

THE EMERGENCE OF CONTENTFUL EXPERIENCE

MARK H. BICKHARD

15 University Drive, Lehigh University

Bethlehem, PA 18015

mark.bickhard@lehigh.edu

<http://www.lehigh.edu/~mhb0/mhb0.html>

There are many facets to mental life and mental experience. In this chapter, I attempt to account for some central characteristics among those facets. I argue that normative function and representation are emergent in particular forms of the self-maintenance of far from thermodynamic equilibrium systems in their essential far-from-equilibrium conditions. The nature of representation that is thereby modeled — an interactive, pragmatic form — in turn, forces a number of additional properties of mental process, such as consciousness being inherently contentful and from a situated and embodied point of view. In addition, other properties of interactive representation make strong connections with the central nervous system properties that are found to realize mental experience, such as a field organization of oscillatory and mutually modulatory neural processes.

Keywords: *representation, cognition, representational content, consciousness, normative function, Dretske, Fodor, Millikan, mental experience, situated cognition, embodiment, interactivism, far from equilibrium systems, Piaget, goal directedness, pragmatics, information semantics, encodingism, asymmetric dependence, timing, Turing machines, central nervous system, volume transmitters, silent neurons, modulatory processes*

1 Introduction

There are many facets to mental life and mental experience. Ultimately all of them must be addressed in the overall task of the naturalization of mind. Here I will focus primarily on three aspects of basic consciousness. In particular, basic conscious experience:

1. is a process,
2. that is contentful,
3. from a point of view.

Additional characteristics of mind, such as embodiment and a convergence with functional properties of the central nervous system, emerge in the course of the main line of discussion.¹

Organisms are inherently far from thermodynamic equilibrium; to go to equilibrium is to die. Work must be done in order to maintain the essential far-from-equilibrium conditions, and it must be done in ways and at times that are appropriate to the relevant environmental conditions. Even very simple living systems can exhibit this function of selecting “what to do next”: some bacteria, for example, can swim if they are swimming up a sugar gradient, but tumble if they are swimming down a sugar gradient [22,23]. Together, these interactions with the environment tend to increase the sugar supply available to the system.

I will argue that representation has emerged in the evolutionary answers to such problems of selecting “what to do next”, and that several aspects of both mental experience and central nervous system processing are accounted for by that answer.

¹ Several other aspects — such as perception, motivation, language, development, rationality, sociality, personality, and so on — have been addressed elsewhere [2,4,5,6,8,9,14,16,17,18,24,25,26].

2 Function

The first step in the discussion is a model of the nature and emergence of normative function — function as distinguished from dysfunction. For current purposes, a brief outline of this model will suffice.

Some far-from-equilibrium systems, insofar as they are stable through time at all, depend on external support to maintain that stability. A chemical bath, for example, may be maintained in some far-from-equilibrium condition by pumping various solutions into it, and the maintenance of this activity, in turn, depends on the pumps continuing to work and receive power, and the reservoirs of those solutions remaining full. Some far-from-equilibrium systems, on the other hand, make contributions to their own stability. A candle flame, for example, maintains above combustion threshold temperatures, and, in standard atmospheric and gravitational conditions, induces convection, which brings in fresh oxygen and removes combustion wastes. Far-from-equilibrium systems that make such contributions are, in that sense, *self-maintenant* systems [7].

Such contributions to the maintenance of relevant far-from-equilibrium conditions are *functional* for that system [7]. Conversely, to fail in making such contributions is *dysfunctional* for that system. Functionality, in this sense, is relative to a particular system as reference point: a heart in a parasite may be functional for the parasite but dysfunctional for the parasitized host.

2.1 Etiological Approaches to Function.

This model of function is in contrast to standard etiological approaches [43,59,60]. The central notion in these approaches is that the heart has a function of pumping blood, instead of, say, making heart beat sounds, because it is the evolutionary descendant of prior hearts that were selected for pumping blood, not for making heart beat sounds. A kidney, then, that does not filter blood is not serving the function that it has — is being dysfunctional — since kidneys in general have the function of filtering blood.

Etiological approaches to function model the having of a function as being constituted in having the right kind of evolutionary history. This has a sometimes counterintuitive consequence: if, for example, a lion were to miraculously pop into existence that was molecule for molecule identical to some lion in the zoo, the science fiction example lion would have no functions for any of its organs, because none of them would have the right kind of evolutionary history.² They have, in fact, no evolutionary history at all. Millikan is willing to accept this consequence [59], but although such counterintuitive consequences for purely science fiction thought experiments may be worth accepting if other successes of the model warrant, this example points to a far deeper problem — one that is, I argue, fatal to all such approaches.

In particular, the lion example exemplifies that function, on the etiological account, cannot be defined in terms of the current state of the system. Two systems can be in identical states, such as the two lions, but one of them will have organs with functions and the other not, depending on their histories. But, physics tells us, only the current state of a system can have causal efficacy. Etiological accounts, then, at best provide an epiphenomenal account of function — an account with no causal importance in the world. That is not a successful naturalization of the notion of function.

Note, in contrast, that function understood in terms of contributions to maintaining relevant far-from-equilibrium conditions is a current state definition. It does make a causal

² This is from a discussion by Millikan [59]. The idea would be, for example, if the atoms in the air were to suddenly converge in such a way that they formed a lion. This, of course, is statistically impossible, even though logically possible. Millikan uses the example simply to demonstrate what she claims is a counter-intuitive, but nevertheless acceptable, consequence of the historical approach to function. I argue that there is a deeper and more important issue at stake here.

difference whether or not this flame or that organism remains in far-from-equilibrium conditions. Function, then, emerges in self-maintaining far-from-equilibrium systems, including in particular living systems.

