

Attractive and In-discrete

*A Critique of Two Putative Virtues of the Dynamicist Theory of Mind**

CHRIS ELIASMITH

University of Waterloo, Department of Philosophy, Waterloo, Canada ON N2L 3G1;

E-mail: eliasmith@uwaterloo.ca

Abstract. I argue that dynamicism does not provide a convincing alternative to currently available cognitive theories. First, I show that the attractor dynamics of dynamicist models are inadequate for accounting for high-level cognition. Second, I argue that dynamicist arguments for the rejection of computation and representation are unsound in light of recent empirical findings. This new evidence provides a basis for questioning the importance of continuity to cognitive function, challenging a central commitment of dynamicism. Coupled with a defense of current connectionist theory, these two critiques lead to the conclusion that dynamicists have failed to achieve their goal of providing a new paradigm for understanding cognition.

1. A Shiny New Paradigm

Cognitive scientists have recently been told that they are doing it all wrong. They have been told that it is time for a Kuhnian paradigm shift in our understanding of cognitive systems and minds. Of course, in order for such a shift to occur, we need to have a paradigm to shift *to*. Those calling for the change have provided one: dynamicism (Thelen and Smith, 1994; van Gelder, 1995; van Gelder and Port, 1995). The ‘dynamicists’ tell us that, rather than thinking of cognitive systems as computing and representing, we need to embrace the embedded, continuous, “attractive”, in short, *dynamical* character of cognition. They claim that in doing so we will discover the true nature of cognitive functioning, from bottom to top.

My goal in this paper is to determine if two of the central virtues of dynamicism (continuity and attractor dynamics) are, even potentially, the boon to our understanding of cognitive function that dynamicists claim. First, I show that the dynamicist reliance on low-dimensional attractor dynamics: 1. raises a tension between a central commitment of the dynamicist and her or his rejection of the symbolicist (i.e., computationalist or classicist) paradigm; and 2. results in deep difficulties for explaining high-level cognition within a dynamicist framework. Second, and more importantly, I argue that the putative continuity of cognitive systems is not relevant to understanding cognitive function. Because continuity is not relevant to cognition, I contend that dynamicist arguments against the computational/representational commitments of current paradigms fail. Given these arguments, I conclude that dynamicism will not be the panacean paradigm its proponents predict.

Before delving into these specific arguments, it is helpful to have some understanding of the historical and theoretical roots of the dynamicist position. Clearly,



Minds and Machines 11: 417–426, 2001.

© 2001 Kluwer Academic Publishers. Printed in the Netherlands.

dynamicists wish to reject both connectionism and symbolicism in favor of a third view of what it is to be a cognitive system. Historically, symbolicism has been their primary target (van Gelder, 1995, 1998; van Gelder and Port, 1995; Thelen and Smith, 1994; Globus, 1992). Indeed, they have provided valuable criticisms of this classical approach to cognitive science (exemplified in, e.g., Newell, 1990). Their critiques provide insight into the kinds of difficulties researchers may expect to encounter in characterizing the mind as a physical symbol system, and are valuable in this role. What I call into question are the dynamicists' strong claims that they can provide successful arguments for a wholesale *dynamicist* revolution to understanding cognition.

The theoretical roots of dynamicism are derived from the mathematical theory known as 'dynamical systems theory', which uses sets of difference or differential equations to describe the evolution of a system through time. This theory has been successful in revolutionizing our understanding of many phenomena, including the weather, animal population dynamics, and economic change. Dynamicists feel that such a revolution concerning cognition can also be affected by dynamical systems theory. Certain strengths of these mathematical tools have inspired Timothy van Gelder to formulate what he calls the "Dynamicist Hypothesis" (1995, p. 4):

Natural cognitive systems are certain kinds of dynamical systems, and are best understood from the perspective of dynamics.

Under this hypothesis, the concepts of dynamical systems theory, which include 'state space', 'attractor', 'trajectory', and 'deterministic chaos' are used to explain the internal processing that underlies a cognitive system's interactions with its environment. Furthermore, clarifying precisely which "certain kinds" of systems are important to cognition has resulted in dynamicists embracing a mandate to "provide a *low-dimensional* model that provides a scientifically tractable description of the same qualitative dynamics as is exhibited by the high-dimensional system (the brain)" (van Gelder and Port, 1995, p. 28, italics added). Thus, dynamicist models should "rely on equations using many fewer parameters" (ibid., p. 27) than those typically used by connectionists. Without this corollary, it is no easy task to distinguish dynamicism from connectionism. Without that distinction, dynamicists will not succeed in establishing a new paradigm, as is their ambition (Thelen and Smith, 1994; van Gelder and Port, 1995).

