Computation and Dynamical Models of Mind

CHRIS ELIASMITH
Philosophy-Neuroscience-Psychology Program, Department of Philosophy, Washington University in St. Louis, Campus Box 1073, One Brookings Drive, St. Louis, MO 63130-4899, U.S.A.

Abstract. Van Gelder (1995) has recently spearheaded a movement to challenge the dominance of connectionist and classicist models in cognitive science. The dynamical conception of cognition is van Gelder's replacement for the computation bound paradigms provided by connectionism and classicism. He relies on the Watt governor to fulfill the role of a dynamicist Turing machine and claims that the Motivational Oscillatory Theory (MOT) provides a sound empirical basis for dynamicism. In other words, the Watt governor is to be the theoretical exemplar of the class of systems necessary for cognition and MOT is an empirical instantiation of that class. However, I shall argue that neither the Watt governor nor MOT successfully fulfill these prescribed roles. This failure, along with van Gelder's peculiar use of the concept of computation and his struggle with representationalism, prevent him from providing a convincing alternative to current cognitive theories.

Key words: Dynamical systems theory, computation, motivational oscillatory theory, representation, Turing machine, van Gelder, Watt governor.

1. Introduction

In his article 'What might cognition be, if not computation?' van Gelder attempts "to describe and motivate the *dynamical* conception sufficiently to show that it does in fact amount to an alternative conception of cognition, and one which is currently viable, as far as we can now tell" (van Gelder, 1995, p. 347, italics added).[1] However, even this modest claim involves theoretical difficulties which preclude its simple acceptance. In order to locate the *dynamicist* conception of cognition within the domain of cognitive science, van Gelder identifies relations amongst the current principal contenders: "the computational, connectionist, and dynamical conceptions"[2] (p. 345). He maintains connectionism to be a subcategory of the wider class of dynamicist systems (p. 370). Furthermore, he claims that the dynamicist agenda is indeed a novel, currently viable research program which is strictly opposed to the traditional computational/representational approach (p. 345).

I will argue that van Gelder's dynamicist program does not convincingly subsume connectionism, nor is it a plausible alternative to either connectionist or computational approaches. However, van Gelder does succeed in illustrating the importance of dynamical systems theory as a mathematical tool for describing complex, systemic behaviors – exactly those behaviors exhibited by many connectionist networks. In other words, van Gelder's larger project fails, but his discussion highlights considerations relevant to current cognitive science.

Minds and Machines **7:** 531–541, 1997.

In evaluating van Gelder's position, I will demonstrate the effects of his misapplication of the term *computational*. Subsequently, I will show that he constructs an ambiguous relation between the centrifugal governor (i.e. the dynamicist exemplar) and the human mind. From this relation van Gelder proceeds to identify important characteristics of the dynamicist conception of cognition, particularly the role of representation. Additionally, I will criticize van Gelder's reliance on the *motivational oscillatory theory* (MOT) of Busemeyer and Townsend (1993) as an empirical exemplar of dynamicism.

2. Defining Computation

Van Gelder's unorthodox use of the term *computational* permits him to introduce artificial theoretical distinctions. Specifically, he bisects cognitive science into *computational* and *noncomputational* camps. His reliance on this distinction prompts his assertion that cognition is "the behavior of some (*noncomputational*) dynamical system" (p. 358, italics added). However, once we free ourselves from van Gelder's unconventional application of the concept of *computation*, this assertion is patently false.

Van Gelder defines a computational system as follows (p. 366):

> A concrete computational system – a computer – is any system realizing an abstract computational system.

Though this explicit definition is neither eccentric, nor particularly contestable in its own right, his application of it in both the computational approach and connectionism is.

Van Gelder seeks to use his definition of a computational system to characterize the work of cognitive researchers such as Newell, Simon, Chomsky, Minsky and Anderson (Van Gelder and Port, 1995). In particular, Newell and Simon (1976) are cited by van Gelder as having definitively stated the 'computationalist hypothesis' as (p. 365):

> Natural cognitive systems are intelligent in virtue of being physical symbol systems of the right kind.

