism and its environment—they address how such an *observer* might determine that something in the organism that is 'supposed' to represent X is in error because what is in fact in the environment is (a) Y: how an observer might determine that a representation that is supposed to represent a cow is in fact "representing" a horse on a dark night (e.g., Fodor, 1990a, 1990c). This restriction to and reliance on an external observer violates the third desideratum:

Organism-Detectable Error

Representation must be more than an ascription by an observer—representation must be representation *for* an organism. A limitation to ascription fails to account for the representations of the ascripter—the observer. Furthermore, representation for an organism, for example, prior to the evolution of any ascriptive 'observers', must be capable not only of being true or false, but of being true or false from the organism's perspective, and, still further, at least in principle, *capable of being detected as true or false by that organism*.[3]

This might be considered to be a more complex desideratum, one that can and should be postponed till simpler issues concerning representation are solved or resolved. But this will not do: If organism-detectable error is not possible, then error-guided action and learning are

[2] Emergence is itself a problematic metaphysical notion, especially for the possibility of the emergence of various kinds of normativity. For an (anti-reductive) argument that ontological emergence does occur, and, in fact, is ubiquitous, see Bickhard (2009a).

[3] 'At least in principle' because there may be species that make minimal use of such error detection possibilities. But there *are* species that make major use, and, therefore, any model that makes such organism-detectable error impossible in principle cannot be correct. On the other hand, even bacteria can correct swimming *down* a sugar gradient (away from the sugar) by ceasing to swim and tumbling for a moment (Campbell, 1974). Note that an ability to detect error does not guarantee that such detections are themselves correct: a bacterium will continue to swim up a saccharin gradient as well as a sugar gradient. This functional error is not detected. And there is no a priori reason why an error detection could not be itself in error: "detecting" an error where none exists. *All* such functional processes are fallible, subject to the possibility of dysfunction.

This document is copyrighted by the American Psychological Association or one of its allied publishers.
This article is intended solely for the personal use of the individual user and is not to be disseminated broadly.

not possible.[4] Error-guided action and learning occur, therefore any model that renders this impossible is thereby refuted. Any model that cannot address this is thereby at least incomplete.

Representation Still Resists Naturalism

I will argue that none of the major models of representation available today satisfy these desiderata, with a focus on the third: organism-detectable error.

Fodor

Fodor advocates a version of information semantics. Representation is constituted in nomological relations between activities in the nervous system and phenomena in the environment, for example, ‘transduction’ of light in the retina into encodings of properties of the light (Fodor & Pylyshyn, 1981; Fodor, 1990a, 1990b, 1990c, 2003). This encounters the problem of the possibility of error, which Fodor proposes can be handled via a relationship of asymmetric dependency from false activations of a representation to true or correct activations. The idea is that a COW representation, for example, might be falsely activated by a horse on a dark night in way that is *dependent on* the fact that cows will activate COW, but that cows can activate COW *without* there being any dependency on horses being able to activate it. There is an unreciprocated dependency from the possibility of false activations to the possibility of true activations.

This does capture something of the intuition of the possibility of error: it is dependent on—parasitic on—the possibility of success. But it is not specific to representation at all: a ‘poison’ molecule might activate a neural receptor in a way that is parasitic on the possibilities of the correct neurotransmitter activating that receptor, but there is no representation here (Levine & Bickhard, 1999).

Most fundamentally, however, the model fails the criterion of organism-detectable error (Bickhard, 2009a). First, the ‘correct’ content of a representation in this model is constituted in complex asymmetric dependencies between counterfactual possibilities of false activations and correct activations. No organism has any access to such counterfactual ‘facts’ about itself, and, thus, no organism has access to its own contents to be able to compare them with what is actually present in the environment to be able determine truth or falsity.

Worse, comparing any such content with what is being represented requires epistemic access to what is being represented independent of the organism’s own representations—stepping outside of oneself to check what is being represented in the world and comparing it to the content of one’s representation. Epistemically stepping outside of oneself is impossible, thus organism-detectable error, and, thus, error-guided action and learning are impossible on this account. The account, therefore, cannot be correct—it is refuted.

Information Processing

An information processing framework is a general form of information semantics that dominates, for example, research in psychology, especially cognitive psychology, as well as cognitive science (e.g., Bermúdez, 2010). The foundation of the notion of information processing is the technical definition of information in terms of correlation or covariation: *X* carries information about *Y* if *X* is correlated with *Y*. There is no conceptual or theoretical difficulty with this technical notion of information per se, but it is commonly assumed that such informational relationships constitute representational relationships. This constitutes an equivocation between information as covariation and information as representational semantics. There are ubiquitous correlations all over the universe—for example, every causal relationship constitutes an informational relationship—but none of them are representational per se. If someone *knows* about such a correlational relationship, and *knows* about *X*, then they can infer things about *Y*, and thereby render *X* as an *interpretable* representation of *Y*, but this requires that all of the relevant variables (and the crucial rela-

[4] There are models that construe some kinds of passive responses to pure input, with no error, as learning (e.g., in unsupervised connectionist systems). There are important considerations missing from such passive learning construals (Bickhard & Terveen, 1995), but, even if these are accepted, such ‘passive learning’ would have to encompass *all* of learning to avoid the necessity for organism-detectable error. That is, in any case, error-guided learning exists, so organism-detectable error has to be possible.

This document is copyrighted by the American Psychological Association or one of its allied publishers.
This article is intended solely for the personal use of the individual user and is not to be disseminated broadly.

tionships among them) be already represented: interpretation is already deeply embedded in the realm of the representational. It cannot constitute nor yield foundational representation.

Furthermore, such a correlational model of information semantics cannot account for the possibility of error: If the correlation exists, then the 'representation' exists and is correct, whereas if the correlation doesn't exist, then the 'representation' doesn't exist, and there is no third possibility. And the information processing framework cannot even *address* the problem of organism-detectable error; it is thereby refuted as a model of representation.[5]

Millikan

One central reason that information semantics models have difficulty in accounting for the possibility of error is that, in such models, representational relationships only have two parts: the representation itself—presumably some condition or process in the organism—and the represented. With only two parts, if the relationship exists, then the represented exists, and so the representation is true.[6]

Millikan's model has a property that is crucial for accounting for the possibility of error: it is a three part model (Millikan, 1984, 1993). In particular, it models the existence of *content*—a specification of what the representation is *supposed* to represent—in addition to, and distinct from, representation and represented. This aspect or component of representation makes it possible for the content to misapply to the represented, and, therefore, for the instance of the representation to be false for the represented.