2. Attractive

As noted in the introduction, dynamicist models characterize cognitive function using the mathematical concepts of dynamical systems theory. Central among these is the notion of an *attractor*. Simply put, an attractor is a point or set of points in state-space towards which starting points of the system will tend over time. Each such attractor will have a *basin of attraction* that defines a region of the state-space. A system at any point in that region will tend to move towards the associated attractor in a noise free system.

When pushed on the applicability of the dynamicist hypothesis (and related models) to higher cognitive processes such as language, proponents often claim that such attractors and their basins can be taken to be concepts stored in the system (Zeeman, 1965; Amit, 1995). However, as David Mumford (1997) has recently noted, “[t]his makes concepts basically Boolean and discrete: the dynamical system cannot fall partly into two such basins of attraction, so the model is closer to classical logic than to fuzzy logic or to probability models” (p. 247). Of course, this comparison to classical logic is not something that would sit well with a dynamicist. Classical logic forms the basis for *symbolicism* and operates over *discrete symbols*; both of which have been rejected by the dynamicist. But, there is nothing in the dynamicist hypothesis that introduces the ‘fuzziness’ necessary to avoid this problem (and chaos will not do the trick), so Mumford is correct. Thus, there is a tension between the dynamicist commitment to attractor dynamics for explaining higher cognition and their rejection of classical, discrete symbol systems.

Furthermore, because dynamicists are committed to low-dimensional descriptions, their models (unlike classical symbol systems) are not flexible enough to capture the richness of our conceptual life. A system with only a few dimensions generally has fewer and less complex attractors. So, low dimensionality typically means less flexibility. Considering (1) that the average adult has a vocabulary of well-over 15,000 words (that combine to encode far more concepts), (2) that these are learned at a rate of about 2–3 per day around the age of 2, and (3) that we rapidly manipulate and encode complex structural relations (e.g., in generating analogies (Eliasmith and Thagard, 2001)), it seems highly unlikely that a *low-dimensional* description of cognitive function will be adequate. In other words, low-dimensional attractor dynamics probably can’t do the explanatory work that many dynamicists assume they can.

Now, admittedly, when it comes to cognition it just isn’t clear one way or the other how ‘high’ the dimensionality has to be. But, models taken to be typical exemplars by dynamicists use around 3–8 dimensions for explaining capacities like “motivation” and “decision making” (see, e.g., Busemeyer and Townsend, 1993; van Gelder, 1995). While dynamicists *claim* such models are cognitive models (van Gelder, 1995, p. 359), it simply isn’t possible to satisfactorily explain the complex processes involved in a typical instance of making a decision with this kind of model (Eliasmith, 1997). A high-level description of the alternatives alone would prove more than such a model could handle. It is, then, *highly* unlikely that low-dimensionality of this kind will do when it comes to understanding cognition.

Despite a reliance on dynamical analyses, connectionism is in a far different position. For one, connectionist models are high-dimensional. Thus, connectionists do not face the same problems regarding either flexibility or the encoding of many concepts, structures, etc. Unlike low-dimensional dynamicist models, connectionist models can be more sensitive to a quickly changing environment. For example, high-dimensionality allows for the construction of temporary attractors that correspond to cognitive structures such as relations, analogical mappings, and

property bindings. As well, a high-dimensional state space has ‘room’ for many easily distinguishable vectors, which may be concepts, memories, schemas, etc. (see, e.g., Eliasmith and Thagard, 2001). In addition, many connectionist models are inherently probabilistic, and embrace uncertainty as central (Yuille and Geiger, 1995). Under such conditions, ‘conceptual attractors’ lose their strong discreteness. Of course, we can wonder if dynamicists could simply include probability in their models to similarly avoid the discreteness problem. If they do so, however, they will limit the number of distinguishable attractors in their system even further; the wider (i.e., ‘fuzzier’) the boundary between concepts, the fewer *different* ones can be encoded.¹ Thus, the high-dimensionality of connectionist systems allows them to include *both* uncertainty and large numbers of distinguishable states. These properties of connectionist models allow for more realistic modeling of our conceptual life than is possible if we adopt the dynamicist hypothesis.