3 Representation

Self-maintaining systems make contributions to their own maintenance, but those contributions are fixed. There is no ability to change to making different kinds of contributions if the environment were to change so that some such change in self-maintaining contributions would be appropriate. Candle flames, for example, cannot shift into a “hunt for fuel” mode when the candle is getting low.

The bacterium, however, *can* make such shifts. Swimming if moving up a sugar gradient but tumbling if moving down a sugar gradient is precisely to do different things in different circumstances so as to contribute to far-from-equilibrium maintenance in ways appropriate to those changing conditions. Such systems tend to maintain their condition of being self-maintaining — they are, in that sense, *recursively self-maintaining* [7].

The key point to note is that such selections on the part of a recursively self-maintaining system are anticipatory in nature, and that, as such, they can be in error. They are anticipatory in that they anticipate that the consequences of engaging in the selected activity, under these conditions, will in fact serve the function of self-maintenance. They can be in error because such anticipations depend on, among other things, the environment, and the environment may not cooperate. The bacterium will swim up a saccharin gradient just as readily as it will swim up a sugar gradient.

This is not a standard usage of “anticipate” because it is not meant in any necessary sense of deliberate or explicit anticipation. It is, instead, a functional sense of anticipate. Some functions — contributions to self-maintenance — depend for the success of their functional contributions on particular things working out, or being the case, in the future as the functional process proceeds. To indicate that such a functional process will be appropriate, or to initiate such a functional process, then, functionally or implicitly anticipates that those necessary supporting conditions will obtain.

Such anticipations constitute the most primitive emergence of representational truth value: There is, first of all, a truth value in the anticipation itself — it is either correct that the activity will be self-maintaining or it is not. Second, that truth value is about the environment: the anticipation constitutes an implicit predication about the environment, viz., this is an environment in which the selected activity is appropriate. And third, it has representational content: the anticipation implicitly defines whatever those environmental properties are that would support the selected activity being successful toward self-maintenance. This is implicit definition in a dynamic generalization of the sense in which a set of axioms implicitly defines the class of models for those axioms [12,48,55]. So, there emerges content, which is about the environment, and which has truth value; this is representation, however primitive.³

3.1 Evolutionary Elaborations.

Such representation, however, is quite primitive. It fits, perhaps, bacteria or paramecia, but what about more complex representation, such as in human beings? I will turn to several ways in which primitive representation can be elaborated, each such elaboration improving the adaptability of the organism.

First, notice that the “selection” of what to do next in the bacterium is a kind of triggering. Under specified conditions — conditions that normally detect sugar gradients

³ This is a pragmatist model of representation, rather than the standard encoding or empiricist models (e.g., [50,65,69]).

— do X, swim perhaps, or do Y, tumble perhaps. Under more complex conditions, there may be more than one potentially appropriate next interaction, and a selection *within* some set of possibilities must be made.

One basic manner in which this more complex kind of selection can be accomplished involves three interrelated innovations beyond the triggering model. First, the relationship to potentially appropriate next interactions must be some sort of *indicative* or *pointer* relationship, not a simple triggering. Second, there must be some basis for making a selection within a set of such indicated potentialities. In general, that basis will involve information about the anticipated *outcomes* of the indicated interactions, should they be selected. That is, choose the interaction on the basis of its expectable outcomes.

But those outcomes, at least in the logically primitive sense, cannot be *represented* outcomes, on pain of a circularity in the basic model of representation. If, however, they are internal outcomes, internal states, perhaps, that are indicated in association with various interaction possibilities, then that information is functionally available inside the system, and does not require a circularity of modeling representation in terms of representation.⁴

Third, there must be some process for using such outcome indication information in the service of selecting next interactions from among those indicated. The basic process architecture within which this can take place is a goal directed system, that selects interactions from among those indicated on the basis of their fit to a current goal. (Goal directedness, however, can also involve architectures that are much less explicit: [18].)

Again, however, a potential circularity threatens. If goal conditions must themselves be represented, then again the model of representation has made necessary use of representation. Goal conditions, however, do not have to be represented in order to be functional (though clearly they can be so represented once representation as a function is already available). Goal conditions need only be *detected*. A goal of raising blood sugar, for example, need only yield a continuation of potentially appropriate activities so long as blood sugar is in fact below some threshold. No representation of blood sugar level is necessary. The bimetallic strip in the classic thermostat example does not represent temperature, but it does detect it; and the set point in the thermostat similarly does not represent temperature, but it does detect when the actual temperature has reached the set point temperature. Such functional relationships of detection are all that are necessary for goal directedness, so this potential circularity too is avoided.

So, the first evolutionary elaboration beyond simple triggerings of activities is the evolution of the ability to make use of information about interaction outcomes in the selection of next interactions. Note that such indications of anticipated outcomes not only make possible the selection of next interactions in a way much more sophisticated than simple triggering, they also permit the system to detect whether or not those indicated outcomes are in fact obtained — they permit the system to detect the truth value of its (still primitive) representations. Such system detectable error, in turn, can be quite useful in guiding further behavior, and is essential for error guided learning [15,18].⁵ Indicated outcomes, then, ground the task solutions for both interaction selection and interaction evaluation.