Models of content as a third aspect of representation have predecessors in Frege's Sinn (Beaney, 1997) and the Brentano-Twardowski notion of content (Jacquette, 2004a, 2004b; Poli, 1996; Smith, 1994). But these models don't really model content; instead they point out the functional necessity for content, and leave its nature unclear. Content is the normative aspect of representation—that which generates the normative values of truth and falsity; that which the representation is *supposed* to represent—and the long-standing difficulties in accounting for any kind of normativity within a naturalistic framework[7] have worked against models making attempts to account for content. Witness, for example, the pretzel twists in Fodor's attempt within a nominally two part information semantics.

Millikan, in contrast, has a genuine three-part model, and, therefore—insofar as the model works in other ways—it has no in-principle problem accounting for the possibility of representational error. There are, however, other problems.

In this model, a representation of *X* is something that *has the function* of indicating *X*. Something having such a function, in turn, is constituted in ancestral organs or processes in the species having been selected during evolution for having such indications (or for producing things that have such indications). For example, a kidney has the function of filtering blood: ancestral kidneys were selected for this filtering process, and having been selected for this filtering consequence is what constitutes the having of such a function by descendant kidneys. So, function is constituted in having the "right" evolutionary selection history, and representation-as-indication is a kind of function.

Consequently, what is *supposed* to be indicated by a given representation—the content of that representation—is constituted in having an appropriate evolutionary history. Content, then, is distinct from what is being represented—this is a three-part model—and representational error seems easily possible.

But, no organism has access to its own evolutionary history, and, therefore, no organism has access to its own content. So, no organism can compare its own content to what it is attempting to represent to determine error or lack of error.

---

[5] This is not to argue that *correlational* information is not useful or necessary: it is both. See Bickhard (2000b, 2014a, 2014b) for discussions.

[6] Brentano made the point that relationships in the "natural" world can exist only between relata that exist, whereas *representational* relationships can be with things that do not exist or are false (e.g., the golden mountain or unicorns; Albertazzi Libardi, & Poli, 1996; Jacquette, 2004a). This problem has, in fact, bedeviled Western thought since the Pre-Socratics: how can a representational relationship, a representational correspondence, exist with something (a relata) that doesn't exist (Gill, 1989; Campbell, 1992)?

[7] For example, the difficulties summarized in the Humean slogans "No ought from is" or "No norms from facts" For an argument that Hume's "reasoning" (actually, he didn't present an argument for the position) is valid but unsound, see Bickhard (2009a).

Finally, to compare one's own content, even if there were access to it, requires independent epistemic access to what is being represented—stepping outside of oneself—and that is impossible. Organism-detectable error, therefore, is impossible, and consequently also error-guided action and learning. The model, thus, is refuted.[8,9]

Dretske

Dretske's model is also an etiological model—a model in which the crucial normativity is constituted in having the right origins, the right history—like Millikan's (Dretske, 1988), except that the required history is a learning history rather than[10] an evolutionary history. Nevertheless, the same problems arise: There is no access to one's own content, and, therefore no way to compare that content to the world being represented. And, in what should by now be familiar, even if the organism did have access to its own content, it could not compare with the world because it does not have independent epistemic access to the world. It cannot step outside of itself.

Organism-detectable error, therefore, is not possible, and the model is refuted.

Cummins

Cummins (1996) also has a three-part model, and, thus, can seemingly account for the possibility of error. He distinguishes between representation and target of the representation, thus error seems possible in a lack of match between representation and target.

But, in this model, representation is constituted in structural isomorphism between representation and represented. Isomorphism, however, is transitive and reflexive, and representation is neither. Isomorphism is also relative to unbounded ranges of ways to determine what structure *is* in real systems; in general, many structures can be defined in any given real system, and Cummins does not specify how the 'correct' structure is to be identified. Furthermore, Cummins' notion of target is relative to the normative phenomena of purpose or goal, and there is no account of such normativity (Bickhard, 2009a).

So, in this model too, there is no way to access content (structure) and, as usual, no way to compare with the represented—the target—even if content could be accessed. Organism-detectable error is not possible, and the model is refuted.

Symbol System Hypothesis

In the original symbol system hypothesis (Newell, 1980; Newell & Simon, 1972, 1987a, 1987b), representation is simply designation or pointing. But there is no model of how designation could possibly work outside of the 'pointers to memory locations' model inside of a computer (Bickhard & Terveen, 1995). Such pointers are functional, not epistemic, and do not address representation at all.

Later, we find a footnote that dismisses problems about representation with the claim that it is constituted in structural isomorphism (Vera & Simon, 1993, 1994). The symbol system hypothesis, thus, does not take the problem of representation seriously; doesn't attempt to address the problem of representational error; and has no way to address comparison of representation with represented. Organism- or system-detectable error is impossible, and the model is refuted.

Connectionism

Connectionist models (e.g., Rumelhart & McClelland, 1986; Bechtel & Abrahamsen, 1990; Churchland, 1995; Bermúdez, 2010) are built out of nodes that are connected in various network topologies (often in layers) by links. This node-link architecture is intended to approximate a neuron/axon-synapse architecture in the brain. The nodes have levels of activation, usually with some subset of input nodes (e.g., a

[8] An additional serious problem is that, in being constituted in having the right history, this model of function makes function causally epiphenomenal: two organisms with identical dynamic/causal properties can differ with respect to their evolutionary histories, thus to having functions, thus with respect to having representation (Bickhard, 1993, 2009a).

[9] Note that Millikan's (2005) notion of pushmi-pullyu representation expands representation to include inducing changes in the world so that the world fits the representation (e.g., imperatives), in addition to the more indicative relation of fitting the representation to the world. This is a major development, but it still relies on the evolutionary history model of representation and content, and encounters problems in-principle in accounting for organism-detectable error for either direction of fit.

[10] Or in addition to Dretske (2000).

