Dynamical Models and Mechanistic Explanations

Carlos A. Zednik (czednik@indiana.edu)
Cognitive Science Program, Indiana University
Bloomington, IN 47406 USA

Abstract

Philosophers of science increasingly believe that cognitive phenomena are explained by describing the mechanisms that produce those phenomena, rather than by subsuming them under laws. With respect to the increasingly influential dynamical approach to cognition, however, it has been claimed that only certain kinds of dynamical models (in particular, connectionist dynamical models) satisfy the requirements for mechanistic explanation. This article evaluates this claim. In particular, here it is argued that in the field of evolutionary robotics, several non-connectionist dynamical models have been developed that provide mechanistic explanations of "minimally cognitive" phenomena. Accordingly, the dynamical approach to cognition is conceived as a scientific approach unified by the common employment of a certain body of mathematical tools and concepts, rather than by a common explanatory method.

Keywords: explanation; mechanism; dynamical systems theory; connectionism; evolutionary robotics.

Introduction

The increasingly influential dynamical approach to cognition (Port & van Gelder, 1995) is to a large extent in the business of developing dynamical models—models that employ the tools and concepts of dynamical systems theory—to explain natural and artificial cognitive phenomena. Although much has been said about the way dynamical modelers *think of* cognitive phenomena—as temporally continuous, non-representational, embodied, and embedded (Beer, 1995; Clark, 1997; van Gelder, 1995)—relatively little has been said about the way they go about *explaining* such phenomena. What is the nature of dynamical explanation?

In an influential paper, William Bechtel (1998) compares two prominent philosophical accounts of scientific explanation: *mechanistic explanation* and *covering law explanation*. With respect to the dynamical approach to cognition, Bechtel suggests that this distinction can be mapped onto another distinction, due to van Gelder & Port (1995), between *connectionist* and *non-connectionist* dynamical models. The suggestion is that, whereas the former are typically involved in mechanistic explanations of cognitive phenomena, the latter more closely resemble covering law explanations.

In what follows, I evaluate Bechtel's claim. In particular, I suggest that some non-connectionist dynamical models—those which resemble Randy Beer's (2003) model of a simulated "minimally cognitive" agent—are in fact involved in mechanistic explanations. Accordingly, these models serve as counterexamples to Bechtel's claim: the distinction between connectionist and non-connectionist dynamical modeling cannot be straightforwardly mapped onto the distinction between mechanistic and covering law explanation. I conclude with a brief discussion of the nature of the dynamical approach, arguing that it should be viewed as a scientific approach unified by the common employment of a particular body of mathematical tools and concepts, rather than by a common explanatory method.

Mechanistic and covering law explanation

In order to determine how the dynamical approach goes about explaining cognitive phenomena, it is important to consider the nature of scientific explanation in general. Two accounts of scientific explanation are particularly relevant: the mechanistic account and the covering law account.[1]

What is mechanistic explanation? Although no "standard" account exists at present, at least two core ideas can be identified. The first is the idea that mechanistic explanations distinguish between the phenomenon being explained and the particular system that underlies that phenomenon (Bechtel & Abrahamsen, 2006; Glennan, 2005). The goal of mechanistic explanation is to "show how" a phenomenon arises from such a system by describing the *mechanism* which (a) is instantiated by the system, and (b) is responsible for producing the phenomenon. What is a mechanism? Typically, mechanisms are defined as *organized structures of parts and operations* (Bechtel & Richardson, 1993; Bechtel & Abrahamsen, 2005; Machamer, Darden, & Craver, 2000). Accordingly, a phenomenon will have been mechanistically explained if an organized structure of parts and operations can be described in a way that "shows how" that structure is instantiated by a particular system and is responsible for producing the phenomenon of interest.