Goal directed processes are an important elaboration of basic triggered system activities. Another important development occurs with respect to the conditions under which various interactive potentialities are indicated — the processes of detection. The most general manner in which such detections can occur is by *interactive differentiation*. If

⁴ Such a circularity will yield an infinite regress if the circle is followed in an attempt to find some foundational level that breaks out of the circle. Since there is no such level, the unboundedness of the regress follows. Even prior to generating such a regress, however, such a definitional circularity is unacceptable because defining representation in terms of representation does not contribute to the task of understanding representation.

⁵ One aspect of the emergence of such primitive representation is the concomitant emergence of equally primitive motivation [9].

a subsystem engages in interaction with the environment, the internal course of that interaction — and, therefore, the internal outcome of that interaction — will depend in part on the environment being interacted with. Some environments will yield the same internal outcome, while other environments will yield some different internal outcome. The set of possible internal outcomes serves to differentiate the class of possible environments into those that yield outcome A, say, versus those that yield outcome, or final state, B. Such environment differentiations, in turn, can serve as the conditions for further indications of potentiality. Arriving at outcome A, for example, might indicate that interaction Q is possible, while arriving at outcome B might indicate that interaction R and interaction S are both possible.

The set of environments that would yield final state A as outcome are *implicitly defined* by the interaction subsystem that engages in the relevant interaction. As before, this is a dynamic generalization of the sense in which a set of formal sentences implicitly defines its class of models [11,48,55]. Differentiation and implicit definition, then, are duals of each other. Final state A of some subsystem implicitly defines A-type environments, and arriving at A differentiates the current environment as being of type A.

An interactive subsystem with possible final states, therefore, is the basic manner in which conditions for indications of potentiality are set up. But the interactive potentialities that are indicated as possible are themselves interactive subsystems with associated possible final states: the two are the same kinds of system organization. Any interactive subsystem, then, will differentiate environments in accordance with its possible final states — actually engaging in the interaction and arriving at one of the final states differentiates the environment as being of the type implicitly defined by that final state — and any interactive subsystem can be indicated as possible if appropriate prior differentiations have occurred.

This suggests the next important elaboration: indications of interactive potentiality can branch and can iterate. A given differentiation can evoke indications of potentiality of multiple further possibilities: final state A might indicate the potentialities of both P and Q. So the indicative relationships can branch. And if P is engaged, arriving, say, at final state D, that might serve to indicate the potentialities of R, S, and T. Such branched and iterated organizations of indications of interactive potentialities can, in more cognitively sophisticated organisms, be quite complex, forming vast webs of potentiality indications.

It is such webs that constitute the basis for more familiar forms of representation, such as of objects, and do so in a generally Piagetian manner (e.g., [2,62]). The representation of abstractions, such as of electron or the number six, requires still further architectural machinery, but will not be pursued here [23,24]. The most important properties of interactive representation that I will develop for current purposes are those of temporal and functional continuities, which underlie aspects of both phenomenology and central nervous system functioning.

3.2 *Information Semantics.*

First, however, a detour to compare the interactive model with the approach to representation that is dominant in contemporary cognitive science: information semantics. Consider an interactive differentiation that takes place with no outputs. This is no longer a full *interaction*, but a passive processing of inputs. When differentiations can be performed in this manner, they are less costly of time and energy, and such forms of differentiation are ubiquitous in complex organisms. One major class of examples is the sensory tracts and associated “information processing” as neural activity progresses along those tracks [27]: the outcomes of such processing, at any level, implicitly define the environments that would yield those outcomes if encountered.

The important point for current purposes is that such passive differentiation processes are the paradigm of what information semantics approaches to modeling representation submit as examples of *representation*. A differentiation, passive or not, does create an informational — and, perhaps, a nomological and causal — relationship with various

properties in the environment: those properties that support arriving at that internal state. Information semantics would have those properties be the content of the representation that is constituted by that final state. The states involved in the sensory information processing are said to “encode” the environmental properties that they differentiate. The interactive model, in contrast, does not attribute content to such differentiations. Instead, the differentiations are the contentless differentiations upon which contentful indications of further potentialities may be based.⁶

The comparison being made here is with standard models which attribute representational content to “mere” differentiations, especially passive differentiations, such as in so called “sensory encodings”. My claim, in contrast, is that such differentiations, passive or not, do not have any content — they are contentless differentiations. But, such differentiations may serve as the basis for setting up indications of further interactive potentiality, and those indications *can* have content — the content that is implicitly defined in the supporting conditions for those further potentialities. That is, such differentiations may differentiate in fact those kinds of environments in which the indicated interactive potentialities will work. But such a differentiation is not and need not be a representation of whatever the conditions are that will support those indicated interaction potentialities: a detector need not be a representation of what is detected — a differentiator need not be a representation of what is differentiated.

What’s wrong with modeling the differentiations themselves as possessing content? This stance is of millennia-long standing. It is a current version of assuming that representation is constituted by correspondences between the representation and what it represents [33,34,36,37,38,39,46,57,67]. *External* examples of representation do seem to fit this approach: Morse code, blueprints, maps, ciphers, and so on. They form the basis for the never ending appeal of modeling purported *mental* representation in the same mold. But such external representations require an interpreter to know and interpret the correspondences involved, while mental representation cannot require such an interpreter on pain of a classic infinite regress of interpreters interpreting the results of previous interpretations.

