

**Radical Embodied Cognitive Science.** Anthony Chemero. Cambridge, Massachusetts: MIT Press, 2009, 244 pages, \$30.00 hardcover.

*Reviewed by Rick Dale, The University of Memphis*

---

Anthony Chemero's *Radical Embodied Cognitive Science* is obviously boldly entitled. It has bold goals, and cuts across an impressive range of topics. It is filled with diverse interweaving threads — topics of interest to the full disciplinary range of the cognitive sciences, from psychology to philosophy. For example, any cognitive scientist interested in a basic summary of philosophical theories of representation would find Chapter 3 invaluable, as it is one of the clearest reviews of this conceptually challenging area that I know of. As another example, Chemero draws out the consequences of the radical embodied approach for the philosophy of mind, a rare agenda (see also, e.g., Noë, 2005). Its inclusion of both conceptually challenging philosophical theories, along with technically sophisticated empirical review (lots of it a review of Chemero's own empirical work), makes the book a cognitive science book *par excellence*.

For the sake of full disclosure, this review writer is a psychologist (with a smattering of philosophical training) whose graduate training leans thoroughly inside the radical embodied camp (e.g., Spivey and Dale, 2004). I am familiar with Chemero's work, philosophical and empirical, and have a very high opinion of it. This review will therefore seem quite enthusiastic in many places. With full disclosure now in place, I do this shamelessly. Yet, I will indeed compensate for this tendency in critical sections of this review, and identify what I feel to be key weaknesses of radical embodied cognitive science. And this is not just for the sake of argument. I believe the radical embodied approach is limited by fundamental meta-theoretical attitudes that keep it from becoming a more common player in the domains of cognitive science, empirical and philosophical.

Having now cleared the air, this review has three primary goals. The first is some good old exegesis: I summarize the book's primary claims and emphasize what are, in my view, their strongest characteristics. The second is to identify points of weakness in radical embodied cognitive science as they are expressed in the book and elsewhere. I do this with the help of some tongue-in-cheek quantitative textual analysis. The final goal is to advance a perspective that I feel addresses these problems, yet still retains some aspects of the radical embodied agenda that Chemero lays out in his great book.

### Hegelian Arguments

Before we consider some central aspects of the book, I let readers in on a dirty little secret that the book's title conceals: Chemero is no radical. The book does develop a certain kind of radical theoretical system for cognitive science, but he warns the reader in many places that it is just one account among many. The purpose of the book is to offer some promising glimpses of how far one can take such a perspective in understanding the phenomena of cognitive science. At the book's very outset, Chemero states (with remarks of humility) that he sees his book as doing for "non-representational, embodied, ecological psychology what Fodor's *The Language of Thought*" did for "rationalist, computationalist psychology" (p. ix). But anyone who has followed the sometimes sharp "framework debates" of cognitive science in the past few decades will perhaps be surprised by some aspects of this book that make it an exception to certain expectations: Chemero leaves room for other accounts.

One important example of this is a lengthy and fascinating discussion (pp. 3–16) of what Chemero calls *Hegelian arguments*: "... arguments that purport to show that" some research or theoretical approach "can never succeed; indeed, nearly every book written by a philosopher begins with an argument that the competing approaches are hopeless. Yet, for some reason, we persist" (p. 3). He provides intriguing examples, linking the argument styles of modern cognitive science to nascent scientific and medieval accounts of digestion. Consider an example Chemero supplies. Noam Chomsky is well known for having cast aside any research on learning and language through casual observations of the supposed impoverished language environment children are provided. The conclusion Chomsky draws in this argument is that *learning* a language is simply *not possible* in this linguistic environment. Chemero points out that, in early versions of this argument, Chomsky's premises were offered with virtually no rigorous empirical evidence for them (p. 7). This makes Chomsky's argument a Hegelian one: it seeks to cast aside an entire research program, in this case the application of learning principles to language acquisition, by premises unanchored to empirical work.

The upshot of all this discussion is that, for Chemero, his bet is "that the empirical facts will ultimately show that we need more than one theoretical approach to cognitive science" (p. 16). The book starts and ends with such caveats:

I have not shown that radical embodied cognitive science is the one true story about the mind or cognition or even perception-action. No clever philosophical argument can do that. Indeed, I don't believe radical embodied cognitive science is the one true story, because I doubt there is any one true story. It is, though, a promising step in our ongoing attempt to naturalize cognition, to see human experience as part of the natural order. And I think that it is the best way to do cognitive science right now. (p. 208)

The book is therefore no Hegelian argument. Its aim is to unobtrusively proselytize, while acknowledging a plurality of viewpoints. Chemero invites readers of any theoretical persuasion to consider the promise of the explanatory tools of radical embodied cognitive science. He casts aside any pretensions of winning an ideological war by magically uprooting other frameworks for understanding cognition. The title of the book is more a statement of Chemero's process: let's see how far radical embodied cognitive science can take us. To begin in this review, then, let me consider what radical embodied cognitive science is, exactly.