This document is copyrighted by the American Psychological Association or one of its allied publishers.
This article is intended solely for the personal use of the individual user and is not to be disseminated broadly.

first layer) being activated by inputs from an environment, and the nodes transmit those activations to other nodes that they are linked to. The links, in turn, have weights that effect the transmitted activation levels and thereby influence what recipient nodes receive; furthermore, these weights can be positive or negative. Nodes receiving weighted activation levels, in turn, determine their own activations as a function of the received activations, and send those resultant activations out along whatever links they have to other nodes.

Some special set of nodes, often a final layer of nodes in a layered architecture, is taken to constitute the output nodes, and the output of the system is constituted as the (vector of) activation levels of those output nodes. The output is taken to be a classification of an input category: a given vector of activations classifies some set of input patterns together into the same category. It is thereby taken to be a representation of that category (and of instances of that category).

A connectionist model can be run on a set of training instances in which resultant output vectors are compared with the outputs that are "supposed" to be obtained—that is, to the categorizations of input patterns that the system is "supposed" to implement—and "errors" in such classifications are fed back into the system in a way that induces modifications of the weights along the links between nodes. Ultimately, a set of weights is sought that does yield the classifications intended by the designer/trainer of the system. That is, a connectionist model is supposed to enable the training of "representational" output vectors.

Thus, connectionist representations are defined in terms of (trained) correspondences with (categories of) what is represented. This is a version of information semantics, except that the correspondences (covariations) are trained rather than being 'transduced.'[11]

In spite of the advantages of trainability, distributed activation vectors, and so on, therefore, as a model of representation, connectionism has all of the problems of information semantics. Error is determined by the 'trainer,' using specialized input for error,[12] or from the perspective of the designer or user or observer. Organism-detectable error is not possible, and the model is refuted.[13]

Dynamic Systems—Tracking

Dynamic systems models focus on the (usually continuous) processes manifested by a system, and on the *organizations of spaces of possible temporal trajectories* of the states of such systems. This focus on spaces of system-state possibilities enables the use of powerful mathematics originating in dynamic systems models in physics (Port & van Gelder, 1995; Jost, 2005). Such models constitute a very sophisticated form of dispositional explanations: a system is "disposed" to follow the trajectory that it starts on, and thereby to honor the organization of the space of such trajectories (Bickhard & Campbell, 2003).

Some dynamic systems modelers have argued that representation is not a necessary or relevant notion: the only point of a model is to get the dynamics right. Against this position, Clark and Toribio (1994) argued that some situations are 'representation hungry'—they require tracking. Keeping track of a predator behind a rock, for example. Such tracking would be constituted by some (sub-)state of the system that varies with—is correlated with—something being tracked. Thus, so the argument goes, even dynamic systems require representation-as-tracking as an aspect of getting the dynamics right.

But tracking as constituting representation requires an observer to determine whether the tracking is in error (it is yet another version of

[11] But *transduction*, in one of Fodor's incarnations (Fodor, 1986), is 'just' generation of representation without inference. So connectionist systems are trained transducers in this sense.

[12] And even this input is not an error input from the perspective of the system. It is simply a special input that, by design, evokes special weight adjusting processes (e.g., back propagation). See Bickhard and Terveen (1995).

[13] There is a framework, derived from multiple sources, but including in particular machine learning, that discusses *internal error* and *error minimization*, and that would prima facie seem to be a counter-example to this claim. But this too is 'error' only from the external perspective of the observer, not from the perspective of the system: these models are intrinsically non-normative. See McDermott (1981) for a discussion of similar misguidance from mislabeling. This general framework of action-oriented predictive processing (Clark, 2013) has multiple complexities, but they do not alter the basic point concerning representation and representational error. The framework is discussed further later in the text, and in more detail elsewhere (Bickhard, in press, 2014a, 2014b; Bickhard & Terveen, 1995).

This document is copyrighted by the American Psychological Association or one of its allied publishers.
This article is intended solely for the personal use of the individual user and is not to be disseminated broadly.

information semantics): to compare what is "supposed" to be being tracked with what is actually (if anything) being tracked. This does not violate Clark's stance (nor Dretske's) because they advocate a notion of representation as *constituted in being attributed* by such an observer (Wheeler & Clark, 1999; Clark, 2001; Dretske, 1988, 2000). Representation, in such a model, is not a dynamic reality, but, rather, an attribution involved in an observer's explanation of how and why the organism is doing whatever it is doing. Among other problems, such observer dependency fails to account for the representations of that observer.

More generally, however, tracking either occurs or not. If so, then the tracking-representation exists and is correct. If not, the representation relationship does not exist. There is no account of representation existing but being in error, except from the perspective of an observer who provides the normative content of what the tracking 'representation'—the tracking state—is *supposed* to represent.

Organism-detectable error, thus, is impossible, and the model is refuted.

Dynamic Systems—No Representation

As mentioned, some dynamic systems advocates (including the article from which the title of this article is derived: van Gelder, 1995) have argued that representation is an unnecessary and misleading concept—that the fundamental modeling task is to get the dynamics right, and that this can be done without any representation at all (Port & van Gelder, 1995).

The model that I will outline is a version of a dynamic systems model, but I will argue that complex agents require that certain functions be performed,[14] and that performing these functions requires processing that yields emergent truth value—thus representation. Representation in the sense of *truth-valued* processes or conditions, therefore, will exist in complex agents whether or not the word 'representation' is used. So, avoiding the term 'representation' simply makes talking about the dynamics of real systems much more difficult. It doesn't change the dynamics of those systems at all.

The Radical Skeptical Argument

Before turning to that model, I would like to point out an indication of the depth of the problem of organism-detectable error, and, perhaps, a reason why it is so universally ignored: the problem of organism-detectable error is the focus of the radical skeptical argument. The conclusions of many skeptical arguments are some kind of fallibilism, and I will not address those. But the radical skeptical argument focuses on our inability to step outside of ourselves to compare our representations with the world to see whether they are correct. So, the multiple and ubiquitous failures to model organism-detectable error are part of the long history of failures to resolve radical skepticism (Popkin, 2003; Popkin & Stroll, 2002; Rescher, 1980; Weintraub, 1997).

This history of failures is reason enough, for some, to ignore radical skepticism: it cannot be solved or resolved, so it should simply be ignored. But, error-guided action and learning do occur, so organism-detectable error occurs, so there has to be something wrong with the radical skeptical argument—representation, including derivative error-guided action and learning, cannot be correctly modeled without correcting whatever has produced such a history of failure.