This regress problem is just one of a great many fatal flaws in correspondence approaches to mental representation. I will touch upon only a few of them here (see [7,18]). One derives from the fact that correspondence, informational, nomological, isomorphic, and causal relationships exist profusely throughout the universe, while at best an extremely small fraction of them might constitute representational relationships. So something further must be specified to attempt to pick out the special such relationships that are supposed to be representational. There is no consensus about what that additional special qualification might be, and I argue that none of them on offer works, and that none *can* work [7,18]. One perspective on why this is so is to note that, even if some special additional property did succeed in extensionally picking out only those correspondences that are genuinely representational, that would still not constitute a naturalistic model of the nature of the representational content involved for the organism itself. For example, there is one finer differentiation in the class of correspondences that does pick out representational correspondences: those that are genuine encodings, such as Morse code. But genuine encodings require an interpreter in order to provide those encodings with content. This is not a problem for many purposes, but for the purpose of modeling representation and representational content, it merely pushes the problem off onto understanding and modeling the interpreter, and that was the original task in modeling *mental* representation in the first place.

Another fundamental problem has to do with being able to model the possibility of representational error. The problem arises because, if the special “representational” correspondence — or informational relationship, or lawful relationship, or whatever special kind is picked out by a model — exists, then the representation exists, and it is correct. On

⁶ Note that it is not the interactive *model* that makes contentful indications, but, rather, the organism.

the other hand, if that special correspondence does not exist, then the representation does not exist, and therefore it cannot be incorrect. There are multiple attempts to solve this problem, but none that succeeds, and none that even addresses the basic problem of not just the possibility of representational error, but that of *system detectable* representational error.

One such attempt regarding the “simple” possibility of error is that of Jerry Fodor [36,37,39,57]. The central notion of relevance here is that of *asymmetric dependence*. The idea is that the possibilities of false evocations of a representation are asymmetrically dependent on true evocations of that representation, and this asymmetry in the dependence relationships distinguishes true from false possibilities. If, for example, a horse dimly seen on a dark night happens to evoke a representation of a cow, that evocation should somehow be modeled as being false. Fodor’s point is that such evocations by horses on dark nights are dependent on evocations by cows in the sense that if cows did not evoke the representation, then horses on dark nights would not either. But the dependency is not reciprocated: if horses on dark nights never evoked the cow representation, that has no bearing on cows evoking the cow representation. The dependency between the two possibilities is asymmetric.

There are a number of problems with this kind of an account. Here is one of them: a counterexample. Consider the docking of a neural transmitter molecule, dopamine, perhaps, in a receptor on a cell surface, triggering internal activities in the cell. This constitutes a causal, nomological, informational correspondence between the transmitter molecule and the cell activities, but there is no representation involved. Still further, consider a poison molecule, crank, perhaps, that can dock on the same receptors and trigger the same internal activities. Again, there are all the kinds of correspondence relationships anyone could want, and, furthermore, there is an asymmetric dependence of the crank possibility on the dopamine possibility, but there is still no representation [7,56].

Here is another: there is no way for any organism to know about, to be able to determine, what the various asymmetric dependency relations are among its potential evocations of representational elements. Therefore, there is no way for an organism to possess in any relevant sense what the contents are of its own representations — to know what they are supposed to represent. Still further, to detect error in its representations, an organism would have to compare such content (which it does not possess) with the actual entity or property currently being represented — the current contact with the environment [7,18], or target of representation [30] — to determine that they do not fit each other. But representing the current contact, or target, is precisely the original problem of representation all over again. So, system detectable error is simply impossible on this account. Not all representations are in error; not all that are in error are detected as being in error; not all organisms are capable of detecting such error. But system detection of representational error does occur — it underlies error guided behavior and learning — and Fodor’s model (along with virtually all others) renders it impossible [7,13,18]. They are thereby falsified. Fodor wishes to set aside such issues of the epistemology of representation until the metaphysics of representation is clear [39]. In itself, that is an acceptable strategic move, but Fodor’s metaphysics not only does not address the basic problem of representational epistemology, it makes representational epistemology impossible. Fodor’s metaphysics is thereby refuted [56].

Representation as some special form of correspondence has an ancient provenance, and many different kinds of issues concerning such approaches and elaborations of such approaches have been addressed over the millennia (e.g., [44,66,72]). I will truncate this discussion at this point, however, with having shown that such approaches suffer foundational flaws. The interactive model, note, models the possibility of error and of system detectable error with ease. It requires no interpreter. It is a viable candidate as a model of representation and representational content. I return, then, to the main discussion of elaborating further properties of interactive representational systems.

3.3 Continuities.

I will develop two kinds of continuity involved in interactive representation: a functional continuity and a temporal continuity. Consider again the set of possible final states for a differentiating interactive subsystem. I have provided examples of such sets above, always with only two possible final states — usually A and B — for simplicity of presentation. But there is nothing that precludes such a differentiating set from being large in cardinality, or even infinite. In fact, differentiating sets with the size of the real numbers should be expected to be common. Such differentiating sets could be realized as, for example, levels of activation of some neural process, or wavelength of some oscillatory process, and so on. Infinite differentiating sets will not set up discrete indications of potentiality for each element, but, instead, will function more as the setting of parameters for further activity in the system that might be engaged in by that system.⁷ System activity and control flow in such an architecture will involve a generally smooth process of engaging in current interactions as one aspect of an overall process, of which another aspect will be the exploration and following of smooth manifolds of parameterized indications of further potentiality.

I turn now to another form of continuity in the model. The interactive model is of representation emerging naturally out of action systems: representation offers a joint solution to the problems of action selection and action evaluation. Action and interaction, however, require correct timing in order to be successful. Mere speed is not sufficient: an interaction can fail from being too fast just as easily as from being too slow. Interaction has to be appropriately coordinated, and that includes temporal coordination.