### Radical Embodied Cognitive Science

Radical embodied cognitive science is the culmination of a few related approaches to the mind in the past century or more. Chemero spends several pages detailing the intellectual heritage of this approach (pp. 17–32). Its roots are the eliminativistic and naturalistic strains of thinking that permeate the descendants of American pragmatism, which focus not on “hidden” abstract entities that make up our mind (or anything else in the universe), but focus instead on the contextual practices of human beings, and organisms more generally. One of these descendants, the ecological approach to perception founded in large part by the writings of J.J. Gibson (e.g., 1979), plays a special role in the origin of radical embodied cognitive science. Gibson and related theorists emphasize the organism–environment system as the central aspect of a science of perception and cognition, and that our explanations should come not from latent, abstract unseen entities (mental “representations” or “computations”) but instead from an understanding of the space of possible actions that organisms can carry out, and how these actions relate to the structure of the environment in which the organism is embedded.

This intellectual history brings about a theoretical framework that has the following key claims. First, as a negative thesis of the approach, radical embodied cognitive scientists argue that we should *not* seek computational and representational accounts of cognition (p. 29). A radical embodied approach does not look to explanatory schemes of “mental gymnastics,” as Chemero calls them (p. 18). The historical reasons behind this claim are many, among them is one that may be described as another key claim of the approach, a positive thesis about the role of context in our explanations: cognition is best described as the functioning of the organism–environment system. Computational–representational accounts tuck cognition away into a “sub-personal” abstract realm of mechanistic operation (these days it’s proposed to lie inside our cranium); radical embodied cognitive science instead sees cognition as the unfolding of lawful relations between body and environment. Chemero doesn’t offer this claim in his list of definitive claims (p. 29), but the book’s focus on the organism and environment together makes it a central characteristic of the approach. Finally, as another positive claim regarding tools of this approach, radical embodied cognitive science seeks to employ a set of methodologies (among them the formalisms and concepts of dynamical systems) that can work to explain cognition yet do not invoke mental gymnastics (p. 29).

Chemero is concerned that the first and negative claim sounds too much like a Hegelian argument, and stays consistent with his caveats mentioned earlier by not pursuing it much. Instead, the book thoroughly reviews the insights offered by the two positive theses. First, the contextual approach to cognition offers explanations for a promisingly rich array of cognition, primarily through the ideas of ecological psychology (such as affordances, see below). Second, the “tools of choice” of radical embodied cognitive science provide a *generative* explanatory framework for a similar range of cognitive abilities. They are, as Chemero describes, “guides to discovery” (p. 79). Key among these tools in the past few decades is the so-called Haken–Kelso–Bunz model for coordinated oscillators (famous in psychology as a model for coordination of “two-finger twiddling”; Haken, Kelso, and Bunz, 1985; Kelso, 1995). I consider these guides to discovery next.

#### Chemero’s Chosen Guides to Discovery

Chemero, after providing an interesting case study of the debate between Mach and Boltzmann regarding the atomic theory of matter (pp. 78–82), argues that theories

survive and thrive by generating “guides to discovery”: “Atomism, then, is the best methodology for physics because it provides a *guide to discovering* new equations that describe the phenomena more accurately; by assuming that there are atoms, one is led to testable predictions of new phenomena” (p. 79, his emphasis). The radical embodied account thus needs its own guides to discovery. Chemero identifies two important ones just mentioned: Gibson’s ecological psychology, and a suite of powerful tools untainted by mental gymnastics. I detail both here, starting with Gibson. His ecological approach to perception already had the key features of radical embodied cognitive science. For example, in relation to the negative thesis mentioned above: Gibson “disagreed with the tradition that took the purpose of visual perception to be the internal reconstruction” of the external world (p. 106) and “rejected . . . views of perception in which [it] results from the addition or processing of information in the mind . . .” (p. 106).

Yet, more importantly for positive “guides to discovery,” Gibson developed two core concepts that have founded a whole school of thought in the study of perception and action: direct perception and affordances. Direct perception describes the working theoretical assumption that our perceptual systems pick-up information in an unmediated manner. It is most intuitively understood by considering the structure of light information available for vision (the primary domain Gibson concerned himself with was vision). Light energy reflected off the surfaces of objects in the environment informs the organism about what can be done in that environment — explaining visual perception is about understanding how that light energy specifies the environment in important ways (p. 109).

The second core concept relates to this idea: perception of the environment is perception of affordances. The various surfaces, relative placement of objects, etc. provide opportunities for action in the environment. The relationship between the object and the organism is thus an “affordance.” In this way, direct perception of objects is perception for how the objects “afford the organism” opportunities for action. The handle of a hammer affords hammering; the seat of a chair affords sitting; staircases afford stair climbing, etc. Gibson argued for the tight intertwining of perception and action, and Chemero describes some classic studies showing that characteristics of one’s body influence various aspects of visual perception (e.g., Warren, 1984).