Radical skepticism is not just an armchair philosophical problem. It manifests a fundamental error in the framing of the problem of the nature of representation that has persisted for millennia.

Just check consequences? An easy and frequent intuitive response to the radical skeptical argument is to point out that, even though checking on a particular correspondence in the moment may not be possible, all one has to do is to check further consequences to see whether the original representation is correct—for example, walk up to a perceived-cow and discover that it is really a horse instead. I will argue that the intuition of *anticipation* in this response is an important aspect of the solution. As formulated, however, this response simply runs into the same problem with respect to the perceptions of the consequences—how do I know that

[14] *Not* tracking—though the strictly *correlational* information involved in tracking can be quite useful, it cannot in itself bear truth value, and, so, cannot be representational.

This document is copyrighted by the American Psychological Association or one of its allied publishers.
This article is intended solely for the personal use of the individual user and is not to be disseminated broadly.

the representations of what is in front of me as being a horse are correct? They may be inconsistent with my earlier representations of a cow in front of me, but there is no more guarantee that these representations of a horse are correct than there was any such guarantee concerning the earlier cow representation.

Thus, I can discover, perhaps, that my representations are not consistent with each other (or that they *are* consistent), but I still cannot bridge the gap from representation to represented. Perhaps, then, it is superfluous to postulate anything on the other end of my representational correspondences; perhaps my representations are all that I have, and there are no such correspondences. My world simply *is* the world of my “representations.” Such concerns have been a strong source of idealisms, and of arguments for idealism: defeat correspondence and conclude idealism. Note that such reasoning by elimination (eliminate correspondence and idealism is all that is left) presuppose that there is no other alternative to standard correspondence models. I argue otherwise below.

The problem in this naïve response to the radical skeptical argument is that the consequences to be “checked” themselves require perception, and, thus, simply reencounter the basic argument. I will suggest that a form of anticipation that is *internal*, thus can be *functionally* checked, and, thus, does *not* require further representation that ‘simply’ encounters radical skepticism again, can resolve this problem. The basic idea is that of functional anticipations of *organism internal* processes, which processes in part depend on interactions with the world, and, thus, in anticipating such *internal* flows of process, there is an *indirect* anticipation of how the processes of the world flow that are engaged in the interaction with the organism.

An Interactivist Pragmatic Model of Representation

I will adumbrate a model of representation—an interactive, pragmatic model—that satisfies all three of the desiderata for models representation mentioned above, including that of organism-detectable error.[15,16]

A complex agent must select and guide its interactions. To do so, it must have some sort of functional indications of what interactions are possible in its current situation, so that it can select among them (Bickhard & Richie, 1983; Bickhard, 2009a, 2009b). It cannot successfully engage in an interaction that is not supported in the current environment—I cannot open the fridge to get a drink if I’m in a forest.

Such indications of interactive potentiality are anticipatory, and, thus, capable of truth value: the indicated interaction, if selected and engaged, may or may not proceed as anticipated—the actual environment may or may not support that interaction: the anticipation may be true or false.[17] Truth value is the central problem of representation—the normative aspect of representation[18]—so anticipations of what interactions are possible yield emergent (primitive) representation as a natural result of the evolution of complex agents.

Content, in this model, is implicit, not explicit as in encoding models. It is constituted in the implicitly defined, implicitly presupposed conditions in the environment that would support the interactive anticipation if those conditions were to obtain. It is that the environment manifests sufficient such supporting conditions that is implicitly predicated of that environment in anticipating that an interaction would flow within anticipated bounds in that environment (Bickhard, 2009a, 2009b).

[15] I will not address here some crucial aspects of the model, such as that of the emergence of normative function in a way that is not subject to the problems of etiological models of function. See Bickhard (1993, 2009a, 2014b).

[16] For early presentations of the model, see Bickhard (1978, 1980a/1973, 1980b, 1980c, 1982). For more developed versions, see Bickhard (1993, 2005, 2009a, 2009b, 2014a, 2014b).

[17] If ‘indications’ of interactive potentialities must themselves be represented, then the model is circular. That such a circularity is not forced is demonstrated by the fact that such indications could be realized as pointers, which are strictly machine-theoretic functional, not representational. For a model of how the brain realizes such anticipation (*not* by pointers), see Bickhard (2011, in press, 2014a, 2014b).

[18] Normativity here means the possibility of being true or false; it does not in itself entail that anything is guaranteed to be true (or false), though logical truths may be modeled in ways that yield such guarantees. Representational normativity—truth and falsity—is part of a large family of normativities, including function/dysfunction; rational/irrational; ethically good/bad; and many others.

This document is copyrighted by the American Psychological Association or one of its allied publishers.
This article is intended solely for the personal use of the individual user and is not to be disseminated broadly.

This document is copyrighted by the American Psychological Association or one of its allied publishers.
This article is intended solely for the personal use of the individual user and is not to be disseminated broadly.

Interactivist Representation

Representation, thus, emerges in future oriented, pragmatist, anticipations of possible interactions, not in past oriented 'spectator models' (Dewey, 1960; Tiles, 1990) attempting to look backward in time down the stream of inputs (e.g., to "see" where they came from).

Indications of potential interactions are primitive loci of emergent truth value, but they do not look much like canonical forms of representation. So here is an outline of some of the resources available in such a model for more complex representation.

More Complex Representation

Consider first a frog, which might have several flies and a worm that it could flick its tongue toward in anticipation of eating—indications of interaction possibilities can *branch* into multiple possibilities, perhaps widely. Furthermore, if the frog moves over a little, perhaps another worm becomes reachable—indications of interaction possibilities can *iterate*, with some interactions setting up the conditions for the possibilities of other interactions. *If* you walk into the kitchen, *then* you can open the refrigerator to get a drink. Such branchings and iterations can generate vast complex webs of interaction indications, especially in the most complex agents.

Consider now a small toy block, for an infant or toddler. The toddler can engage in many interactions, such as manipulations, visual scans, dropping, chewing, and so on. The block affords a special kind of subweb of interaction indications with two special properties: (a) any interaction possibility in the subweb, such as a particular visual scan, is reachable from every other possibility, such as a different visual scan, via some intermediary interaction, such as rotating the block—so the subweb is internally completely reachable; and (b) that subweb remains invariant under many things that the toddler can do and under many things that can happen in the environment, such as throwing the block, putting it into the toy box, leaving it on the floor and leaving the room, and so on. It does not, however, remain invariant if the block is, for example, burned or crushed.