The second idea common to most accounts of mechanistic explanation concerns the way such mechanisms are typically identified and described: via a *decompositional* strategy (Bechtel & Richardson, 1993; Cummins, 1983). On the one hand, the phenomenon is functionally decomposed into several component operations, each simpler than the phenomenon itself. On the other hand, the underlying system is structurally decomposed into several interacting parts. Provided that each operation can be associated with ("localized in") one or more component parts, such decomposition will "show how" the system produces the phenomenon. Accordingly, decomposing a phenomenon and an underlying system

[1] Note that I do not mean to suggest that these two accounts together span the space of (actual or possible) scientific explanations. Rather, I am merely contributing to a debate which has assumed that most explanations offered under the banner of the dynamical approach can be associated with either one of these two accounts. I also do not mean to suggest that mechanistic and covering law explanations necessarily stand in *competition* with one another. Indeed (as I briefly discuss in the conclusion), I suspect that much can be gained from *combining* the explanatory insights achieved from mechanistic and covering law explanations of the very same cognitive phenomena.

allows us to mechanistically explain that phenomenon as arising from that particular system.

In contrast, covering law explanations depend neither on distinguishing between the phenomenon being explained and the underlying, mechanism-instantiating system, nor on employing a decompositional strategy. Rather, covering law explanations subsume phenomena under *natural laws*: a phenomenon is explained when a description of that phenomenon can be derived from a specification of one or more natural laws and relevant initial conditions (Hempel, 1965). Although it remains controversial what exactly constitutes a natural law, it is generally agreed that laws are specified as mathematical or statistical regularities which allow us to reason counterfactually about the phenomena being explained. Notably, such regularities usually apply to phenomena directly—it is not necessary to refer explicitly to the underlying systems from which those phenomena arise. Finally, unlike mechanistic explanations, which are involved with describing *particular* systems, covering law explanations strive to be wholly *general*: they apply to any phenomenon which can be subsumed under the posited natural laws.

What kind of explanations—mechanistic or covering law—are provided in the dynamical approach to cognition? In his influential article "Representations and Cognitive Explanations: Assessing the Dynamicist Challenge in Cognitive Science", William Bechtel follows van Gelder and Port (1995) in distinguishing *connectionist* dynamical models from *non-connectionist* dynamical models. Bechtel's principal claim in that article is that, although the former are frequently involved in mechanistic explanations of cognitive phenomena, the latter are more appropriately associated with the covering law account of scientific explanation. Is this claim accurate? In the next two sections, I will elaborate on the differences between connectionist and non-connectionist dynamical modeling, and will discuss the relation of each of these to the two accounts of scientific explanation just outlined. After that, I will challenge Bechtel's claim by discussing a particular non-connectionist dynamical model which after all seems to provide a mechanistic explanation of a "minimally cognitive" phenomenon.

Connectionist dynamical modeling

A connectionist dynamical model consists of two things: a particular connectionist model that reproduces some cognitive phenomenon, and a dynamical analysis (i.e. an analysis that employs the conceptual tools of dynamical systems theory) of that model's behavior. Although there exist no commonly accepted necessary and sufficient conditions for distinguishing connectionist models from other kinds of models, certain characteristics can be considered prototypical. First, the formal structure of connectionist models is typically *high-dimensional* and *homogeneous*: they are specified as large networks of interconnected "neural" units, some of which "may be designated as inhibitory and others as excitatory, but beyond that, are rarely differentiated" (Rumelhart & McClelland, 1986: 137). Second, connectionist models usually describe *neural processes*: they represent the "internal" workings of a cognitive system, receiving input from the "external" environment (or from other parts of the larger system), and producing output usually interpreted as the cognitive system's overt behavior. Third, the behavior exhibited by connectionist models is typically understood by asking (and answering) a particular set of questions about their information-processing capacities: What, if anything, do the hidden unit activations represent? How do the model's connection weights encode information about the training data? Which input-output function is "learned" by the network? Although it is important to remember that none of these prototypical characteristics—the first syntactic, the second semantic, and the third pragmatic—constitute necessary or sufficient conditions, I take them to be the features most commonly exhibited by models developed under the general connectionist banner.

A certain subset of connectionist models—*recurrent neural network models* (Elman, 1990)—have been described as "welcome participants in the dynamical approach" (van Gelder & Port, 1995: 34). Unlike traditional feed-forward networks, recurrent neural networks specify feedback loops that allow them to "remember" their own activations from previous moments in time. That is, the behavior of such networks is determined not only by momentary inputs, but also by internal states. This feature is particularly useful for modeling some of the aspects of cognitive phenomena most emphasized by proponents of the dynamical approach: temporal continuity, sensitivity to temporal detail, and dependencies between later states and current states (van Gelder & Port, 1995). Moreover, such networks call for a particular conceptual toolkit—the analytic toolkit of dynamical systems theory—with which to understand the way they work. In particular, the questions concerning hidden-layer representations and information-processing capacities are posed and answered in terms of state space trajectories, limit sets, and sensitivity to continuously changing system parameters. In short, connectionist dynamical modeling involves a particular kind of connectionist model in addition to involving a uniquely dynamical language of description.