Although dynamicist models concerned with cognition involving concepts are rare (if existent), it is important to consider what resources are available for capturing higher level cognition given the dynamicist hypothesis. Language and conceptual analysis may be on the back burner for dynamicism, but if dynamicist commitments are inapplicable to such cognitive behaviors it is unclear why dynamicism should be considered a cognitive theory, let alone a paradigm.

The considerations in this section by no means *prove* that dynamicism is unable to capture high-level cognition. However, they do highlight shortcomings given current dynamicist commitments. Perhaps those commitments will change to incorporate high-dimensional, probabilistic models. If this happens, however, dynamicism has simply become connectionism and offers nothing new to cognitive science.

3. In-Discrete

One of the reasons Mumford’s (1997) criticism of dynamicist models should be taken seriously by dynamicists is that dynamicists have often stressed the putative fact that cognitive systems are not discrete. In fact, they have relied on this ‘fact’ to establish some of their more controversial claims. Perhaps the best example is the dynamicist contention that cognitive systems are, contra connectionism and symbolicism, noncomputational (Globus, 1992, p. 304; van Gelder, 1995).

Notably, connectionists and dynamicists disagree on how best to understand what computation *is*. So, the strong dynamicist claim against connectionism seems to be, in its most obvious form, merely a result of differences in definition. Dynamicists hold that for a system to be computational, its evolution must be specifiable by means of rules of *symbol* manipulation (Globus, 1992; van Gelder, 1995). Here, a symbol is a discrete object that stands in a representational relationship with some state of affairs. This definition would sit well with most symbolicists, and is strongly influenced by the serial, digital computer as a paradigm case (Newell,

1990). So, in rejecting this notion of computation, dynamicists are again rejecting a classical symbol systems approach to understanding cognition.

However, connectionist definitions tend to be somewhat more general. For example, Churchland and Sejnowski (1992) define a computer as: "a physical device with physical states and causal interactions resulting in transitions between those states" (p. 66) and feel that "once we understand more about what sort of computers *nervous systems* are, and how they do whatever it is they do, we shall have an enlarged and deeper understanding of what it is to compute and represent" (p. 61). Notably, then, Churchland and Sejnowski seem to be providing a much weaker account of computation than the dynamicists.² Rather than defining computation outright, connectionists adopt this account and take it as part of their purpose to derive a fuller account of computation by figuring out what properties are shared by systems that we can usefully understand as computing. In any case, both parties agree that the nervous system is "quite unlike the serial, digital machines on which computer science cut its teeth" (ibid., p. 7). So connectionists hold that classical symbol systems aren't the *right kind* of computer for understanding cognition.

The central difference between the dynamicist and connectionist definitions is the conspicuous absence of the notion of a 'symbol' from the connectionist definition. However, connectionists do take there to be representations. So, connectionists do not embrace symbols *per se*, but they do speak of representations. Dynamicists, in contrast, tend to reject both. So, the dynamicist definition of computation does conflict with connectionism if we replace 'symbol' with 'representation' understood more generally. And, dynamicists are quite happy to make this stronger, *nonrepresentational* claim (Globus, 1992; Thelen and Smith, 1994; van Gelder, 1995, 1998).

So, dynamicists think cognitive systems aren't computational because they think cognitive systems don't traffic in symbols or representations. Connectionists do think cognitive systems are computational (and representational) because they take computation (and representation) to be useful notions for describing cognitive systems.

So much for definitional differences. The important question is what kinds of arguments do dynamicists muster for their rejection of computation (and representation)? If we can understand how these arguments work (or fail) regardless of the definitions being employed, we will have a deeper understanding of the strengths (or weaknesses) of dynamicism.

In general, dynamicist arguments against computation (and representation) rely heavily on the purportedly continuous nature of cognitive systems. They posit, as seems reasonable, that *if* the cognitively relevant level of a system is continuous in time *then* discrete symbols/representations (and hence computation, by definition) will not be adequate for understanding cognition. It is quite clear that dynamicists wish to affirm the antecedent of this conditional, making the consequent only a *modus ponens* away. How then, does the dynamicist argue for the antecedent, i.e. continuity of cognitive systems?