From the epistemic perspective of the toddler, the block *is* whatever affords such a special subweb of potential interactions.[19] This is essentially Piaget's model of the representation of small manipulable objects, translated into interactivist terms (Piaget, 1954). More generally, the interactivist model has resources for addressing complex representation.[20]

Two Comparisons

I would like to address two contemporary approaches in a little more detail to demonstrate that the above critique still applies to these frameworks: they involve the same problems as outlined more broadly above. The two approaches are those of 'enactivism' and a family of related kinds of models that are designated variously 'the predictive brain,' 'the Bayesian brain,' and related names.

Enactivism

Enactivism, descended from Maturana and Varela's autopoietic model (Maturana & Varela, 1980, 1987; Varela, Thompson, & Rosch, 1991), incorporates aspects of dynamic systems frameworks, ecological psychology, and embodied cognition. It shares several intuitions with the interactivist model, perhaps the most basic being that cognition is inherent in the nature of living beings (Maturana & Varela, 1980; Varela, 1997; Weber & Varela, 2002). Compare "knowing as explicated above is an intrinsic characteristic of any living system" Bickhard (1973, p. 8; also Bickhard, 1980a, p. 68).

[19] The notion of affordance is used here with deliberate allusion to Gibson and contemporary ecological psychology research (e.g., Gibson, 1966, 1977, 1979; O'Regan & Noë, 2001). For further discussion of Gibson's model, see (Bickhard, 2005, 2014b; Bickhard & Richie, 1983). Note that any such model of the interactive processes of perceiving is *not* a model of sensory encoding, or feature representations, or transduction, or sensations, or any of multiple other equivalent notions. Inputs *participate* in interactions, and interactions (and interaction indications) can enable and constitute representations, but inputs are never themselves representational. To assume that they are or could be is the primary seduction into the information semantics error—What could they represent? How could they do so (for the organism)? How could they possibly be wrong? How could any such error be organism detectable?

[20] There is much more to be addressed, such as abstractions (e.g., numbers or fictions), language, and so on, but the model has resources for addressing (or at least attempting to address) those too (Bickhard & Richie, 1983; Bickhard & Terveen, 1995; Bickhard, 2009a, 2014b).

This document is copyrighted by the American Psychological Association or one of its allied publishers.
This article is intended solely for the personal use of the individual user and is not to be disseminated broadly.

With regard to representation and cognition more broadly, however, there are problems that trace back to the basic definition of autopoiesis. In particular, autopoiesis focuses on an *internal* perspective on a system, emphasizing the (re-) production of the components of that system. This kind of process will only occur in far from thermodynamic equilibrium conditions, but there is no mention of such conditions (till Di Paolo, 2005), nor any mention of *any* kind of essential relation to the system's environment. With no intrinsic relations to an environment, there is no basis in the framework for modeling representational or cognitive relations with that environment.

This problem is manifest in the basic discussion (Maturana & Varela, 1980). Relations between a system and its environment are discerned (only) by an observer of the system in its environment (though how the observer's cognitions are to be modeled is not clear).[21] But later in the book (roughly pages 38–40), an internal *recursive* process somehow becomes an internal *observer* of the internal system, and magically thereby becomes an observer of the system's relations with its environment—with the perspective of an *external* observer. This nonsequitur is deep in the conceptual framework of autopoiesis and enactivism. Recognition of some of these basic limitations of the framework has been slow (Varela, 1997; Weber & Varela, 2002; but see Bickhard, 2000a, 2014b; Christensen & Hooker, 2000; Di Paolo, 2005; Moreno, Etxeberria & Umerez, 2008).

With respect to the basic issue of representation, enactivism has no model. Enactivist discussions often allude to cognition in some broad sense, but there is no attempt to account for representational truth value or representational error, and there cannot be any successful such attempt so long as the basic autopoietic definitional framework is used. Consequently, there is no, and can be no, model of organism-detectable error, thus error-guided behavior and error-guided learning.[22]

Action-Oriented Predictive Processing[23]

The predictive processing framework proposes that the brain attempts to generate predictions—or expectations or anticipations—of *inputs*.[24] The 'action orientation' derives from recognition that 'predictive' signals might, instead of being matched against incoming actual signals to determine whether the predictions hold, be directed at muscles, thus influencing the environment from which the inputs originate, and, thus, potentially *inducing* the 'expected' inputs instead of just passively receiving them (Clark, 2013; Friston & Stephan, 2007; Friston, 2012).[25]

These are computational models,[26] and the differences between actual inputs and 'predicted' inputs are called 'error,' so it might seem that models in this family would be an exception to the claim that no models on offer in the current literature can account for error and for system detectable error.

However, these models are modern-day variants of classical sensory encoding empiricisms—approaches that originated millennia ago—and these kinds of models haven't succeeded yet. Perhaps the proper conclusion is that it is a fundamentally flawed approach, including in its contemporary action-oriented predictive processing guise. In fact, as outlined above, there are in-principle reasons, in addition to historical reasons, to conclude that encodingism in all of its variants is at root incoherent: among other problems, it cannot resolve the radical skeptical problem. *And* there are alternatives: pragmatist-oriented action models, such as the interactivist model.

[21] This, in fact, commits to an observer idealism, in spite of later disavowals of such an idealism (Maturana & Varela, 1987).

[22] Di Paolo (2005) argues that autopoiesis cannot in itself address adaptivity, and, thus, must be enhanced. One version of his notion of adaptivity is learning ('plastic adaptivity'). This, I would agree, is precisely the right direction (see also Moreno, Etxeberria, & Umerez, 2008). But I would also argue that adaptivity should not be simply added on top of autopoiesis as standardly defined: the definition has to be altered so as to focus on the system's maintenance of the far from thermodynamic-equilibrium conditions that are necessary to the system's existence—(recursive) self-maintenance of its conditions of existence (see Bickhard, 1993, 2009a).

[23] A term adopted from Clark (2013).

[24] For an early notion of outputs *anticipating* inputs, see MacKay (1956, 1969).

[25] For an early notion of outputs *controlling* inputs, see Powers (1973).

[26] Though not necessarily 'computational' in the sense of 'symbol manipulation.' More commonly, they are (computable) mathematical models of purported semantic information processing.