A particularly prominent example of connectionist dynamical modeling is Jeffrey Elman's well-known model of word-prediction (Elman, 1995). This model specifies a simple recurrent network which, via a standard backpropagation algorithm, is trained to predict the next word given the first *n* words of a sentence. Notably, the network's performance depends on a feedback loop which "remembers" the linguistic context in which individual words appear. This context contains revealing information about the grammatical and semantic role played by individual words in a sentence, without which robust word-prediction would be impossible. In virtue of taking into account this context by storing its hidden-unit activation from previous moments in time, the network is able to predict the grammatical role as well as the semantic category of subsequent words.

For current purposes, it is important to consider how Elman goes about *explaining* the behavior of the network. In order to understand exactly how the network solves the word-prediction problem, Elman first performs a functional decomposition of the task. Among others, he identifies the operations of correctly encoding verb argument structure (different verbs require different kinds of objects), and of ensuring noun/verb agreement within clauses (plural nouns require appropriately modified verbs). He then performs a principal component analysis of the network's hidden-unit activation to demonstrate that each of these linguistic rule-following operations is in fact instantiated by the system. To this end, the operations are visualized as trajectories in a principle-component state space, different regions of which correspond to different grammatical and semantic categories:

> "The network has learned to represent differences in lexical items as different regions in the hidden unit state space. The sequential dependencies which exist among words in sentences are captured by the movement over time through this space as the network processes successive words in the sentence." (Elman, 1995: 215)

Although the principal component analysis does not identify the underlying network's component parts (because the hidden-unit representations are distributed rather than local), it demonstrates that each of the linguistic operations is in fact instantiated by the system. That is, he "shows how" the network instantiates a mechanism which is responsible for accurate word classification and prediction. Therefore, insofar as Elman's model exemplifies the class of connectionist dynamical models, this class can be associated with the mechanistic account of scientific explanation.

Non-connectionist dynamical modeling

Non-connectionist dynamical modeling differs from connectionist dynamical modeling along each of the syntactic, semantic, and pragmatic dimensions discussed above. First, the formal specification of non-connectionist dynamical models is typically *low-dimensional* and *heterogeneous*: non-connectionist dynamical models usually consist of a small and structurally distinct set of equations defined over a small number of state variables. Second, these models typically capture the features of embodied and embedded cognitive systems: the state variables, parameters, and equations of a single non-connectionist dynamical model can be used to capture global or local features of the cognitive agent's nervous system, its body, and even objects in the environment. Third, the pragmatic role played by most non-connectionist dynamical models is quite different from the role played by their connectionist kin. Rather than asking questions about hidden-unit representations, non-connectionist dynamical modelers tend to focus on determining the precise temporal relationships between individual features of a cognitive phenomenon, on describing the possible behavioral trajectories a cognitive system might traverse, and on specifying the way the shape of those trajectories depends on changes in one or more system parameters.[2]

A particularly well-known non-connectionist dynamical model is the Haken-Kelso-Bunz (HKB) model of spontaneous bimanual coordination (Kelso, 1995). Given the task of oscillating left and right index fingers to the rate of an oscillating metronome, subjects invariably settle on one of two between-fingers phase relations: in-phase or antiphase. Near a certain frequency however, subjects reliably settle on the in-phase relation only. To explain this spontaneous phase-switching behavior, the HKB model specifies a single differential equation:

$$\dot{\phi} = -asin\phi - 2bsin\phi$$

where a and b are system parameters which capture the oscillation rates of each finger, and ϕ is a "collective" variable denoting the between-fingers phase relation.