Computationalist models, in contrast, are based on computer models, and, ultimately, on Turing machines. But Turing machines cannot model temporal coordination. They cannot model timing. Turing machines function with respect to a *sequence* of actions, but the timing involved in the sequence is arbitrary. Timing per se makes no difference to the Turing machine properties and is invisible to any possible Turing machine processes. If the first step required ten seconds, the second ten centuries, the third ten nanoseconds, and so on, nothing about the Turing machine per se would be different from any other timing [17].

Actual computers, of course, do involve timing, and, in that sense, go beyond Turing machines per se. But they do so with a central clock driving myriads of lock step processes. This is a viable design architecture, but an impossible evolutionary architecture: every evolutionary change in the central nervous system would have to involve simultaneous well-coordinated changes in the processing architecture and in the timing architecture. This is vanishingly improbable even once; it is not possible (in any but a strictly logical sense) for evolutionary time spans of change.

So, the brain does it differently. Put clocks everywhere, and render all functional relationships as relationships among the clocks. This sounds odd when put in terms of clocks, but, if it is recognized that clocks are “just” oscillators, it becomes: make all processes oscillatory and render functional relationships as modulatory relationships among those oscillatory processes. In such an architecture, timing is ubiquitous. It is available anywhere that it is useful, and can be ignored if not. Note that such a framework for an architecture is at least as powerful as a Turing machine: a limit case of one process modulating another is for one process to turn the other on and off, that is, to switch the other on and off. But switches are sufficient for building a Turing machine, so oscillatory and modulatory principles have at least the power of Turing machines. They have, in fact,

⁷ Note that setting parameters does not, in general, in itself suffice to specify a system process or interaction. Parameters blend with each other in influencing further activity; they do not build together like bricks. Parameters are not interaction units out of which more complex such units might be constructed. Instead, they join and blend like themes of interaction [17]. This suggests that themes should constitute a major aspect of functional processing in a complex interactive system.

greater power in that they intrinsically capture timing while Turing machines cannot [17,18].

Brain processes are commonly modeled in terms of the current technological models available. From switch boards to symbol manipulations to connectionist nets, studies of the central nervous system have tended to follow the technological lead. This yields currently, for example, a dominant model of neurons as threshold elements that fire or not depending on incoming activations and inhibitions. The paradigmatic neuron is the classic dendritic arborization leading to the extended axon, with the cell body as an appendage [27]. Of course, there are other kinds of neurons, but they are left out of the general functional picture of by what principles the brain might work.

Much of what we know about how neurons function, however, is not easily accommodated by such models. A large population of neurons never fires — the so called “silent” neurons [20,64]. Neurons and neural circuits can exhibit base line oscillatory, or firing, rates, independent of incoming influences [32,42,51,52,68]. Some neurotransmitters are not restricted to a synaptic cleft, but diffuse throughout a local population of neurons — they are “volume” transmitters [1,40,47,71]. Some neurotransmitter release is not all or none, but is “graded” in accordance with the “not all or none” oscillatory ionic waves reaching the terminal buds [20,41]. Some neurons influence others via “gap junctions” that involve no neurotransmitter at all [32,45,61]. Even the glia seem to be involved in influencing neural activity [47,73]. And so on. All of this deviation from paradigm must be construed as merely implementational on standard accounts, though it is not at all clear why evolution would have crafted so many modes of influence if all that was functionally relevant were threshold switches.

But such a tool box of modulatory relationships among oscillatory processes is precisely what would be expected if the functional principles by which central nervous system operated were those of oscillatory processes modulating each others’ activity. Gap junctions provide an extremely fast and spatially localized influence. Traditional synapses are slower and less localized. Volume transmitters are much slower and affect significant local populations. Silent neurons don’t have to fire in order to modulate other activity. And so on. The interactive model puts timing at the center of any interactive system’s functioning, and timing puts oscillatory and modulatory relationships at the center of the processing architecture of such an interactive system. And the central nervous system manifests multiple properties that are perplexing and at best superfluous on standard views, but are simply an evolutionary toolbox for modulatory relationships from the perspective of the interactive model.

Processes in a complex interactive system, then, can be expected to manifest at least two forms of continuity: functional and temporal. Mental processes that might be emergent in such processes, therefore, should be expected to manifest similar continuities.

4 Brain and Mind: Some Relations

Mental life is a process. It is a process that is inherently contentful: it involves intentionality or “aboutness”. The interactive model generates a model of that process as having an ongoing execution of interaction as one aspect and an ongoing consideration of further potentialities as another aspect.⁸ But the “consideration” of further process potentialities *is* the consideration of representational content. It is the consideration of the contents involved in those anticipations of further potentialities. The interactive model, then, captures mentality as a contentful process.

⁸⁸ This aspect is elsewhere called microgenesis. Microgenesis itself offers a powerful model both for characteristics of central nervous system functioning and for important cognitive capabilities, such as metaphor and heuristic problem solving [15].

Mental process involves continuity in both functional and temporal aspects. Oscillatory processes continuously distributed throughout the central nervous system will manifest the properties of an oscillatory *field*. Mental process, then, should be emergent in fields of processes in the brain. That is, consciousness, at least in its most basic form, should be emergent in central nervous system processes organized as fields [53,54]. Mental life manifests properties of this field organization in levels of activity of the field, fineness of differentiations engaged in, coherence (or lack thereof) of the contents being processed, and truncations of experience corresponding to truncations of field processes, such as in cases of neglect [53,54].