Chemero’s goals in detailing these Gibson-derived guides to discovery is to develop them such that they can apply more broadly in a conceptually coherent way. For example, Chemero provides a conceptual analysis of “information” (in Gibson’s, not Claude Shannon’s, sense) so that the theory of direct perception can apply to such things as individuals of categories (e.g., perceiving a *particular* tree), social individuals (e.g., *your* mother), and linguistic information and conventions (pp. 112–113). Such applications extend direct perception and affordances to higher-levels of cognition, to which many still suspect radical embodied cognitive science cannot apply. In the case of defining “information,” Chemero extends the reigning perspective on Gibsonian information (which he dubs the Turvey–Shaw–Mace approach, p. 109) by employing John Barwise and colleagues’ situation semantics (p. 116; Barwise and Perry, 1981). Chemero’s sophisticated framework specifies local, contextual relations that define when visual information specifies or *informs* about states of affairs (the example he uses is the visual presence of a beer can informing about the presence of beer: a contingent, cultural conventional relation that is not a matter of natural law; pp. 118–120). Situation semantics embraces how *specific* individuals play a role in this causal relationship as well (e.g., a *particular* can of beer), allowing direct perception to function in specific contexts rather than at the level of lawful, natural categories

of things, as was originally described in early theories (e.g., “stand-on-able” as a natural relational class between organism and perceived object). In short, it is permissible in this framework to say a person can “directly perceive beer-presence” in the presence of a beer can (p. 120).

In an application to a problem of high-level cognition, Chemero details some previous work he and colleagues conducted on a simple neural network model of analogical reasoning (pp. 131–133). Their basic input–output neural model was trained to identify whether an array of visual patterns all contained the same sub-pattern, or contained at least one distinct pattern (akin to *Sesame Street*’s “one of these things is not like the other”). The network simply responded “same” or “different” through activation of one of two output neurons. Chemero and colleagues obtain the expected result: the simple network (known as a simple perceptron) can do this task. The conclusion Chemero draws is that it “suggests very strongly that pigeons and baboons perceive sameness and difference by directly perceiving entropy” (p. 133) because the neural network does not have a hidden layer (a layer of neurons mediating between input and output). [A problem here is that no neural network modeler would say that a perceptron — even one with a single layer — is somehow “directly” perceiving anything. Its computational properties are well-known, and given its limitations (that Chemero discusses briefly on p. 131), its standing as an archetypal “direct-perceiving” system would mean that computational problems in vision are virtually never unmediated as the inability of simple perceptrons to capture crucial properties of perception have been well-known since Minsky and Papert (1969). Chemero tucks this point into a footnote, but it is important to point out here.]

I don’t want to give away all of the interesting pursuits found in Chapters 6 and 7, where Chemero lays out the Gibson-derived guides to discovery, and demonstrates their application. I should mention that Chemero details the nuanced debate that has occurred on the nature of affordances, and offers his own account of an affordance that is more dynamic compared to its predecessors (p. 150). His comprehensive and flexible approach also makes room for ways in which organisms, through their actions, can also select and shape the structure present in an organism’s niche. A detailed example of affordances is presented in the haptic domain, in the fascinating area of study known as “dynamic touch” (pp. 154–157; see Turvey and Carello, 1996). In all, Chemero’s developments of these two concepts, direct perception and affordances, appears to make them more amenable to the diverse range of timescales involved in explaining cognition (e.g., evolutionary concerns) and appears to have the *potential* to better integrate high-level cognitive issues (though see below for problems with this) all while remaining unanchored to mental gymnastics.

Gibson’s ideas are undergoing continued development by a range of scholars who are well known in perception and action psychology for both their theoretical and empirical extensions of them (Michaels and Carello, 1981; Shaw, Turvey, and Mace, 1982; Turvey, 1977, 1992; see Chemero, Chapters 6–7 for further and more recent review). In addition to these Gibson-derived guides to discovery, Chemero details a few additional guides that fall under a general class of methodological and mathematical tools that do not invoke mental gymnastics. A primary example of this, discussed in detail in the book, is the Haken–Kelso–Bunz (HKB) model for coordination dynamics (Haken et al., 1985). Chemero likens this model to the kind of hypothesis-generating guide to discovery that the atomic theory succeeded in being for Boltzmann when he argued against Mach’s phenomenological physics. Chemero reviews a range of impressive applications of the HKB model, including what should be an eye-popping paper by Schmidt and Turvey (1995) showing that ten interesting predictions have been made

by HKB, and confirmed in experiments on interlimb coordination (Chemero reviews other applications beyond limbs as well on pp. 86–93).

In summary, radical embodied cognitive science has a rich and still-growing theoretical tradition, and some powerful tools to generate predictions and discoveries in cognitive science. After reading through these sections of Chemero (Chapters 5–7) one easily comes away with this feeling of enticement towards this radical theoretical approach. Yet, when stepping back into my own lab where my graduate students study, among other things, attention, sentence processing, social interaction, and deception, I realized that these guides to discovery are a double-edged sword. On the one edge they do suggest new avenues of exploration for radical embodied cognitive science. As just described, the HKB model makes predictions about the dynamics of two stable patterns of behavior and how these change as the task constraints change. But on the other edge, the guides to discovery also serve as explanatory *constraints*. They are a framework for exploration, and any framework has boundary conditions that must be satisfied in order for empirical phenomena to fit inside them (Chemero is aware of this, given his pluralistic tendencies). The result of this is that, looking back on those chapters, Chemero's emphasis is on a range of phenomena that fit best inside the radical embodied boundaries, and extending these boundaries to embrace higher-level cognitive phenomena, while having much promise, is still promissory. This second edge is the first problem of three facing radical embodied cognitive science, as it is instantiated in the book and more generally. I consider a detailed description of this boundary problem first: the scope of radical embodied cognitive science still seems limited, even despite Chemero's heroic efforts in developing and expanding them.