Temporal continuity, for dynamicists, is an obvious property of cognitive systems. Van Gelder and Port (1995) find their evidence in an analogy between cognition and the motion of an individual's arm: "No matter how finely time is sampled, it makes sense to ask what position your arm occupies at every sampled point. Now, the same is true of cognitive processes" (p. 14). Thus, continuity, for them is "just an obvious and elementary consequence of the fact that cognitive processes are ultimately physical processes taking place in real biological hardware" (p. 15). The dynamicist commitment to continuity is reflected in their talk of "flow", "participatory, unpredictably harmonizing self-evolution", "covariation", "self-generating dynamic evolution", "state-space evolution" (Globus, 1992; Thelen and Smith, 1994; van Gelder, 1995; van Gelder and Port, 1995). Reliance on this sort of a vocabulary shows their strong belief in the relevance of continuity to providing accurate descriptions of cognitive systems. In his seminal paper on the virtues of dynamicism, van Gelder puts the point strongly: "the system's entire operation is smooth and continuous; there is no possibility of nonarbitrarily dividing its changes over time into distinct manipulatings, and no point in trying to do so" (van Gelder, 1995, p. 354).

What is it that makes continuity so obvious, and obviously important? The obviousness, as van Gelder remarks above, is simply due to the fact that cognitive systems are real physical systems. Our best physical theories tell us that space and time are continuous. Physical systems, by definition, exist in space and time and are thus continuous. Of course, the difficulty with such reasoning is that all putative or potential cognitive systems (including the loathed serial, digital computer) are subject to these claims. So, the mere obviousness of the continuity of cognitive systems isn't very interesting. This means that the really interesting question is: Why is continuity so obviously important for understanding *cognitive systems* as distinct from other kinds of systems?

Dynamicists reason that because the brain is continuous it processes analog signals, and these are best described by real numbers (as opposed to integers, or rationals). They presume that the important sets of numbers for describing brain state evolution lie in a continuous interval (or a set of intervals) on the real number line. Clearly, a discrete system such as a digital computer can only represent such numbers with a finite precision, and in a discontinuous manner. This finite, discrete representational capacity limits the areas of the real number line 'accessible' to such a representer (i.e., only the rationals can be represented). Because dynamicists suppose that continuous parts of the number line are important to understanding the brain (i.e., the system which underlies all of our cognitive functions), discrete descriptions are deemed inadequate for explaining cognition (van Gelder, 1998, pp. 618, 620).

To render this argument less abstract, suppose that a neuron's spike train (i.e., the set of nearly identical rapid voltage changes, or 'action potentials') encodes some signal of interest. It is not immediately clear how much information is encoded by a given spike. If we take the distance between spikes in the train to be

the basis of neural signal passing, then it is conceivable that such spiking patterns are describable only by real numbers. This is so because the *precise* distances between spikes can only be expressed by a real number (since time and distance are continuous). In other words, if the analog properties of the neuron are central to information passing, then it is possible that neurons are sensitive to a degree of precision not achievable by digital computers.

However, this line of reasoning misses the important role of noise and uncertainty in any physical system that propagates information. Assuming that the *exact* distance between any two neural spikes is the relevant measure of information entails that an infinite amount of information has been encoded by the 'sending' neuron (and can be decoded by the 'receiving' neuron). This result is entailed by that assumption because a real number can only be precisely represented by an *infinite* bit string. Now, it may seem highly unlikely that a real neuron could actually pass or use an infinite amount of information. But, even more problematic is the fact that if there is *any* expectation of noise or uncertainty in the signal being passed from one neuron to the next, the actual precision of a neural code will drop dramatically.

There are a number of good biological reasons to think that neurons are operating in an uncertain environment. For example, synapses have been found to be rather unreliable in their release of vesicles into the synaptic cleft given the presence of an action potential in the presynaptic axon (Stevens and Wang, 1994). As well, the amount of neurotransmitter in each vesicle can vary significantly, as can the ability of the presynaptic neuron to release the vesicles (Henneman and Mendell, 1981). And lastly, axons themselves have been shown to introduce jitter into the timing of neural spikes (Lass and Abeles, 1975). So, even the 'wires' used to pass the signal introduce noise. Nevertheless, neurons have been shown to reproduce and respond similarly (though not identically) to similar signals (Baer and Koch, 1994; Gallant, 1994).