Content in this model is always grounded in differentiation processes and possibilities. Differentiations are inherently indexical and deictic. They are relative to the organism making those differentiations in several senses:

1. They are differentiations that, insofar as they are spatial, are spatial in body centered coordinates — they are differentiations *produced* by interactions that that body engages in, and for the subsequent potential use in the interactions that that body engages in. For example, the toy block is just in front of me. Less indexical location representation requires more sophisticated elaborations of invariance representations. The toy block is behind me, or in my room.
2. They are differentiations only as fine as the organism is capable of making and has found to be useful in further processing. Frogs, for example, typically do not differentiate narrowly enough to distinguish flies from small pebbles tossed in front of them. Frogs have not much needed finer differentiation in their evolutionary history. On the basis of such differentiations, frogs will process the potentiality of tongue flicking and eating⁹, along with other relevant possibilities should they exist, such as mating or the potentialities indicated by differentiating the shadow of a hawk overhead.

Mental life, then, is from *a point of view*, both spatially and functionally. Mental life arises in the framework of the view of the organism on all of its further potentialities, spatial, interactive, goals, values, and so on.¹⁰ Mental life is from a point of view most fundamentally because content is from a point of view. The context independent notion of encoded content is a myth. It is impossible because mental representation cannot fundamentally be constituted as encodings. Achievement of relative context independence, of greater scope of invariance, *is* an achievement, on both an individual as well as a cultural level — in science, for example [49].

That is, mental life is inherently *situated*. It is relative to the situation of the organism, again most fundamentally because content is situated. Similarly, mentality is *embodied*. Interaction cannot take place except by some body or another. Mentality is not possible in an inherently passive system — such as a computer that only processes inputs. Mental point of view, then, is situated in the entire representable realm of its further interactive potentialities; it is situated spatially and functionally and relative to the embodiment in which that mental process is taking place.

5 Conclusions

Mental life is a process that is inherently contentful, inherently embodied, and inherently from a situated point of view. The interactive model accounts for these properties as intrinsic aspects of interactive processes. In fact, once the relevant aspects of the interactive model are elaborated, the emergence of these corresponding aspects of mentality is automatic and completely natural.

⁹ Note that the frog's content is that of tongue flicking and eating, not that of “fly” or “pebble” or “fly or pebble” [15].

¹⁰ See Campbell & Bickhard [24] for a model of the emergence of values within interactive systems.

The interactive model also accounts for otherwise puzzling characteristics of the central nervous system processes in which mind is emergent. In particular, the field characteristics of functional and temporal continuity, and the underlying biochemical level of oscillatory processes engaged in mutual modulations, together with the elaborate neural modulatory tool kit, are also automatic and completely natural from the interactive perspective.

The interactive model, thus, accounts in a very natural way for multiple properties of both mind and brain. There are, of course, important characteristics not addressed here, such as those of qualia, emotions, reflexivity, and others,^{11,12} but the naturalness with which the interactive model connects with multifarious properties of both phenomenology and brain processes encourages exploration of further mental characteristics within the interactive framework.

¹¹ The vast and rapidly growing recent literature addressing the phenomena of consciousness includes: Block, Flanagan, Güzeldere [19], Chalmers [28], Cohen & Schooler [29], Dennett [31], Flanagan [35], Marcel & Bisiach [58], Revonso & Kamppinen [63], and Tye [70].

¹² Mind is not emergent in all of its properties at once from underlying functional and physico-chemical processes. This is evident, for example, from a consideration of evolution and non-human animals: not all animals are capable of reflective consciousness; not all are capable of emotions; not all are capable of learning. Necessarily, then, at least these properties must be differentiable from mind in its simplest form. Nevertheless, there is still a strong vestige on the contemporary scene of Cartesian dualism, not in an explicit dualism per se, but in the presupposition that mind differs from the non-mental in some kind of singular gulf [12]. Instead, mind seems to have evolved through a complex trajectory, involving learning, perception, emotions, reflective consciousness, and so on. If so, then these mental phenomena must be modeled as emergent in evolutionary elaborations of simple mental awareness [3,24].