This document is copyrighted by the American Psychological Association or one of its allied publishers.
This article is intended solely for the personal use of the individual user and is not to be disseminated broadly.

Problems With Predictive Encoding

First, note that the 'errors' in such models are *differences* between 'predicted' or 'expected' inputs and actual inputs. They are *errors about the world* only insofar as the inputs constitute the ground for all representations of that world—as aggregates or organizations of input sensations, for example.[27] Second, such inputs have to be construed as Fodor transductions (Fodor & Pylyshyn, 1981; Bickhard & Richie, 1983) to ground the computational inferences that are based on them—*representational* sensations have to be somehow generated from strictly factual, *causal* inputs. This presumed interface from fact to representation, of course, has defeated every attempt to account for it throughout history. It is the fundamental problem of presumed sensory encodings.

Third, there is in fact no normativity at all in these models, except in the views of the designers of the models. They are strictly causal and factual, and they tend to settle in conditions in which differences between generated signals and input signals are minimized. Such differences are 'errors' only from an observer perspective, just as the inputs per se are 'representations' only from an observer perspective.

The interactivist model, however, claims to model the emergence of representational truth value in terms of anticipations, and these predictive encoding models also involve what are called 'predictions,' 'expectations,' and sometimes even 'anticipations'—so perhaps they can model representation in the same way? Despite the use of the same or similar terms, however, there are fundamental differences.

Insofar as there are anticipations in predictive encoding models, they are anticipations of inputs. All representation, all prediction, all anticipation, ultimately is defined relative to the *global-organism input-output interface*. In contrast, in the interactive model, anticipation is anticipation of the (*local*, *internal*) *flow* of interactive processes. These flows will be *influenced by* inputs, and outputs, but it is not the inputs or outputs per se that are anticipated. *It is the anticipating (of process flow) that is representational*—truth valued—*not the inputs*.

Relatedly, in predictive encoding models, the representational connection with the world is mediated strictly by sensory encodings—predictive encoding models are incarnations of semantic information sensory empiricism—encountering all of the problems of encodingism. In the interactivist model, the representational relationship is via interactive implicit definitional presuppositions about that world. This ground of future-oriented dynamic implicit definition instead of direct explicit input-encoding is historically novel, and is what transcends classical problems.

So, in spite of using terms such as 'error,' predictive encoding models are no exception to the conclusion that there are no major models on offer that can account for the possibility of organism detectable representational error.[28]

Enactivism and Predictive Processing

Fundamentally, there is no representational-transduction interface between the mind and the world, or the brain and the world[29]: the sensory encoding framework is fundamentally flawed, and the complexities and sophistications of predictive models do nothing to alter that. Conversely, neither the rejection of representation by some enactivists, nor the observer idealism of the autopoietic model, nor the closed system perspective of the autopoietic model, can frame an acceptable model of representation and the possibility of organism detectable representational error.

Neither enactivism nor predictive encoding frameworks constitute or provide exceptions to the general critiques offered above.

So, What Could Cognition Be?

Cognition is a representational dynamics, in which representation is constituted in pragmatic future oriented indications of potentialities for interaction in complex interactive agents. This is a dynamic model, but it goes beyond dynamic

[27] In many of these models, the resources for constructing higher levels of representations out of sensations—such as of objects, numbers, electrons, the concept *chair*, and so on—are even weaker than this: the models generate *parameters of probability distributions over inputs*, and it has to be assumed that our world is actually represented solely in terms of such parameters (of probability distributions, etc.).

[28] For further discussion of predictive encoding models, see Bickhard (in press).

[29] "Transduction, remember, is the function that Descartes assigned to the pineal gland." (Haugeland, 1998, p. 223)

systems per se, for example, in requiring functionally normative indications of interaction possibilities (Bickhard, 2007).

Interactive representation is emergent in particular kinds of dynamic system processes, and, thus, is emergent whenever those kinds of system dynamics occur (or are constructed)—it satisfies the criterion of emergence.

Interactive representation is anticipatory, and the anticipations of interactive potentialities may be true or may be false—the model has no difficulty accounting for the possibility of representational error.

Interactive representational anticipations are not only capable of being false, engaging such a false anticipation will functionally encounter the failure of the anticipation in the *internal* processes of the organism—thus organism-detectable error is possible, and, therefore, so are error-guided interaction and learning.

References

Albertazzi, L., Libardi, M., & Poli, R. (Eds.). (1996). *The school of Franz Brentano*. Dordrecht, The Netherlands: Kluwer. doi:10.1007/978-94-015-8676-4

Allen, J. W. P., & Bickhard, M. H. (2013a). Stepping off the pendulum: Why only an action-based approach can transcend the nativist-empiricist debate. *Cognitive Development, 28,* 96–133. doi:10.1016/j.cogdev.2013.01.002

Allen, J. W. P., & Bickhard, M. H. (2013b). Beyond principles and programs: An action framework for modeling development. *Human Development, 56,* 171–177. doi:10.1159/000351140

Beaney, M. (1997). *The Frege reader*. Oxford, UK: Blackwell.

Bechtel, W., & Abrahamsen, A. (1990). *Connectionism and the mind: An introduction to parallel processing in networks*. Cambridge, MA: Blackwell.

Bermúdez, J. L. (2010). *Cognitive science*. Cambridge, UK: Cambridge University Press. doi:10.1017/CBO9780511781322

Bickhard, M. H. (1973). *A model of developmental and psychological processes*. PhD Dissertation, University of Chicago.

Bickhard, M. H. (1978). The nature of developmental stages. *Human Development, 21,* 217–233. doi:10.1159/000271586

Bickhard, M. H. (1980a). A model of developmental and psychological processes. *Genetic Psychology Monographs, 102,* 61–116.

Bickhard, M. H. (1980b). Approaches to the study of mind. *Cognition and Brain Theory, 3,* 146–152.

Bickhard, M. H. (1980c). *Cognition, convention, and communication*. New York, NY: Praeger.

Bickhard, M. H. (1982). Automata theory, artificial intelligence, and genetic epistemology. *Revue Internationale de Philosophie, 36,* 142–143, 549–566.