However, van Gelder's 'computational hypothesis' is a misnomer for a specific conjecture Newell and Simon themselves refer to as the *Physical Symbol System Hypothesis*. Rather than equating computation with symbol manipulation, as van Gelder has done, Newell and Simon understand a *symbol system* to be a specific *form* of a universal *computational system* (e.g. a Turing machine). A universal *computational system* is defined by Newell to possess characteristics including memory, symbols, and operations (1990, p. 77). However, natural cognizers (i.e *symbol systems*) are defined precisely by symbolicists, not simply as *computational systems*, but as: "symbol systems that are at least modest approximations of knowledge systems" (Newell, 1990, p. 113). In other words, Newell distinguishes the *type* of computer deemed necessary to exhibit cognition (i.e. one that will "provide a means to build representation systems") from the *family* of universal computers (Newell, 1990, p. 68). These representational systems are, for Newell, strictly *symbol systems* – not *computational systems* (Newell, 1990, p. 76). Therefore, it is more accurate to refer to van Gelder's *computationalists* as *symbolicists*.

A further consequence of this odd characterization of computation is van Gelder's assignment of the connectionist approach to the *noncomputational* camp (p. 345). Connectionist adherents unanimously construe themselves as committed to a computational view of mind; quintessential connectionists Churchland and Sejnowski (1992) appropriately titled their book *The Computational Brain*. In this work, Churchland and Sejnowski detail the connectionist commitment to complex dynamics and nevertheless candidly address *computational* problems: "Using the dynamical framework, we can begin to bring nonlinear networks to heel; that is, to understand their capabilities, and most important, to give us insight into how best to design networks to solve particular computational problems" (Churchland and Sejnowski, 1992, p. 89). Rather than resorting to the use of *computation* in the limited sense of van Gelder, they employ a more conventional conception – though it is one which does not contradict van Gelder's explicit definition – and define computers as "a physical device with physical states and causal interactions resulting in transitions between those states" (Churchland and Sejnowsi, 1992, p. 66). Under this conception, connectionists are clearly implementing abstract computational systems and thus, by van Gelder's own definition (p. 366), view mind as a computer. Therefore, where symbolicists suggest cognizers are *symbol manipulating* computers, connectionists posit *dynamic, distributed* networks in an effort to determine "what sort of computers *nervous systems* are" (Churchland and Sejnowski, 1992, p. 61).

Evidently, both connectionists and symbolicists define computation much more broadly than has van Gelder. Both camps claim to be examining a specific *type* of computation, and both regard the other as computational[3] – their disagreements arise as a result of a difference in *type* and not as a result of a difference in *kind*. Though van Gelder's definition of a computational system is not itself misleading, his application of it to cognitive approaches is: there is *no* such computational/noncomputational distinction applicable to connectionism and symbolicism. Therefore, connectionism is *not* naturally included in the noncomputational class of systems somehow delineated by the centrifugal governor as van Gelder claims (p. 370).

3. The Centrifugal Watt Governor

Van Gelder draws heavily on an analogy between the dynamical, centrifugal Watt governor and the functioning of the human mind to establish the validity of his dynamicist approach (p. 358). The Watt governor is a steam engine speed controller invented by James Watt in the late eighteenth century. It consists of a central spindle linked to the engine's main flywheel. The spindle has two hinged arms attached to it which end in small masses. As the flywheel turns the spindle, the masses move outward due to the effects of centrifugal force. The motion of the arms controls the throttle of the engine. Thus, engine speed is reduced the faster the spindle is turned and the farther outward the arms move. The result is that as the engine begins to increase in speed, the arms of the spindle are forced outward, closing the throttle and slowing the engine down. The slowing of the engine in turn allows the arms to move inward, opening the throttle and increasing engine speed. The overall effect of the Watt governor is to maintain a constant engine speed despite changes in load. Van Gelder intends the centrifugal governor to be a paradigm example of the class of dynamical systems to which the mind belongs (pp. 367, 369):

> The dynamical hypothesis in cognitive science, then, is the exact counterpart to the computational hypothesis: cognitive systems such as people are *dynamical* systems in the sense just laid out [i.e. as per the centrifugal governor], and cognition is state-space evolution in such systems. Alternatively, dynamicists are committed to the claim that the best model of any given cognitive process will turn out to be drawn from the dynamical sub-category of state-dependent systems.

A critical examination of the dynamical hypothesis exposes a series of difficulties with this conception of cognition. The most pressing theoretical issue to be resolved can best be phrased as a question: what relation exists between the centrifugal governor and natural cognizers? In other words, how is the brain supposed to be *like* the governor? It is in no way obvious that the centrifugal governor or its properties can somehow scale to the complexity of the human brain. However, there an three obvious ways in which the governor may be related to human cognition: the governor may play a role comparable to that of the symbolicist *Turing machine*; it may be a foundational unit for modeling cognition as is the connectionist *neuron*; or it may play an *analogical role* in motivating possibly new conceptions of cognition.