### Problem 1: Limited Epistemological Scope

Chemero does not advance Hegelian arguments (pp. 3–10), and so cannot defend what he calls a “metaphysical claim” (p. 67) for radical embodied cognitive science: averring that cognitive systems are as a matter of fact ones that never have representations. Instead, he recommends the “epistemological claim”: that the best explanations for the functioning of cognitive systems will come from accounts that do not include representations. This claim requires repeated empirical demonstrations of the value of non-representational accounts of kinds of cognition, and I have described some of his case examples above (e.g., analogical reasoning).

Yet, when we survey the case studies of actual research in the book, a basic problem with the epistemological claim becomes fairly plain: there is little high-level cognition to support any significant scope. I exemplify this here with some simple and superficial quantitative analysis of the book. Consider Table 1. The first, left column is major “labels” for cognitive processes found in textbook topics. These include what are commonly termed “low-level” processes (e.g., perception, action, and attention) up to “higher-level” cognition (e.g., problem-solving and language). The next two columns reflect the appearance of these labels in the 2009 *Annual Meeting of the Cognitive Science Society*. The first is a score of the proportion of the conference abstracts in which the terms appear. Consider “perception”: it appears in 14% of the abstracts. After doing this for each term, I then computed a “relative score” for a term's occurrences in the conference abstracts by calculating relative percentage. This is fairly evenly distributed, with the highest score coming from “memory” (about 22% of the total summed proportions are attributed to “memory”).

As a reviewer of this book and an acquaintance of its author, I was generously supplied with a pre-publication electronic version to provide comments, which I shamelessly subjected to the same analysis, this time using “paragraphs” as the occurrence unit instead of abstracts (Dear publisher: this version has been duly deleted). In the next two columns, Chemero’s proportions and percentage scores are shown for the same terms. This (loosely) lays out the various cognitive processes that Chemero uses to demonstrate the power of non-representational accounts of cognition. There are a few examples of high-level cognition (e.g., “problem-solving”), but the vast majority of them stay within the realm of the regular bread and butter of the radical embodied theoretical agenda: basic perception and action. On the surface (and granted, this is just the surface), the epistemological claim has a limited number of case studies from which to make inferences to the best explanation. This is a significant problem of scope, if what is desired is to extend radical embodied cognitive science beyond perception and action and into the realms of problem solving, decision-making, language processing, and so on. The rightmost column shows the proportion of Chemero’s percentages to those of the conference abstracts. A clear emphasis on perception and action is still present.

This problem is clearer when we carefully consider a human’s actual *cognitive ecology*. The space of information humans deal with is vastly linguistic, and conventionalized, such as the spoken and written words which inundate us daily. This is exemplified by another well-known theorist’s work that served to influence the development of modern ecological thinking. Roger Barker’s famous *One Boy’s Day* demonstrates this (Barker and Wright, 1951), in which Barker and his colleagues detail the goings-on of a seven-year-old boy in Oskaloosa, Kansas (referred to as *Midwest* in initial publications). Radical embodied theories may be seamlessly developed for low-level patterns of engagement with the environment, such as the boy’s various moments of navigation and ambulation taking place as he moves about *Midwest* and during school recess. Yet, the boy experiences vast numbers of events of a “high-level” sort, from writing in a journal, chatting with adults and other children, listening to class instruction, and so on. Whatever landscape of affordances underlies these linguistic behaviors, it is

Table 1

Common terms for cognitive processes, and their appearance in the texts. Right column represents Chemero’s divided by CogSci’s score, and in bold are overrepresentations in the book compared to the CogSci proceedings abstracts. See text for details.

Term	CogSci 2009		Chemero, 2009		Score
Perception	0.14	10.2	0.079	37.4	<b>3.7</b>
Attention	0.19	13.9	0.01	4.7	0.3
Action	0.15	11.0	0.07	33.1	<b>3.0</b>
Recognition	0.11	8.0	0.0052	2.5	0.3
Neur- (on/al)	0.1	7.3	0.028	13.2	<b>1.8</b>
Memory	0.3	21.9	0.0074	3.5	0.2
Linguistic	0.14	10.2	0.003	1.4	0.1
Category	0.19	13.9	0.0067	3.2	0.2
Problem-solving	0.049	3.6	0.0022	1.0	0.3

likely hierarchical and multi-scale (Barker and Wright, 1955; discussed in Heft, 2001), such that a theory of information specifying various aspects of the day in the manner detailed in Chemero's upgraded theory of direct perception, information, and affordances (Chapters 6 and 7) does not yet seem prepared to integrate more fully.