On the basis of this non-connectionist dynamical model, qualitative and quantitative predictions can be made about how and when spontaneous phase-switches occur. Moreover, the HKB model supports counterfactual reasoning in the form of determining what happens if one or both fingers are briefly forced out of their regular oscillatory motion. For these reasons, it has been suggested that the HKB model does in fact provide an explanation of *some* kind (Chemero, 2000; Clark, 1997). Indeed, what the model's differential equation amounts to is a law-like regularity that describes the way the relative phase between index fingers changes over time, given a particular set of initial conditions (i.e. the initial phase and oscillation rate)—it fits the characterization of covering law explanation given above. Moreover, the equation describes a property of the phenomenon itself, without referring to the underlying system of muscle tissue and neural processing which underlies that phenomenon. Finally, the authors also hypothesize that the mathematical regularity they describe is a special case of a more general regularity that applies to all relevantly similar phenomena of self-organized coordination, including gait transitions in two and four-legged animals, speech production, and even certain coordination behaviors between organisms (Kelso, 1995: 69-95): the HKB model can be seen as subsuming a particular cognitive phenomenon under a *general* law of coordination dynamics. For all these reasons, this non-connectionist dynamical model appears to provide a genuine covering law explanation of the phenomenon of spontaneous bimanual coordination.

[2]Notably, like the "prototypical characteristics" of connectionist dynamical models discussed above, these features common to many non-connectionist dynamical models constitute neither necessary nor sufficient conditions. Indeed, just like connectionist dynamical models might sometimes be relatively heterogeneous and (at least in principle) describe features of brain, body, and the environment, non-connectionist dynamical models might themselves be quite homogeneous, and be used to describe the behavior of purely "internal" processes. Therefore, although syntactic and semantic considerations should not be discounted, the most important differences are likely to be pragmatic—how the models are used, thought of, and understood. In a sense, it is not always particularly helpful or illuminating to think of a model in connectionist terms.

The rest of this discussion will be concerned with a particular claim, originally due to William Bechtel (1998), that identifies all such non-connectionist models with the covering law account of scientific explanation:

> "The difference and differential equations in these models are intended to describe patterns of linked change in the values of specified [variables] in the course of the system's evolution over time. [They] do not correspond to components of the system which interact causally. They are, rather, features in the phenomenon itself...In this respect, these DST explanations better fit the alternative, covering law model of explanation." (Bechtel, 1998: 5)

More generally, Bechtel suggests that the distinction between connectionist and non-connectionist dynamical modeling straightforwardly maps on to the distinction between mechanistic and covering law explanation:

> "the distinction van Gelder and Port drew between two sorts of DST explanations, connectionist and non-connectionist, represents a bigger gulf—that between mechanistic explanations and covering law explanations." (Bechtel, 1998: 5)

To what extent is this suggestion accurate? Over the next two sections, I will discuss an apparent counterexample to Bechtel's suggestion from the field of evolutionary robotics—a non-connectionist dynamical model which provides a mechanistic explanation of a "minimally cognitive" phenomenon.

Non-connectionist dynamical modeling and mechanistic explanation

Recent research in the field of evolutionary robotics has harnessed the tools of dynamical systems theory to develop and study a variety of simulated agents, artificially evolved to perform a range of "minimally cognitive" tasks. One particularly prominent example of this kind of research is presented by Randy Beer (2003). Beer introduces a class of simulated agents that are evolved to "visually" categorize simple objects according to their shape, and to manifest their decisions by catching or avoiding them as they approach. Interestingly, the most successful agent that emerges from the artificial evolutionary process (97.08% accuracy on 10,000 random trials) adopts an unexpected categorization strategy: it *actively scans* the object by moving side to side a number of times before either catching it at its center or avoiding it to the left or to the right. How might this behavior be explained?