Van Gelder himself seems to favor the first option as he roughly equates the centrifugal governor's role with that of the Turing machine for the symbolicist view of mind (p. 358):

> Perhaps, that is, cognitive systems are more relevantly similar to the centrifugal governor than they are similar. . .[to] the Turing machine.

However, the centrifugal governor is very unlike the Turing machine (see Table I).

These differences reveal the centrifugal governor's inadequacy for fulfilling a central role for the dynamicists that parallels the Turing machine's importance to symbolicists. Fundamentally, the governor does not provide a means to quantitatively define a related class of cognitive systems. Without this ability it is impossible to determine whether or not a given system is 'relevantly similar' to the centrifugal governor, as there is no definitive, formal relationship between it and cognitive systems. Rather, van Gelder's previous claim would be better expressed as:

> Perhaps cognitive systems are of the class of systems which display complex dynamics that cannot be effectively captured by a Turing machine.

Table I. Comparison of Turing machine and centrifugal governor.

Turing machine	Centrifugal Governor
mathematical object	physical object
infinite	finite
formally proven to be strictly equivalent to other members of its class	anology or exemplar
representational	nonrepresentational

This proposal is satisfactory as it no longer implies any unlikely similarity between the roles of the Turing machine and the centrifugal governor, yet remains true to van Gelder's project. The inability of the governor to fulfill a Turing machine-type role for dynamicism leaves two options for understanding the relation between the mind and the centrifugal governor.

The second possibility – that the governor is comparable to a connectionist unit – can be dismissed. Not only does van Gelder himself explicitly reject this potential role but there is no biological or other justification for choosing the centrifugal governor as a basic unit for modeling cognitive behavior (p. 371). It is simply not a tenable option. The possibility remains that the centrifugal governor should play the part of an analogy to, or an exemplar of, natural cognitive systems.

Van Gelder's attempts to strengthen the role of the governor by comparing it to the Turing machine are indicative of his reluctance to assign the governor to the logically weak position of a simple example. Nonetheless, van Gelder has *only* this option remaining; eventually he must relegate the governor to the status of a dynamicist exemplar. As such it is intended to suggest properties which the class of dynamical systems appropriate for modeling cognitive behavior should have. Van Gelder implicitly places the governor in this role through his discussion of what he refers to as 'morals' that may be drawn from an analysis of the centrifugal governor problem. In brief they are (p. 358):

1. Fundamentally different kinds of systems can perform the same tasks
2. The belief that a cognitive task *must* be performed by a computational (i.e. symbolicist) system is false.
3. Cognitive systems may be dynamical in nature.

However, none of these 'morals' is revolutionary. Van Gelder's analysis of the governor problem has *not* provided novel insight into cognitive functioning. Without exception, these contentions have been frequently disputed in the cognitive science literature:

1. Symbolicist and connectionist systems are fundamentally different, and model similar cognitive tasks (Newell, 1990; Churchland and Sejnowski, 1992).

2. Connectionists believe cognitive tasks are performed by connectionist networks, not symbolicist systems (*ibid*; Pollack, 1990; Pouget and Sejnowski, 1990).
3. Many connectionists, including the Churchlands, Pollack, Meade, Traub, Hopfield and Smolensky, are committed to a dynamical view of cognition (*ibid*; Smolensky, 1988; Churchland, 1989).

Consequently, there is no need to introduce the centrifugal governor to bring any of these 'morals' to the attention of the cognitive science community – they have become evident since the re-introduction of connectionism in the late 1970s.

Even in its final possible role of exemplar, the governor does not present a novel conception of cognition. Inevitably the centrifugal governor fails to provide either theoretical insights or a formal basis for accepting van Gelder's dynamicism as an "alternative conception of cognition" (p. 347).

4. Dynamicist Representation

Despite the limitations of the centrifugal governor, van Gelder relies on characteristics of the centrifugal governor's behavior to offer what he refers to as another "way to understand cognitive processes" (p. 345). One of the most intriguing aspects of van Gelder's discussion is his insistence that the governor, as a cognitive exemplar, is *not* representational: "arm angle and engine speed are of course intimately related, but the relationship is not representational" (p. 351).