Admittedly, Chemero's book does not have the goal of *completely* integrating high-level cognitive concerns with radical embodied cognitive science. Yet readers might, as I did, expect some thorough consideration of levels of analysis that goes beyond the standard grist for the dynamical mill. These levels include the terms in the bottom of Table 1: language production and comprehension, semantic categorization, problem solving and reasoning, and so on. The table does indicate that Chemero does *not* ignore these levels. In fact, he reviews some intriguing work that does accomplish first steps towards this integration of high-level cognition. For example, he discusses the processes of social perception and how they may figure into direct perception and affordances (e.g., Figure 6.3, p. 121). He also considers analogical reasoning of a very simple sort (pp. 131–133). (Though, as I discussed above, this does not seem to escape the use of mental gymnastics in a computational model.) The Chemero columns also indicate a discussion of brain matters, a topic sometimes ignored by theorists in the radical embodied tradition (see pp. 174–181). On this point, he demonstrates that any attention paid to brain regions and their functions must integrate a study of behavioral context (this discussion is particularly compelling and interesting). He also describes (pp. 93–95) a recent paper by Stephen, Dixon, and Isenhower (2009) on a problem-solving task. I agree with Chemero that this is a fascinating and important new course of study that Damian Stephen, J. Dixon, and their colleagues at the University of Connecticut are pursuing. It is relevant to another problem with radical embodied cognitive science, so I consider it next.

### Problem 2: Kinds of Explanation

Stephen, Dixon, and Isenhower (2009) have shown that when participants in a study come upon a new solution to the gear problem, the movement of their hands transitions into new patterns by traversing what may be termed a “phase transition”: a moment of instability, in which there is higher disorder in the movement while participants transition from one stable strategy to another. The gear problem is a simple visual puzzle, likely familiar to readers, in which participants follow a series of gears and decide in which direction the last gear is going to rotate. In their study, the participants had their hand tracked while trying to solve the problem. At some point, (many) participants discover an alternation rule that can be used to solve the problem.

They found that at about this point of discovery, the dynamic structure of hand movement has higher entropy. This indicates that the new stable problem-solving strategy, what Stephen et al. call a “new representation,” engenders a dynamic property of bodily movement that is predicted from a perspective of self-organizing systems. If the human cognitive system is a self-organizing system, then finding a new solution to a problem will involve “decomposition” of one stable strategy and reorganization into a new task set or strategy to solve a task. Stephen et al. (2009) see this in the arm movement data collected while participants have an “a-ha” moment in this task.

Chemero spends a few pages describing this fascinating finding (pp. 93–95) and I agree with him that this experiment is groundbreaking in its implications. However, we differ to some extent regarding what the study contributes. To Chemero, “[t]his work is a nonrepresentational dynamical explanation of a manifestly cognitive, representation-hungry activity” (p. 95). But from another more tempered perspective, the



study tells us that (1) what we commonly term new “strategies” or “rules” for problem-solving and high-level cognition may be two stable modes of a self-organizing system and (2) the distinct “entropic” or disordered transition between stable modes is the “dynamic signature” in how transitions take place. Transitions are not “flip-flopped” bits in a computational system, or similar classical concepts. Instead, stable properties of behavior are adaptive, self-organizing wholes with their own special properties, but a system must transition through a distinct phases of critical instability as it “carries out” this transitioning.

This is fascinating because it tells us about *the structure of change* in the human cognitive system. It offers a general and elegant window into the manner in which a cognitive system undergoes transitions between stable modes. What it does not do, however, is tell us about *any specific strategy or mode* itself. In other words, it is a generalization about self-organizing cognitive systems, not about the particular problems under study. Given some problem P, and two general strategies used to solve that problem, S1 and S2, Stephen et al. show that the system’s behavior as it transitions from S1 to S2 has a particular character (i.e., more “entropic” fluctuation). But what are S1 and S2? What is P? Does P belong in some subclass of problem for which this is relevant, at the exclusion of others? Chemero and Stephen et al. (2009) intend this basic finding to be a generalization across all P’s and SX’s. But in cognitive science, there are *fundamental problems of context* that must be considered to account for any particular P.

To clarify this point, let’s consider a quick example. Of course, there are thousands upon thousands of “problems” humans face (though they can fall into classes: Eysenck and Keane, 2005). In all, it is important to consider the “content” of any individual problem in order to understand it. Consider one problem with direct real-world applied relevance, P1:

P1: Learning to add any two numbers.

This is a problem faced daily by children seeking mastery, but its solution is not likely a single strategy, but a set of stable “processes” that children use, the most prominent of them may be “carry the 1,” but of course the others are the basic arithmetic operations between digits 1–9. Suppose we find that children show entropic transitions in certain moments of realization when they have finally stably adopted whatever instructions the teacher is attempting to instill. This still does not *explain* the problem space itself. Pedagogical problems invoke domain concerns, from mathematics to history to science. This is well known in psychological studies on pedagogy (Shulman, 1986). In general, the idea that the specific *content* of a problem is key to explaining the discovery of solutions to it has formed entire literatures in cognitive science, from smaller insight problems to domain expertise (see Sternberg and Ben-Zeev, 2001 and Eysenck and Keane, 2005, for reviews).