The first component of Beer's study is a 16-dimensional dynamical model: a 14-unit continuous-time recurrent neural network (CTRNN) "brain" coupled to a 2-dimensional specification of the agent and its environment:

$$\tau_i \dot{s}_i = -s_i + I_i(x, y; \alpha) \quad i = 1, \ldots, 7$$

$$\tau_i \dot{s}_i = -s_i + \sum_{j=1}^{7} w_{ji}\sigma(g(s_j + \theta)) + \sum_{j=8}^{12} w_{ji}\sigma(s_j + \theta_j) \quad i = 8, \ldots, 12$$

$$\tau_i \dot{s}_i = -s_i + \sum_{j=8}^{12} w_{ji}\sigma(s_j + \theta) \quad i = 13, 14$$

$$\dot{x} = 5(\sigma(s_{13} + \theta_{13}) - \sigma(s_{14} + \theta_{14}))$$

$$\dot{y} = -3$$

where x and y are the horizontal and vertical coordinates of the object relative to the agent; $I(x, y; \alpha)$ is the "visual" input due to an object with shape a at location (x, y) relative to the agent; s is the state of each of the fourteen "neurons" with activation function σ and neural parameters τ, g, and θ; and w_{ji} is the strength of the connection between units i and j. Notably, whereas the first fourteen equations specify the three layers of the agent's CTRNN "brain", the other two equations specify the motion of the agent and of the object, respectively. In other words, the whole model is a model of a distributed 16-dimensional brain-body-environment system.

What kind of dynamical model—connectionist or non-connectionist—is this? The involvement of a 14-unit CTRNN in determining the agent's behavior might suggest that this model exemplifies the practice of connectionist dynamical modeling. This appearance is arguably misleading, however. First, consider the model's syntactic structure. It is true that Beer's model, like most connectionist dynamical models, is relatively high-dimensional. Nevertheless, it incorporates two equations with a formal structure quite unlike the formal structure of the other fourteen equations. Second, unlike most connectionist dynamical models which (implicitly or explicitly) describe neural processes only, Beer's model is designed to capture "internal" as well as "external" features of a single distributed brain-body-environment system. Importantly, this explicit reference to the environment is taken to be a crucial aspect of the model: Beer himself has strongly suggested that the systems which underlie cognitive phenomena must be modeled and understood as crossing the boundaries between brain, body, and environment:

> "an agent's behavior arises not simply from within the agent itself, but rather through its interaction with its environment...Since properties of the coupled system cannot in general be attributed to either subsystem individually, an agent's behavior properly resides only in the dynamics of the coupled [brain-body-environment system] and not simply in the dynamics of [the agent] or [the environment] alone." (Beer, 1995: 132; see also Clark, 1997; van Gelder & Port, 1995)

Insofar as connectionist dynamical models do not (in general) capture the dynamics of coupled brain-body-environment systems, but merely describe "internal" neural dynamics, they differ substantially from Beer's dynamical model.

Going beyond syntactic and semantic considerations, it is important to consider the way Beer proposes to use and understand his model:

> "Rather than assigning representational content to neuronal states, the mathematical tools of dynamical systems theory are used to characterize the structure of the

> space of possible behavioral trajectories and the internal and external forces that shape the particular trajectory that unfolds." (Beer, 2003: 210)

In other words, the questions Beer asks about the model and its behavior are not questions about the CTRNN's information-processing capacities, but questions about the coupled temporal evolution of the sixteen state variables, about the model's dependencies on changing parameter values, and about the way real-time interactions between individual elements of the model together produce the observed behavior. Although it is surely *possible* to think of one part of the model (viz. the fourteen CTRNN equations) in connectionist terms, it should not be assumed this would be the most insightful way of understanding its contribution to the observed behavior. Therefore, not only does Beer's model *look* different from connectionist dynamical models, nor does it merely *stand for* things other than neural properties, but in addition, Beer *thinks of* the model in a way that differs from traditional connectionist thinking. In this sense, Beer's model of a "minimally cognitive" categorization agent and its environment differs from connectionist dynamical models in all previously mentioned syntactic, semantic, and pragmatic respects. Therefore, considering this model to be a genuinely connectionist dynamical model appears to be, if not impossible, at least contrary to the intentions of its author.