Bickhard, M. H. (1993). Representational content in humans and machines. *Journal of Experimental and Theoretical Artificial Intelligence, 5,* 285–333. doi:10.1080/09528139308953775

Bickhard, M. H. (2000a). Autonomy, function, and representation. *Communication and Cognition—Artificial Intelligence, 17,* 111–131.

Bickhard, M. H. (2000b). Information and representation in autonomous agents. *Journal of Cognitive Systems Research, 1,* 65–75. doi:10.1016/S1389-0417(99)00007-8

Bickhard, M. H. (2005). Consciousness and reflective consciousness. *Philosophical Psychology, 18,* 205–218. doi:10.1080/09515080500169306

Bickhard, M. H. (2007). Mechanism is not enough. In Q. Gonzalez, M. Eunice, W. F. G. Haselager, & I. E. Dror (Eds.), *Mechanicism and autonomy: What can robotics teach us about human cognition and action?* Special issue of *Pragmatics and Cognition, 15,* 573–585.

Bickhard, M. H. (2009a). The interactivist model. *Synthese, 166,* 547–591. doi:10.1007/s11229-008-9375-x

Bickhard, M. H. (2009b). Interactivism. In J. Symons & P. Calvo (Eds.), *The Routledge companion to philosophy of psychology* (pp. 346–359). London, UK: Routledge.

Bickhard, M. H. (2011). The dynamics of acting. *Humana Mente, 15,* 177–187.

Bickhard, M. H. (2014a). *Toward a model of functional brain processes: Central nervous system functional architecture*. Manuscript in preparation.

Bickhard, M. H. (2014b). *The whole person: Toward a naturalism of persons—contributions to an ontological psychology*. Manuscript in preparation.

Bickhard, M. H. (in press). The anticipatory brain: Two approaches. In V. C. Müller (Ed.), *Fundamental issues of artificial intelligence*. Berlin, Germany: Springer (Synthese Library).

Bickhard, M. H., & Campbell, D. T. (2003). Variations in variation and selection: The ubiquity of the variation-and-selective retention ratchet in emergent organizational complexity. *Foundations of Science, 8,* 215–282. doi:10.1023/A:1025046917589

Bickhard, M. H., & Richie, D. M. (1983). *On the nature of representation: A case study of James Gibson's theory of perception*. New York, NY: Praeger.

Bickhard, M. H., & Terveen, L. (1995). *Foundational issues in artificial intelligence and cognitive*

This document is copyrighted by the American Psychological Association or one of its allied publishers.
This article is intended solely for the personal use of the individual user and is not to be disseminated broadly.

science—Impasse and solution. Amsterdam, The Netherlands: Elsevier Scientific.

Campbell, D. T. (1974). Evolutionary epistemology. In P. A. Schilpp (Ed.), *The philosophy of Karl Popper* (pp. 413–463). LaSalle, IL: Open Court.

Campbell, R. J. (1992). *Truth and historicity*. Oxford, UK: Oxford University Press.

Christensen, W. D., & Hooker, C. A. (2000). Autonomy and the emergence of intelligence: Organised interactive construction. *Communication and Cognition—Artificial Intelligence, 17,* 133–157.

Churchland, P. M. (1995). *The engine of reason, the seat of the soul: A philosophical journey into the brain*. Cambridge, MA: MIT Press.

Clark, A. (2001). *Mindware*. Oxford, UK: Oxford University Press.

Clark, A. (2013). Whatever next? Predictive brains, situated agents, and the future of cognitive science. *Behavioral and Brain Sciences, 36,* 181–204. doi:10.1017/S0140525X12000477

Clark, A., & Toribio, J. (1994). Doing without representing? *Synthese, 101,* 401–431. doi:10.1007/BF01063896

Cummins, R. (1996). *Representations, targets, and attitudes*. Cambridge, MA: MIT Press.

Dewey, J. (1960/1929). *The quest for certainty*. New York, NY: Capricorn Books.

Di Paolo, E. A. (2005). Autopoiesis, adaptivity, teleology, agency. *Phenomenology and the Cognitive Sciences, 4,* 429–452. doi:10.1007/s11097-005-9002-y

Dretske, F. I. (1988). *Explaining behavior*. Cambridge, MA: MIT Press.

Dretske, F. I. (2000). The nature of thought. In F. Dretske (Ed.), *Perception, knowledge and belief* (pp. 227–241). Cambridge, UK: Cambridge University Press.

Fodor, J. A. (1975). *The language of thought*. New York, NY: Crowell.

Fodor, J. A. (1981). The present status of the innateness controversy. In J. Fodor (Ed.), *Representations* (pp. 257–316). Cambridge, MA: MIT Press.

Fodor, J. A. (1986). Why paramecia don't have mental representations. In P. A. French, T. E. Uehling, & H. K. Wettstein (Eds.), *Midwest studies in philosophy X: Studies in the philosophy of mind* (pp. 3–23). Minneapolis, MN: University of Minnesota Press.

Fodor, J. A. (1990a). *A theory of content*. Cambridge, MA: MIT Press.

Fodor, J. A. (1990b). Information and representation. In P. P. Hanson (Ed.), *Information, language, and cognition* (pp. 175–190). Vancouver, Canada: University of British Columbia Press.

Fodor, J. A. (1990c). Fodor's guide to mental representation. In J. A. Fodor (Ed.), *A theory of mental content and other essays*. Cambridge, MA: MIT Press. [original 1985]

Fodor, J. A. (2003). *Hume variations*. Oxford, UK: Oxford University Press.

Fodor, J. A., & Pylyshyn, Z. (1981). How direct is visual perception? Some reflections on Gibson's ecological approach. *Cognition, 9,* 139–196. doi:10.1016/0010-0277(81)90009-3

Friston, K. J. (2012). A free energy principle for biological systems. *Entropy, 14,* 2100–2121. doi:10.3390/e14112100

Friston, K. J., & Stephan, K. E. (2007). Free-energy and the brain. *Synthese, 159,* 417–458. doi:10.1007/s11229-007-9237-y

Gibson, J. J. (1966). *The senses considered as perceptual systems*. Boston, MA: Houghton Mifflin.

Gibson, J. J. (1977). The theory of affordances. In R. Shaw & J. Bransford (Eds.), *Perceiving, acting and knowing* (pp. 67–82). Hillsdale, NJ: Erlbaum.

Gibson, J. J. (1979). *The ecological approach to visual perception*. Boston, MA: Houghton Mifflin.