Van Gelder devotes much of Section I to affirming the non-representational nature of the centrifugal Watt governor. He ascertains that "the notion of representation is just the wrong sort of conceptual tool to apply" (p. 353) to a description of the centrifugal governor's dynamic behavior. The alternative to a representational description of the governor is, of course, a dynamicist one which relies on mathematics: "there is nothing mysterious about this relationship; it is quite amenable to mathematical description" (p. 353). Thus, van Gelder contends that an elegant description of the governor's behavior can be provided by a branch of mathematics referred to as *dynamical systems theory*.

As van Gelder's conclusion relates to the centrifugal governor, he is right. However, it is *not* obvious that these same arguments apply in an analogous manner to cognition. It has already been established that the relationship between the centrifugal governor and a cognitive agent is rather ambiguous. Thus, drawing any strong explanatory links between the governor's behavior and that of cognizers is tenuous. The danger of such analogical reasoning lies in the apparent attribution of *all* characteristics of the analogy's source (i.e. the centrifugal governor) to its target (i.e. a cognitive agent). Van Gelder falls victim to this danger with the concept of representation.

Van Gelder is aware of the fate of behaviorism[4] in the early 1960s and so realizes the peril of simply attributing the nonrepresentational nature of the governor to cognizers: "Insofar as the dynamical approach abjures representation completely, or offers some less powerful representational substitute, it may seem doomed" (p. 376). So, van Gelder is obligated to offer a solution to this difficulty. He does, though it is one which directly contradicts one of his most valued characteristics of the centrifugal governor. The solution: include representation in dynamicist models (p. 376).

With this suggestion, van Gelder has undermined his arguments in favor of a *noncomputational* view of cognition. It is mysterious how his desire to capture representation in noncomputational models can be reconciled with his conviction that: "These properties – representation, computation, sequential and cyclic operation, and homuncularity – form a mutually interdependent cluster; a device with any one of them will standardly possess the others" (p. 351).[5]

Furthermore, van Gelder's suggested methods of capturing representation in dynamicist models are *ad hoc*: "representations can be trajectories or attractors of various kinds, or even such exotica as transformations of attractor arrangements as a system's control parameters change" (p. 377). With such disparate possibilities as to the nature of dynamicist representation, it would be inconceivably difficult to distinguish representational from nonrepresentational behavior – why would the centrifugal governor *not* be representational under such a broad range of representational possibilities?

Van Gelder may have realized the arbitrariness of such a characterization of dynamical representation, as he offers a second argument to convince us that the dynamicist view of cognition can incorporate representation. However, this attempt is even more suspect: he claims that connectionist models, being representational and a dynamicist subclass, currently implement a solution which has been proven effective (p. 376).

While it is true that connectionist models are both representational and dynamical, they are not *dynamicist* models – they are *connectionist* models and bear little relationship to a centrifugal governor[6]. As early as 1988, Smolensky explicitly outlined a dynamical-representational-connectionist view of cognition (1988, p. 7):

> The intuitive processor is a subconceptual connectionist dynamical system that does not admit a complete, formal, and precise conceptual-level description.

Thus, the claim that connectionist models can be described using dynamical systems theory is hardly surprising; nor is it surprising that connectionist networks can be representational systems.

Van Gelder exclusively relies on the work of connectionists like Pollack and Grossberg (pp. 378, 375 ff.) prove the viability of the dynamicist position for modeling language in particular and cognition in general. However, van Gelder's dependence on connectionist research, and his inclusion of it in the dynamicist class of models, dilutes his initial description of the dynamicist approach to such a degree as to render it nearly indistinguishable from connectionism. In other words, subsuming connectionism under the 'dynamicist banner' (p. 375) is both poorly

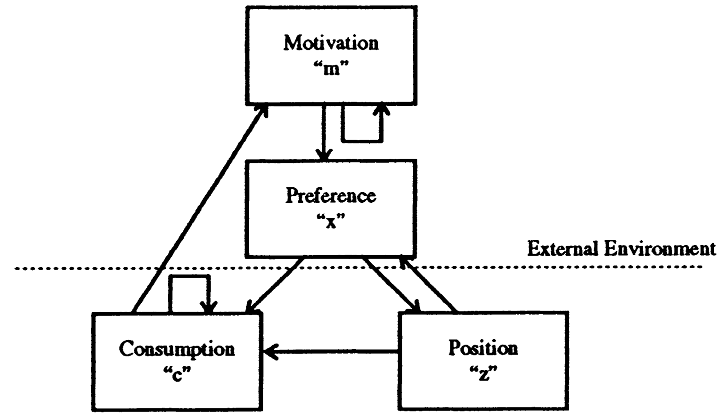

Figure 1. Signal-flow diagram for MOT (adapted from Townsend, 1993)

motivated, since neural networks are not related to the Watt governor, and self-defeating, since including connectionism in dynamicism contradicts fundamental commitments of dynamicism (e.g. a rejection of representation and computation).