6 Bibliography

1. Agnati, L. F., Fuxe, K., Pich, E. M., Zoli, M., Zini, I., Benfenati, F., Härfstrand, A. and Goldstein, M., Aspects on the Integrative Capabilities of the Central Nervous System: Evidence for 'Volume Transmission' and its Possible Relevance for Receptor-Receptor Interactions. In *Receptor-Receptor Interactions*, ed. by K. Fuxe and L. F. Agnati (Plenum, New York, 1987) pp. 236-249.
2. Bickhard, M. H. *Cognition, Convention, and Communication*. (Praeger Publishers, New York, 1980).
3. Bickhard, M. H. A Model of Developmental and Psychological Processes. *Genetic Psychology Monographs* **102** (1980) pp. 61-116.
4. Bickhard, M. H. The Social Nature of the Functional Nature of Language. In *Social and Functional Approaches to Language and Thought*, ed. by Maya Hickmann (Academic, New York, 1987).
5. Bickhard, M. H. A Pre-Logical Model of Rationality. In *Epistemological Foundations of Mathematical Experience*, ed. by Les Steffe (Springer-Verlag, New York, 1991) pp. 68-77.
6. Bickhard, M. H. How Does the Environment Affect the Person? In *Children's Development within Social Contexts: Metatheory and Theory*, ed. by L. T. Winegar and J. Valsiner (Erlbaum, Mahwah, NJ, 1992) pp. 63-92.
7. Bickhard, M. H. Representational Content in Humans and Machines. *Journal of Experimental and Theoretical Artificial Intelligence*, **5** (1993) pp. 285-333.
8. Bickhard, M. H. Intrinsic Constraints on Language: Grammar and Hermeneutics. *Journal of Pragmatics*, **23** (1995) pp. 541-554.
9. Bickhard, M. H. Is Cognition an Autonomous Subsystem? In *Two Sciences of Mind*. ed. by S. O'Nuallain, P. McKevitt and E. MacAogain (John Benjamins, Amsterdam, 1997) pp. 115-131.
10. Bickhard, M. H. Cognitive Representation in the Brain. In *Encyclopedia of Human Biology. 2nd Ed.* ed. by Dulbecco (Academic Press, New York, 1997) pp. 865-876.
11. Bickhard, M. H. A Process Model of the Emergence of Representation. In *Emergence, Complexity, Hierarchy, Organization*, Selected and Edited Papers from the *ECHO III Conference, Espoo, Finland, August 3-7. Acta Polytechnica Scandinavica, Mathematics, Computing and Management in Engineering Series No. 91*, ed. by G. L. Farre and T. Oksala (1998) pp. 263-270.
12. Bickhard, M. H. Levels of Representationality. *Journal of Experimental and Theoretical Artificial Intelligence*, **10** (1998) pp. 179-215.
13. Bickhard, M. H. Interaction and Representation. *Theory and Psychology*, in press.
14. Bickhard, M. H. and Campbell, R. L. Some Foundational Questions Concerning Language Studies: With a Focus on Categorical Grammars and Model Theoretic Possible Worlds Semantics. *Journal of Pragmatics*, **17(5/6)**, (1992) pp. 401-433.
15. Bickhard, M. H., Campbell, R. L. Topologies of Learning and Development. *New Ideas in Psychology*, **14(2)**, (1996) pp. 111-156.
16. Bickhard, M. H. and Christopher, J. C. The Influence of Early Experience on Personality Development. *New Ideas in Psychology*, **12(3)**, (1994) pp. 229-252.
17. Bickhard, M. H. and Richie, D. M. *On the Nature of Representation: A Case Study of James Gibson's Theory of Perception*. (Praeger Publishers, New York, 1983).
18. Bickhard, M. H. and Terveen, L. *Foundational Issues in Artificial Intelligence and Cognitive Science: Impasse and Solution*. (Elsevier Scientific, Amsterdam, 1995).
19. Block, N., Flanagan, O. and Güzeldere, G. *The Nature of Consciousness*. (MIT, Cambridge, MA, 1997).

20. Bullock, T. H. Spikeless Neurones: Where do we go from here? In *Neurones without Impulses*. ed. by A. Roberts and B. M. H. Bush (Cambridge University Press, Cambridge, 1981) pp. 269-284.
21. Campbell, D. T. Evolutionary Epistemology. In *The Philosophy of Karl Popper* ed. P. A. Schilpp (Open Court, LaSalle, IL, 1974) pp. 413-463.
22. Campbell, D. T. Levels of Organization, Downward Causation, and the Selection-Theory Approach to Evolutionary Epistemology. In *Theories of the Evolution of Knowing*. ed. by G. Greenberg and E. Tobach (Erlbaum, Hillsdale, NJ, 1990) pp. 1-17.
23. Campbell, R. L. A Shift in the Development of Natural-Kind Categories. *Human Development*, **35(3)** (1992) pp. 156-164.
24. Campbell, R. L. and Bickhard, M. H. *Knowing Levels and Developmental Stages*. (Karger, Basel, Switzerland, 1986).
25. Campbell, R. L. and Bickhard, M. H. Clearing the Ground: Foundational Questions Once Again. *Journal of Pragmatics*, **17(5/6)**, (1992) pp. 557-602.
26. Campbell, R. L. and Bickhard, M. H. Types of Constraints on Development: An Interactivist Approach. *Developmental Review*, **12(3)**, (1992) pp. 311-338.
27. Carlson, N. R. *Physiology of Behavior*. (Allyn and Bacon, Boston, 1986).
28. Chalmers, D. J. *The Conscious Mind*. (Oxford University Press, Oxford, 1996).
29. Cohen, J. D., Schooler, J. W. *Scientific Approaches to Consciousness*. (Erlbaum, Mahwah, NJ, 1997).
30. Cummins, R. *Representations, Targets, and Attitudes*. (MIT, Cambridge, MA, 1996).
31. Dennett, D. C. *Consciousness Explained*. (Little, Brown, Boston, 1991).
32. Dowling, J. E. *Neurons and networks*. (Harvard University Press, Cambridge, MA, 1992).
33. Dretske, F. I. *Knowledge and the Flow of Information*. (MIT Press, Cambridge, MA, 1981).
34. Dretske, F. I. *Explaining Behavior*. (MIT Press, Cambridge, MA, 1988).
35. Flanagan, O. *Consciousness Reconsidered*. (MIT, Cambridge, MA, 1992).
36. Fodor, J. A. *Psychosemantics*. (MIT Press, Cambridge, MA, 1987).
37. Fodor, J. A. *A Theory of Content*. (MIT Press, Cambridge, MA, 1990).
38. Fodor, J. A. Information and Representation. In *Information, Language, and Cognition*. ed. by P. P. Hanson (University of British Columbia Press, Vancouver, 1990) pp. 175-190.
39. Fodor, J. A. *Concepts: Where Cognitive Science went wrong*. Oxford University Press, Oxford, 1998).
40. Fuxe, K. and Agnati, L. F. *Volume Transmission in the Brain: Novel Mechanisms for Neural Transmission*. (Raven, New York, 1991).
41. Fuxe, K. and Agnati, L. F. Two Principal Modes of Electrochemical Communication in the Brain: Volume versus Wiring Transmission. In *Volume Transmission in the Brain: Novel Mechanisms for Neural Transmission*. ed. by K. Fuxe and L. F. Agnati (Raven, New York, 1991) pp. 1-9.
42. Gallistel, C. R. *The Organization of Action: A New Synthesis*. (Lawrence Erlbaum, Hillsdale, NJ, 1980).
43. Godfrey-Smith, P. A Modern History Theory of Functions. *Nous*, **28(3)**, (1994) pp. 344-362.
44. Graeser, A. (1978). The Stoic theory of meaning. In *The Stoics*. ed. by J. M. Rist (University of California Press, Berkeley, CA, 1978).
45. Hall, Z. W. *Molecular Neurobiology*. (Sinauer, Sunderland, MA, 1992).
46. Hanson, P. P. *Information, Language, and Cognition*. (Oxford University Press, Oxford, 1990).
47. Hansson, E. (1991). Transmitter receptors on astroglial cells. In *Volume transmission in the brain: Novel mechanisms for neural transmission*. ed. by K. Fuxe and L. F. Agnati (Raven, New York, 1991) pp. 257-265.