Gill, M.-L. (1989). *Aristotle on substance*. Princeton, NJ: Princeton University Press.

Haugeland, J. (1998). *Having thought*. Cambridge, MA: Harvard University Press.

Jacquette, D. (2004a). *The Cambridge companion to Brentano*. Cambridge, UK: Cambridge University Press. doi:10.1017/CCOL0521809800

Jacquette, D. (2004b). Brentano's concept of intentionality. In D. Jacquette (Ed.), *The Cambridge companion to Brentano* (pp. 98–130). Cambridge, UK: Cambridge University Press.

Jost, J. (2005). *Dynamical systems*. Berlin, Germany: Springer-Verlag.

Levine, A., & Bickhard, M. H. (1999). Concepts: Where Fodor went wrong. *Philosophical Psychology, 12,* 5–23. doi:10.1080/095150899105909

MacKay, D. M. (1956). The epistemological problem for automata. In C. E. Shannon & J. McCarthy (Eds.), *Automata studies* (pp. 235–251). Princeton, NJ: Princeton University Press.

MacKay, D. M. (1969). *Information, mechanism and meaning*. Cambridge, MA: MIT Press.

Maturana, H. R., & Varela, F. J. (1980). *Autopoiesis and cognition*. Dordrecht, Holland: Reidel. doi:10.1007/978-94-009-8947-4

Maturana, H. R., & Varela, F. J. (1987). *The tree of knowledge*. Boston, MA: New Science Library.

McDermott, D. (1981). Artificial intelligence meets natural stupidity. In J. Haugeland (Ed.), *Mind design* (pp. 143–160). Cambridge, MA: MIT Press.

Millikan, R. G. (1984). *Language, thought, and other biological categories*. Cambridge, MA: MIT Press.

Millikan, R. G. (1993). *White queen psychology and other essays for Alice*. Cambridge, MA: MIT Press.

Millikan, R. G. (2005). Pushmi-pullyu representa-

This document is copyrighted by the American Psychological Association or one of its allied publishers.
This article is intended solely for the personal use of the individual user and is not to be disseminated broadly.

tions. In R. G. Millikan (Ed.), *Language: A biological model* (166–186). Oxford, UK: Oxford University Press.

Moreno, A., Etxeberria, A., & Umerez, J. (2008). The autonomy of biological individuals and artificial models. *BioSystems, 91,* 309–319. doi:10.1016/j.biosystems.2007.05.009

Newell, A. (1980). Physical symbol systems. *Cognitive Science, 4,* 135–183. doi:10.1207/s15516709cog0402_2

Newell, A., & Simon, H. A. (1972). *Human problem solving*. Englewood Cliffs, NJ: Prentice Hall.

Newell, A., & Simon, H. A. (1987a). Computer science as empirical inquiry: Symbols and search. In *ACM Turing Award lectures: The first twenty years, 1966–1985* (pp. 287–313). Reading, MA: Addison Wesley.

Newell, A., & Simon, H. A. (1987b). Postscript: Reflections on the tenth Turing Award Lecture: Computer science as empirical inquiry—Symbols and search. In *ACM Turing Award lectures: The first twenty years, 1966–1985* (pp. 314–317). Reading, MA: Addison Wesley.

O'Regan, J. K., & Noë, A. (2001). A sensorimotor account of vision and visual consciousness. *Behavioral and Brain Sciences, 24,* 939–973. doi:10.1017/S0140525X01000115

Piaget, J. (1954). *The construction of reality in the child*. New York, NY: Basic. doi:10.1037/11168-000

Poli, R. (1996). Kazimierz Twardowski 1866–1838. In L. Albertazzi, M. Libardi, & R. Poli (Eds.), *The school of Franz Brentano* (207–231). Dordrecht, The Netherlands: Kluwer.

Popkin, R. (2003). *The history of scepticism*. Oxford, UK: Oxford University Press.

Popkin, R. H., & Stroll, A. (2002). *Skeptical philosophy for everyone*. Amherst, NY: Prometheus Books.

Port, R., & van Gelder, T. J. (1995). *Mind as motion: Dynamics, behavior, and cognition*. Cambridge, MA: MIT Press.

Powers, W. T. (1973). *Behavior: The control of perception*. Chicago, IL: Aldine.

Rescher, N. (1980). *Scepticism*. Totowa, NJ: Rowman and Littlefield.

Rumelhart, D. E., & McClelland, J. L. (1986). *Parallel distributed processing: Foundations* (Vol. 1). Cambridge, MA: MIT Press.

Smith, B. (1994). *Austrian philosophy: The legacy of Franz Brentano*. New York, NY: Open Court.

Tiles, J. E. (1990). *Dewey*. London, UK: Routledge.

van Gelder, T. J. (1995). What might cognition be, if not computation? *The Journal of Philosophy, 92,* 345–381. doi:10.2307/2941061

Varela, F. J. (1997). Patterns of life: Intertwining identity and cognition. *Brain and Cognition, 34,* 72–87. doi:10.1006/brcg.1997.0907

Varela, F. J., Thompson, E., & Rosch, E. (1991). *The embodied mind*. Cambridge, MA: MIT Press.

Vera, A. H., & Simon, H. A. (1993). Situated action: A symbolic interpretation. *Cognitive Science, 17,* 7–48. doi:10.1207/s15516709cog1701_2

Vera, A. H., & Simon, H. A. (1994). Reply to Touretzky and Pomerleau: Reconstructing physical symbol systems. *Cognitive Science, 18,* 355–360. doi:10.1207/s15516709cog1802_6

Weber, A., & Varela, F. J. (2002). Life after Kant: Natural purposes and the autopoietic foundations of biological individuality. *Phenomenology and the Cognitive Sciences, 1,* 97–125. doi:10.1023/A:1020368120174

Weintraub, R. (1997). *The sceptical challenge*. London, UK: Routledge. doi:10.4324/9780203283868

Wheeler, M., & Clark, A. (1999). Genic representation: Reconciling content and causal complexity. *The British Journal for the Philosophy of Science, 50,* 103–135. doi:10.1093/bjps/50.1.103

Received December 22, 2013
Revision received August 5, 2014
Accepted August 21, 2014 ■

This document is copyrighted by the American Psychological Association or one of its allied publishers.
This article is intended solely for the personal use of the individual user and is not to be disseminated broadly.