5. The Motivational Oscillatory Theory and Dynamicism

To exemplify the dynamical hypothesis, van Gelder draws upon a cognitive model posited by Busemeyer and Townsend referred to as the *motivational oscillatory theory* MOT (Busemeyer and Townsend, 1993; Townsend, 1992). The model is closely related to the decision field theory of Busemeyer and Townsend (1993) and was developed to mimic cyclical milieus in motivated human decision making. The theory was intended to account for eating behaviors and was 'trimmed down' to include the following four aspects: (1) motivation; (2) consumption; (3) preference; and (4) the action in real time and space, based on distance (Townsend, 1992). These subsystems are highly interconnected, described by explicit differential equations, and act in parallel in an environment consisting of a single 'object of desire' (*ibid*.) (see Figure 1).

Van Gelder casts MOT in the role of an exemplar of dynamicism. However, van Gelder's characterization of dynamicism is not adhered to by MOT. Furthermore, accepting MOT as a dynamicist exemplar in no way demonstrates the ability of dynamicist models to effectively describe high-level cognitive processes. Contrary to van Gelder's dynamicist position, Townsend and Busemeyer claim that their model and theory are computational and that they are closely allied to neural descriptions: "computations are assumed to be realized by an underlying neural system" (Busemeyer and Townsend, 1993, p. 444). Not only do Busemeyer and Townsend violate van Gelder's noncomputational criterion but they suggest that such modeling is best used as a dynamical description of connectionist systems. Similarly, van Gelder's assertion that cyclicity is intimate only to symbolicist modeling (p. 351) is violated by a fundamental aim of MOT, as it was explicitly developed "in the hope[s] of mimicking more or less natural cyclical milieus" (Townsend, 1992, p. 220).

Nevertheless, MOT *does* employ low-dimensional differential equations in describing behavior. This characteristic, it seems, is enough for van Gelder to consider it a dynamicist model. He insists that this model will establish the viability of dynamicism being applied to *high-level* cognition:

> Consider the process of coming to make a decision between a variety of options, each of which has attractions and drawbacks. This is surely a high-level cognitive task, if anything is (p. 359).

Even the authors of MOT claim to be modeling a *high-level* decision process (Busemeyer and Townsend, 1993, p. 444):

> When confronted with a difficult personal decision, the decision maker tries to anticipate and evaluate all of the possible consequences produced by each course of action.

Paradoxically, the model that is *actually* offered is one of the decision to eat. Strikingly, this is *not* a "difficult personal decision" nor is it "surely a high-level cognitive task, if anything is." Rather, the model is one which focuses on 'biological drives' (Townsend, 1992, p. 221). As van Gelder later admits (p. 361):

> MOT (motivational oscillatory theory) enables modeling of various qualitative properties of the kind of cyclical behaviors that occur when circumstances offer the possibility of satiation of desires arising from more or less permanent motivations; an obvious example is regular eating in response to recurrent natural hunger.

Every living thing 'experiences' recurrent natural hunger – it seems dubious to call such biological drives 'high-level cognition' – and it is questionable whether such 'decisions' are of the same kind as those we are intrigued by when modeling decision making in humans. Furthermore, experimental support for the derivation of the equations which govern this decision making process were based on animal research of Miller (1959) and human research of Epstein and Fenz (1965); Busemeyer's and Townsend's sources are behaviorist (Busemeyer and Townsend, 1993, p. 442).

Though dynamicists, like Brooks (1991), *have* created dynamicist robots which exhibit impressive behaviors, it is uncertain whether the insect-like reactions of such successes will scale to the complex interactions involved in mammalian cognition. Furthermore, it is completely unknown whether such systems, including MOT, will be able to handle representation use by human cognizers. Van Gelder concludes (p. 362):

> There is thus no question that at least certain aspects of high-level cognitive functioning can be modeled effectively using dynamical systems of the kind that can be highlighted by reference to the centrifugal governor.

This is evidently unsubstantiated.