48. Hilbert, D. *The Foundations of Geometry*. (Open Court, LaSalle, IL, 1971).
49. Hooker, C. A. Physical Intelligibility, Projection, Objectivity and Completeness: The divergent ideals of Bohr and Einstein. *British Journal for the Philosophy of Science*, **42**, (1992) pp. 491-511.
50. Hookway, C. *Peirce*. (Routledge, London, 1985).
51. Kalat, J. W. *Biological Psychology*. 2nd Edition. (Wadsworth, Belmont, CA, 1984).
52. Kandel, E. R. and Schwartz, J. H. *Principles of Neural Science*. 2nd ed. (Elsevier, New York, 1985).
53. Kinsbourne, M. Integrated Field Theory of Consciousness. In *Consciousness in Contemporary Science*. ed. by A. J. Marcel and E. Bisiach (Oxford University Press, Oxford, 1988) pp. 239-256.
54. Kinsbourne, M. (1997). What Qualifies a Representation for a Role in Consciousness? In *Scientific Approaches to Consciousness*. ed. by J. D. Cohen and J. W. Schooler (Erlbaum, Mahwah, NJ, 1997) pp. 335-355.
55. Kneale, W. and Kneale, M. *The Development of Logic*. (Clarendon, Oxford, 1986).
56. Levine, A. and Bickhard, M. H. Concepts: Where Fodor Went Wrong. *Philosophical Psychology*. (in press).
57. Loewer, B. and Rey, G. *Meaning in Mind: Fodor and his critics*. (Blackwell, Oxford, 1991).
58. Marcel, A. J. and Bisiach, E. *Consciousness in Contemporary Science*. Oxford University Press, Oxford, 1988).
59. Millikan, R. G. *Language, Thought, and Other Biological Categories*. (MIT Press, Cambridge, MA, 1984).
60. Millikan, R. G. *White Queen Psychology and Other Essays for Alice*. (MIT Press, Cambridge, MA, 1993).
61. Nauta, W. J. H. and Feirtag, M. *Fundamental Neuroanatomy*. (Freeman, San Francisco, 1986).
62. Piaget, J. *The Construction of Reality in the Child*. (Basic, New York, 1954).
63. Revonsuo, A. and Kamppinen, M. *Consciousness in Philosophy and Cognitive Neuroscience*. (Erlbaum, Mahwah, NJ, 1994).
64. Roberts, A. and Bush, B. M. H. *Neurons without Impulses*. (Cambridge University Press, Cambridge, 1981).
65. Rosenthal, S. B. Meaning as Habit: Some Systematic Implications of Peirce's Pragmatism. In *The Relevance of Charles Peirce*. ed. by E. Freeman La Salle, IL: (Monist, LaSalle, IL, 1983) pp. 312-327.
66. Sanches, F. *That Nothing is Known*. Cambridge University Press, Cambridge, 1988/1581).
67. Stich, S. and Warfield, T. A. *Mental representation : a reader*. (Blackwell, Oxford, UK, 1994).
68. Thatcher, R. W. and John, E. R. *Functional Neuroscience Vol. 1 Foundations of Cognitive Processes*. (Erlbaum, Hillsdale, NJ, 1977).
69. Tiles, J. E. *Dewey*. (Routledge, London, 1990).
70. Tye, M. *Ten Problems of Consciousness*. (MIT Press, Cambridge, MA, 1995).
71. Vizi, E. S. *Non-synaptic Transmission Between Neurons: Modulation of Neurochemical Transmission*. (Wiley, New York, 1984).
72. Wittgenstein, L. *Tractatus Logico-Philosophicus*. (Routledge, New York, 1961).
73. Yuan, L. and Ganetzky, B. A Glial-Neuronal Signaling Pathway Revealed by Mutations in a Neurexin-Related Protein. *Science*, **283**, (1999) pp. 1343-1345.

Biographical sketch:

Mark Bickhard's work focuses on models of the mind and of persons more broadly. A central axis of this work is a model of the nature of representation — a pragmatic, or interactivist, model. The implications of the interactive model of representation ramify throughout all domains of the study of the mind —the classical cognitive backbone of perception, memory, cognition, reasoning, and language, but also, for example, emotions, consciousness, motivation, and so on. A graduate of the University of Chicago, 1973, he was at the University of Texas at Austin until 1990, when he moved to Lehigh University to accept the Henry R. Luce Professorship of Cognitive Robotics and the Philosophy of Knowledge.

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