Assuming that his model is in fact a non-connectionist dynamical model, it is important to consider how Beer employs the conceptual tools of dynamical analysis to *explain* the behavior it exhibits. Notably, this analysis follows a decompositional strategy first articulated in 1995:

> "I will model an agent and its environment as two continuous-time dynamical systems *A* and *E*, respectively...Because an agent and its environment are in constant interaction, *A* and *E* are coupled nonautonomous dynamical systems. This coupling can be represented with a sensory function *S* from environmental state variables to agent parameters and a motor function *M* from agent state variables to environmental parameters." (Beer, 1995: 130)

With respect to the particular agent in the 2003 study, *S* and *M* correspond to the following operations:

> "we will decompose the agent-environment dynamics into: (1) the effect that the relative positions of the object and the agent have on the agent's motion; (2) the effect that the agent's motion has on the relative positions of the object and the agent." (Beer, 2003: 228)

That is, the behavior of the whole 16-dimensional brain-body-environment system can be understood as the interactive operations between two simpler subsystems: the agent's motion *M* as the operation of subsystem *A* (the agent) on the one hand, and the sensory stimulus function *S* as the operation of subsystem *E* (the environment) on the other. Notably, these operations together form a closed causal loop: at every moment in time, *M* (together with the position of the object) determines the state of *E*, while *S* (together with the agent's "neural" state) determines the state of *A*.

This decompositional analysis, I suggest, provides the basis for a mechanistic explanation of this particular system's "visual" categorization behavior: *A*, *E*, *M*, and *S* together constitute a *mechanism* which, when suitably described, "shows how" the active scanning behavior is produced by the underlying brain-body-environment system. Notably, the descriptions of *M* and *S* are entirely articulated in dynamical terms: *M*, the effect of the agent's motion on the state of the environment, is described as the rate of change in state variable *x* (i.e. the agent's horizontal position relative to the object); *S*, the effect of the environment on the state of the agent, is described as a series of "steady-state velocity fields"—phase portraits depicting the way the limit sets of subsystem *A*'s state space change over time with varying values of *x* and *y*. It is the subtle temporal interactions between the agent's limit behavior and the state of the environment which ultimately determine whether and how any particular object is caught or avoided. In this way (and although the details are largely omitted here), I suggest that Beer's dynamical analysis mechanistically explains "visual" categorization via active scanning; it identifies and describes a mechanism which is instantiated by the modeled 16-dimensional brain-body-environment system, and which is responsible for producing the observed behavior.

For the sake of contrast, there are at least two reasons why Beer's non-connectionist dynamical model and analysis should not be identified with the covering law account of scientific explanation. First, the model's state variables and parameters do not correspond to features of the phenomenon itself, but are instead features of the system from which that phenomenon arises. Second, the equations specified in the model are intended to account for a particular system only: Beer denies explicitly that the model is a model of categorical perception *in general* (Beer, 2003: 210). Therefore, the explanation he offers for "visual" categorization via active scanning differs considerably from covering law explanations, and is more plausibly associated with the mechanistic account.

The upshot of this discussion of Beer's dynamical model and dynamical analysis is that there exists at least one non-connectionist dynamical model which provides a genuine mechanistic explanation of a ("minimally") cognitive phenomenon. Insofar as this explanation is prototypical for the explanations offered in the field of evolutionary robotics, it is likely that there exists a large class of mechanistic dynamical explanations. Therefore, it seems wrong to suggest that the distinction between connectionist and non-connectionist dynamical modeling can be mapped onto the distinction between mechanistic and covering law explanation.

Conclusion

A number of general conclusions can be drawn from this discussion of dynamical modeling and scientific explanation.

For one, it highlights the importance, to the field of evolutionary robotics, of supplementing dynamical models of simulated brain-body-environment systems with detailed dynamical analyses of the behaviors those systems exhibit. So far, much research in evolutionary robotics has been content with providing *existence proofs* of mechanisms that exhibit "minimally cognitive" behavior. In contrast, Beer's (2003) study provides a template for how such behavior might be *explained*.

Why is explanation important? From the perspective of the broader cognitive sciences, simulations in evolutionary robotics seem ideally suited not only to provide existence proofs of various kinds of cognitive mechanisms, but also as a testing and development ground for the analytical tools necessary to explain those mechanisms. Because evolutionary robotics is in the unique position (relative to other branches of cognitive science) of having relatively well-understood general-purpose tools (e.g. CTRNNs and simulated environments) for modeling embodied and embedded cognitive systems, most of the research effort can go into the development of new analytic techniques with which to understand how these systems bring about increasingly complex cognitive phenomena. If any one technique turns out to be illuminating with respect to the behavior of some artificial agent, it is possible that the very same technique can be used to explain natural cognitive systems as well. In other words, although it remains to be seen to what extent the *mechanisms* identified in the field of evolutionary robotics relate to the real-world mechanisms that produce natural cognitive phenomena, there should be no question that the *tools* by which such mechanisms can be studied ought to be readily transferable.