In summary for this section, what I have argued is that the approaches of radical embodied cognitive science, identifying the self-organizing properties of change occurring at the juncture of body and environment, are fascinating and contribute something quite substantial to the field. However, they do not gainsay, nor can even make contact with, the particular content that forms the basis of the high-level cognitive concerns themselves. Take the original example of the gear problem. Knowing there was a transition does not tell us “What did it transition *into*?” The answer is “they learned to alternate”: a (perhaps briefly) fascinating discovery to a participant that has “*content*” and not just a dynamic structure. In short, there are two kinds of explanation here. One is about how cognition self-organizes around problem spaces;

the other is what about the problem spaces is relevant, important, etc. to any kind of solution at all. Both are important (e.g., see Bechtel, 1997). Radical embodied cognitive science does wonders with the former. But other approaches do well with the latter. This naturally extends to another problem, which I consider next.

### Problem 3: Hacking's Entity Realism for Mental Gymnastics

Chemero spends a considerable amount of space discussing theories of affordances (indeed a whole chapter, Chapter 7). As the reader of this review now knows, this is an important topic in the realm of radical embodied cognitive science. It can and does bring about many important philosophical questions. One of these is the standard ontological concern that arises when theoretical constructs are offered in scientific theories: Do they "exist"? This question has produced quite a bit of debate in 20th century philosophy of science, and it continues unabated today (e.g., Clarke, 2001). In the case of "affordance," Chemero wonders whether radical embodied cognitive scientists can take a "realist's" perspective about them, which he defines partly as the thought that "objects of our accurate perceptions/thoughts/theories exist in an animal-independent world" (p. 185). (Chemero discusses both realism regarding perceptions, and realism regarding theoretical entities. Here I refer to the latter.)

The philosophical framework for existence that Chemero chooses is Hacking's entity realism (though he acknowledges others may work, pp. 192–193). Entity realism is pretty intuitive, as stated by Hacking himself: "By the time that we can use the electron to manipulate other parts of nature in a systematic way, the electron has ceased to be something hypothetical, something inferred. It has ceased to be theoretical and has become experimental." (Hacking, 1983, p. 283, cited in Chemero p. 192).

Chemero describes three research strands showing that affordances can be seen as kinds of "experimental equipment" (p. 195). This makes them unobservable "tools," and can therefore be endowed with the kind of reality that other scientific concepts have. I agree with Chemero on this note. Even if affordances are relational, and thus unobservable, relations can be "real" in some sense, especially if you like Hacking's entity realism (though, I'm more inclined to agree with Putnam's [2004] recent assessment of ontological matters, which says, in short, ontology doesn't really matter *that much*).

What is curious, though, is that Chemero separates his chapter on mental gymnastics (representations, Chapter 3) from this chapter on the metaphysics of radical embodiment (Chapter 9). To some extent this makes sense, because Chemero's goal is to lay out radical embodied cognitive science, and its own ideas about organisms, environments, and their relations. Yet Hacking's entity realism is clearly also relevant to other cognitive concerns. Are "representations" concepts that are actively, experimentally manipulated? Some would say they clearly are. From the use of word lists to infuse participants with new knowledge to test the "encoding of memories," all the way up to reading research that has participants develop new ideas and expectations about concepts in a text, the concept of "representation" is rich and manipulable as an abstract notion in these high-level cognitive experiments. They are, at the very least, as manipulable as the affordance of Chemero's studies (see pp. 195–197). But if they are not, Chemero does not offer a discussion of this.

In any case, the claim I am making here is that the philosophical support Chemero seeks out for radical embodied cognitive science *may* work just as well when it is applied to domains of cognitive science which radical embodied cognitive science can't explain: the contents of high-level cognitive problem spaces. If this is the case,

and experimenters who work on high-level cognitive science can use as tools abstract concepts that are unobservable, then one wonders whether radical embodied cognitive science is simply too radical. Does it limit itself by its negative thesis? Does it limit itself by its chosen tools? Radical embodied cognitive science and other traditional theoretical paradigms in cognitive science can sometimes have an air of unnecessary purism. Its theoretical practices seem almost like deriving pleasure from washing a messy bathroom with a toothbrush: one insists on “doing it right.” When for different surfaces, different corners, different stains, one can seek out various compounds and tools to achieve different purposes efficiently and effectively. So cognitive science’s occasional theoretical purism is metaphorically something that in military contexts would be more akin to a form of punishment. In the next section, I take seriously the idea that the human cognitive system is complex, concluding that so too should be our cognitive science. In short: sometimes, you’ve gotta put down the toothbrush.

### Conclusion: Taking Complex Systems Seriously

I have summarized the key features of Chemero’s wonderful book that I feel make a compelling case for radical embodied cognitive science. His clever idea to develop “guides to discovery” involves further developing and broadly applying extensions to Gibson’s ecological psychological framework, and further applying elegant and powerful dynamical methods, such as HKB. I did identify, however, three core problems for this approach. In this final section, I would like to briefly propose a solution to the three problems that still embraces important aspects of radical embodied cognitive science.