6. Conclusions

Van Gelder fails to establish dynamicism as an alternative view of cognition. The computational/noncomputational distinction which van Gelder relies on to group connectionism under the 'dynamicist banner' is an artificial construction. It does not serve to clearly distinguish movements within cognitive science nor is it supported by symbolicists, connectionists, or even the authors of a dynamicist exemplar, MOT. Furthermore, van Gelder's attempt to add important theoretical considerations (e.g. representation) to dynamicism force him to contradict his initial characterization of that approach. Consequently, van Gelder's reliance on a peculiar manipulation of the concept of computation and an inadequate exemplar, cause him to construct a cognitive view which, though nurturing important new mathematical tools for modeling the mind, encompasses insupportable claims.

Van Gelder repeatedly relies on the successes of connectionist research to support his assertions about the importance of dynamics in cognitive modeling. Perhaps his article is most usefully regarded as supporting the conjecture that more connectionist researchers *should* be focusing on the complex dynamics of cognition. Rather than dynamicism subsuming connectionism, the natural role for dynamical systems theory is one of describing the high-level and temporal behaviors of connectionist networks. Connectionism is not "perched somewhere in the middle" of symbolicism and dynamicism nor is it an "unstable mixture" of these two approaches (p. 374). Rather, it seems that connectionism is a potent combination of dynamical and representational commitments.

7. Acknowledgements

I would like to thank to Paul Thagard for his critique of previous drafts and Cameron Shelley for his numerous suggestions. I also would like to thank Poco deNada, James Townsend, Jerome Busemeyer and Tim van Gelder for providing valuable clarifications of related material.

Notes

[1] All unreferenced page numbers in the text refer to van Gelder (1995).
[2] Notably, van Gelder's use of term *computational approach* is considerably misleading. However, I will use van Gelder's terminology until the reasons for *not* doing so become apparent.
[3] Indeed, symbolicists even define connectionism as computational: "connectionism is a commitment to a particular neural-like computational technology" (Newell, 1990, p. 484). Thus, van Gelder seems to be alone in his categorization of computational and noncomputational cognitive approaches.
[4] van Gelder notes that the idea of representations "clearly supported the computation conception against a behaviorism which eschewed such resources, however, it was no use against a connectionism which helped itself to internal representations, though rather different in kind than the standard symbolic variety" (p. 346).
[5] Similarly, van Gelder later claims: "If you have a computational state-dependent system, it naturally

implements a system that is representational, sequential, cyclic, homuncular, and so on" (p. 372). So, is it *unnatural* for dynamicist systems to implement representations? If so, it does not seem possible that *we* are dynamicist systems, given that we naturally use representations.

[6]van Gelder himself realizes this when he states: "none of these [connectionist] properties obtains in the case of the centrifugal governor" (p. 361).

References

Brooks, R. (1991), 'Intelligence without Representation', *Artificial Intelligence*, **47**, 139–159.

Busemeyer, Jerome R., and Townsend, James T. (1993), 'Decision Field Theory: A Dynamic Cognitive Approach to Decision Making in an Uncertain Environment', *Psychological Review* **c**, 432–459.

Churchland, P. S., and Sejnowski, T (1992), *The Computational Brain*, Cambridge, MA: MIT Press.

Churchland, Paul (1989), *A Neurocomputational Perspective*, Cambridge, MA: MIT Press.

Newell, A. (1990), *Unified Theories of Cognition*, Cambridge, MA: Harvard University Press.

Pollack, J. (1990), 'Recursive Distributed Representations', *Artificial Intelligence*, **46**, 77–105.

Pouget, A. and Sejnowski, T. J. (1990), 'Neural Models of Binocular Depth Perception', *Cold Spring Harbour Symposia on Quantitative Biology*, Vol. 55, pp. 765–777.

Smolensky, Paul (1988). 'On the Proper Treatment of Connectionism', *Behavioral and Brain Sciences*, **11**, 1–23.

Townsend, James T. (1992), 'Don't be Fazed by PHASER: Beginning Exploration of a Cyclical Motivational System', *Behavior Research Methods, Instruments, & Computers*, **24**, 219–227.

van Gelder, T. (1995), 'What Might Cognition Be, If not Computation?', *The Journal of Philosophy*, **xci**, 345–381.

van Gelder, T. and Port, R. (1995), 'It's About Time: An Overview of the Dynamical Approach to Cognition', in R. Port, and T. van Gelder (Eds.), *Mind as motion: Explorations in the Dynamics of Cognition*, Cambridge, MA: MIT Press.