I will conclude this discussion with a brief reflection on the nature and promise of the dynamical approach to cognition. Bechtel's (1998) claim that non-connectionist dynamical models provide covering law explanations of cognitive phenomena is a special case of the more general claim that the dynamical approach offers a particular *kind* of explanation. The current discussion calls this idea into question by suggesting that dynamical models and dynamical analyses may be involved in both covering law and mechanistic explanations—what matters is not that dynamical models are used, but *how* they are used. The view of the dynamical approach that emerges is the view of a scientific approach unified by common appeal to a particular set of mathematical tools and concepts, not by a common explanatory method. Indeed, it is likely that the dynamical approach to cognition provides a suitable mathematical framework in which to *combine* the relative advantages of mechanistic, covering law, and other forms of scientific explanation.

References

Bechtel, W. (1998). Representations and Cognitive Explanations: Assessing the Dynamicist's Challenge in Cognitive Science. *Cognitive Science*, *22*(3), 295–318.

Bechtel, W., & Abrahamsen, A. (2005). Explanation: a mechanist alternative. *Studies in History and Philosophy of Biol & Biomed Sci*, *36*(2), 421–441.

Bechtel, W., & Abrahamsen, A. (2006). Phenomena and mechanisms: Putting the symbolic, connectionist, and dynamical systems debate in broader perspective. In *Contemporary debates in cognitive science.* Oxford: Basil Blackwell.

Bechtel, W., & Richardson, R. (1993). *Discovering Complexity: Decomposition and Localization as Strategies in Scientific Research*. Princeton University Press.

Beer, R. (1995). Computational and dynamical languages for autonomous agents. In R. Port & T. van Gelder (Eds.), *Mind as motion: Explorations in the dynamics of cognition* (pp. 121–147). MIT Press.

Beer, R. (2003). The Dynamics of Active Categorical Perception in an Evolved Model Agent. *Adaptive Behavior*, *11*(4), 209–243.

Chemero, A. (2000). Anti-Representationalism and the Dynamical Stance. *Philosophy of Science*, *67*(4), 625–647.

Clark, A. (1997). *Being There: Putting Brain, Body, and World Together Again*. MIT Press.

Cummins, R. (1983). *The Nature of Psychological Explanation*. Cambridge, MA: MIT Press.

Elman, J. L. (1990). Finding structure in time. *Cognitive Science*, *14*(2), 179–211.

Elman, J. L. (1995). Language as a Dynamical System. In R. Port & T. van Gelder (Eds.), *Mind as motion: Explorations in the dynamics of cognition* (pp. 195–225). MIT Press.

Glennan, S. (2005). Modeling mechanisms. *Studies in History and Philosophy of Biol & Biomed Sci*, *36*(2), 443–464.

Hempel, C. (1965). *Aspects of scientific explanation and other essays in the philosophy of science*. Free Press.

Kelso, J. (1995). *Dynamic Patterns: The Self-Organization of Brain and Behavior*. MIT Press.

Machamer, P., Darden, L., & Craver, C. (2000). Thinking about Mechanisms. *Philosophy of Science*, *67*(1), 1–25.

Port, R., & van Gelder, T. (Eds.). (1995). *Mind as Motion: Explorations in the Dynamics of Cognition*. Cambridge, MA: MIT Press.

Rumelhart, D., & McClelland, J. (1986). PDP Models and General Issues in Cognitive Science. In D. Rumelhart, J. McClelland, & The PDP Research Group (Eds.), *Parallel Distributed Processing, Vol. 1* (pp. 110–146). MIT Press.

van Gelder, T. (1995). What Might Cognition Be, If Not Computation? *The Journal of Philosophy*, *92*(7), 345–381.

van Gelder, T., & Port, R. (1995). It's About Time: An Overview of the Dynamical Approach to Cognition. In R. Port & T. van Gelder (Eds.), *Mind as motion: Explorations in the dynamics of cognition* (pp. 1–43). MIT Press.