I’m fairly sure Chemero would accept at least some version of what I’m proposing (and have evidence for it: Dale, Dietrich, and Chemero, 2009). The solution begins by embracing one of the key features of the radical embodied approach: the primacy of context in our understanding of how cognition works. Consider again Barker’s *One Boy’s Day* (1951) and his ecobehavioral approach to psychology. Central to the boy’s daily cognitive functioning is his *materials for cognition*: the sea of language, spoken and written, the objects to be described and/or manipulated, the various observable movements of the boy’s conspecifics, and so on. The boy’s behavior and environment are probably, in a paraphrase of Chemero’s position, nonlinearly coupled together in a crucially constitutive way (pp. 31–32). The properties of individual cognitive agents are therefore functions of agent–environment interactions. This much the radical embodied cognitive scientist would accept with a grin and a nod.

But this nonlinear coupling constitution leads too readily to a rejection of talk of computation and representation. It turns out that the boy’s cognitive ecology is replete with computation and representation stuff. For example, you, the reader, have eye movements at this very moment that are hop-scotching across space-delimited words that are ordered according to fairly regular rules. A great debate in cognitive science is about whether the rules “in the head” to process these words are so regularly inscribed (e.g., Christiansen and Chater, 1999); even deeper than this, the radical embodied theorist would ask whether it is helpful to say there are rules there at all! As I have argued, the guides to discovery do not yet have much material to help with this *particular* linguistic domain (see above, section on limited epistemological scope), so we are left with a trilemma: (1) Does the radical embodied theorist *ignore* these complex domains of cognition while at the same time rejecting as hokey the computational–representational theories being developed for them (a common move)? (2) Does she throw her hands up in awe at the complexity of human mental matters and

accept that, in some complex domains, human *do* have representations in the traditional sense of Miller, Chomsky, and Fodor (this move almost never happens)? (3) Or, perhaps the radical embodied theorist can speak of the human cognitive system *behaving in such a manner as to employ latent computational and representational capacities* (there exist some versions of this idea, e.g., Clark, 1997; Dennett, 1998, among others)?

I would argue (only briefly with space here) that (3) is the most sensible move, but requires a bit more elaboration to clarify why it is not trivial. You could call this elaboration the matter of *taking complex systems seriously*. Complex systems are ones with many interacting units whose joint behavior generates some collective properties of the overall system that the individual units do not possess (e.g., Anderson, 1972). What tends to occur in very large complex systems is a hierarchical nesting of spatial or temporal scales, each layer being composed of units at a new scale that may also become interdependent and produce further unique properties. Indeed, exciting research in the tradition of radical embodied cognitive science has shown this multi-scale interdependency in human experimental work (Van Orden, Holden, and Turvey, 2003). One may argue (see Dale, 2008) that the complexity of the human cognitive system, nonlinearly embedded in its ecological context, has evolved to produce the so-called *emergence* of new, diverse cognitive characteristics. And an obvious emergent property of the human cognitive system is computational–representational forms of behavior.

Does this mean “representations exist”? When we embrace the contextual character of radical embodied cognitive science, I would argue that it indeed does: representations are theoretical entities associated with specific human–environment transactions. They are experimentally manipulable in the way Hacking describes for abstract theoretical constructs. In addition, representations are partly decouplable from the transactions in which they are born (for a discussion of this, see Chemero’s Chapter 3), but even in such cases they reflect *capacities* for forms of behavior (in the sense of Cartwright, 1999). These capacities may be, in a certain simplified way, attributed *constitutively* to the stuff inside the cranium, though all representational hypotheses may not turn out to be “internal control structures.” (It is important to note that there is a vast literature on the particular roles of neural control structures that should be considered in any serious discussion of this issue. I finesse this here.)

Through such a strategy, the radical embodied cognitive scientist can speak of the emergent capacities for computational–representational behavior (some endeavors in basic physics have no problem with this: Crutchfield, 1998; Lloyd, 2006). She can loosen her necktie somewhat and get her hands dirty with the *contents* of the particular contextual capacities humans exhibit. In this way, by bridging levels of explanation, the fascinating discoveries of Stephen et al. (2009) may shed *direct* light on these contents and how they change. Through the coupling of humans to their ecological cognitive context, this complex system produces emergent forms of behavior. This perspective does not suffer from any of the three problems described above. In addition, this perspective has an appropriate level of agnosticism regarding mental gymnastics that radical embodied folks would probably like, and that mental gymnastic theorists ought to more commonly apply.

## References

- Anderson, P.W. (1972). More is different. *Science*, 177(4047), 393–396.  
 Barker, R.G., and Wright, H.F. (1951). *One boy’s day: A specimen record of behavior*. New York: Harper & Row.

- Barker, R.G., and Wright, H.F. (1955). *Midwest and its children: The psychological ecology of an American town*. New York: Harper & Row.
- Barwise, J., and Perry, J. (1981). Situations and attitudes. *The Journal of Philosophy*, 78(11), 668–691.
- Bechtel, W. (1997). Representations and cognitive explanations: Assessing the dynamicist's challenge in cognitive science. *Cognitive Science*, 22(3), 295–318.
- Cartwright, N. (1999). *The dappled world*. Cambridge: Cambridge University Press.
- Christiansen, M.H., and Chater, N. (1999). Toward a connectionist model of recursion in human linguistic performance. *Cognitive Science*, 23(2), 157–205.
- Clark, A. (1997). *Being there: Putting brain, body, and world together again*. Cambridge, Massachusetts: MIT Press.
- Clarke, S. (2001). Defensible territory for entity realism. *The British Journal for the Philosophy of Science*, 52(4), 701–722.
- Crutchfield, J.P. (1998). Dynamical embodiments of computation in cognitive processes. *Behavioral and Brain Sciences*, 21(5), 635.
- Dale, R. (2008). Pluralism and the future of cognitive science [special issue]. *Journal of Experimental & Theoretical Artificial Intelligence*, 20(3), 153–267.
- Dale, R., Dietrich, E., and Chemero, A. (2009). Explanatory pluralism in cognitive science. *Cognitive Science*, 33, 739–742.
- Dennett, D.C. (1998). *The intentional stance*. Cambridge, Massachusetts: MIT Press.
- Eysenck, M.W., and Keane, M.T. (2005). *Cognitive psychology: A student's handbook*. New York: Psychology Press.
- Gibson, J.J. (1979). *The ecological approach to visual perception*. Boston: Houghton Mifflin.
- Hacking, I. (1983). *Representing and intervening: Introductory topics in the philosophy of natural science*. Cambridge: Cambridge University Press.
- Haken, H., Kelso, J.A.S., and Bunz, H. (1985). A theoretical model of phase transitions in human hand movements. *Biological Cybernetics*, 51(5), 347–356.
- Hefl, H. (2001). *Ecological psychology in context: James Gibson, Roger Barker, and the legacy of William James's radical empiricism*. Mahwah, New Jersey: Lawrence Erlbaum.
- Kelso, J.A. (1995). *Dynamic patterns: The self-organization of brain and behavior*. Cambridge, Massachusetts: MIT Press.
- Lloyd, S. (2006). *Programming the universe*. New York: Knopf.
- Michaels, C.F., and Carello, C. (1981). *Direct perception*. Englewood Cliffs, New Jersey: Prentice-Hall.
- Minsky, M., and Papert, S. (1969). *Perceptions*. Cambridge, Massachusetts: MIT Press.
- Noë, A. (2005). *Action in perception*. Cambridge, Massachusetts: MIT Press.
- Putnam, H. (2004). *Ethics without ontology*. Cambridge, Massachusetts: Harvard University Press.
- Schmidt, R.C., and Turvey, M.T. (1995). Models of interlimb coordination — equilibria, local analyses, and spectral patterning: Comment on Fuchs and Kelso (1994). *Journal of Experimental Psychology: Human Perception and Performance*, 21, 432–432.
- Shaw, R.E., Turvey, M.T., and Mace, W. (1982). Ecological psychology: The consequence of a commitment to realism. In W. Weimer and D. Palermo (Eds.), *Cognition and the symbolic processes* (Volume 2, pp. 159–226). Hillsdale, New Jersey: Erlbaum.
- Shulman, L.S. (1986). Those who understand: Knowledge growth in teaching. *Educational Researcher*, 15(2), 4–14.
- Spivey, M.J., and Dale, R. (2004). On the continuity of mind: Toward a dynamical account of cognition. In B.H. Ross (Ed.), *The psychology of learning and motivation: Advances in research and theory*, Volume 45 (pp. 87–142). San Diego: Elsevier Academic Press.
- Stephen, D.G., Dixon, J.A., and Isenhower, R.W. (2009). Dynamics of representational change: Entropy, action, and cognition. *Journal of Experimental Psychology: Human Perception and Performance*, 35(6), 1811–1832.
- Sternberg, R.J., and Ben-Zeev, T. (2001). *Complex cognition: The psychology of human thought*. Oxford: Oxford University Press.
- Turvey, M.T. (1977). Preliminaries to a theory of action with reference to vision. In R. Shaw and J. Bransford (Eds.), *Perceiving, acting and knowing* (pp. 211–265). Hillsdale, New Jersey: Erlbaum.
- Turvey, M.T. (1992). Affordances and prospective control: An outline of the ontology. *Ecological Psychology*, 4(3), 173–187.
- Turvey, M.T., and Carello, C. (1996). Dynamic touch. *American Psychologist*, 51(11), 1134–1152.

- Van Orden, G.C., Holden, J.G., and Turvey, M.T. (2003). Self-organization of cognitive performance. *Journal of Experimental Psychology: General*, 132(3), 331–350.
- Warren, W.H. (1984). Perceiving affordances: Visual guidance of stair climbing. *Journal of Experimental Psychology*, 105(5), 683–703.