Given the empirical fact of the matter concerning the noisiness of the neurons' environment and their ability to extract and pass signals, severe limits have been found on the precision of neural codes. In fact, it seems that neurons tend to encode approximately 3–7 bits of information per spike (Bialek and Rieke, 1992, see also Rieke et al., 1997). The technicalities of arriving at this number are beyond the scope of this paper, but it is important to note that these results *do not* rely on discretizing the neural spike train. In other words, this limit is clearly not a result of instrument limitations or preprocessing of spiking behavior, it is a limitation of neurons themselves. These kinds of information theoretic results are quickly becoming central to many analyses in computational neuroscience (Bower, 1998). Given this sort of evidence, the continuous nature of neurons is not relevant to the information they process. Three bits of information is far more information per spike than some have claimed (e.g. Cummins, 1980, p. 189) but it is far less than the infinite amount of information needed to encode a real number. Information

processing in the brain, then, can be equally well described as continuous and noisy, *or* discrete.

What does all this mean to the dynamicist position? It means that continuity just isn't relevant to understanding cognitive systems. *If* dynamicists are materialists *and* they think that continuity is central to cognition *then* there is empirical evidence against their position. We can safely assert the antecedent, and *modus ponens* our way to the consequent. In other words, the effects of noise on encoding precision show that a central claim of dynamicism is wrong. Continuity isn't important to understanding cognitive systems.

It would be unfair, however, to claim that only dynamicists fall prey to this empirical result. One of the best known proponents of connectionism, Paul Churchland (1995), had made similar claims in his book *The engine of reason, the seat of the soul* (p. 243):

Genuinely parallel implementation is important for the further reason that only then will the values of all of the variables in the network . . . have open to them every point in the mathematical continuum. So-called "digital" or *discrete-state* computing machines are limited by their nature to representing and computing mathematical functions that range over the *rational* number . . . This is a potentially severe limitation on the abilities of a digital machine . . . Therefore, functions over real numbers cannot strictly be computed or even represented within a digital machine. They can only be approximated.

From his brief discussion, it remains unclear how parallel computation is key to providing representations of all reals.³ What is clear is his claim that a limitation of digital computers (i.e., discrete-state machines) is that they do not have the necessary access to the real number line. However, from the neurophysiological data, it is evident that a well-chosen discrete-state machine does have the same access to the real number line as cognitive systems do since only about three bits of information need be represented per neural spike.

The mistakes of individuals aside, do dynamicism and connectionism, *as theories*, suffer differently from a rejection of the importance of continuity? I think so. Because dynamicists wish to reject computation and representation on the basis of arguments from continuity, empirical evidence to the contrary makes these arguments unsound. If dynamicists are unable to establish the importance of continuity, they cannot *modus ponens* their way to a rejection of computation. Connectionists, of course, have no need for such arguments; they embrace both representation and computation. So, dynamicism *alone* is significantly less plausible as a cognitive theory for having misidentified the relevant properties of cognitive systems.

4. Conclusion

It is ironic that the putative virtues of attractor dynamics and continuity leave dynamicism so unconvincing. In fact, the failing of both virtues leave dynamicists less able to distinguish themselves from the symbolicists they ridicule. First, because

attractor dynamics are discrete, dynamicist accounts of concepts are not easily distinguishable from symbolicist ones. As well, because dynamicist models are restricted to being low-dimensional, they don't have the representational capacity to account for high-level cognitive phenomena. Worse yet, any attempt by dynamicists to introduce 'fuzziness' into concept representation (in order to distinguish themselves from symbolicists) will reduce that representational capacity of dynamicist systems even further. Second, because continuity is not relevant to cognitive function, dynamicist arguments to noncomputationalism and nonrepresentationalism become unsound, disallowing their reasons for rejecting some symbolicist commitments.

In contrast, connectionism falls prey to neither of these problems. Although individual connectionists may make similarly mistaken theoretical claims (e.g., Churchland's claims about continuity), as a cognitive paradigm connectionism does not share these commitments. Indeed, it is when dynamicists attempt to distinguish their position from contemporary connectionism that many of their theoretical difficulties arise.

In conclusion, the dynamicist preoccupation with continuity and low-dimensionality are not convincingly motivated. Furthermore, relaxing these dynamicist commitments leaves the position indistinguishable from current connectionism. So, in the face of the limitations of attractor dynamics as currently conceived, and in the face of evidence for the finiteness of information capacities of real neural systems, dynamicism does not present a compelling new cognitive paradigm.

Acknowledgements

This research was supported by the Social Sciences and Humanities Research Council of Canada and the McDonnell Foundation.

Notes

¹This is a straightforward consequence of information theory (see Reza, 1994).

²Or so it initially seems. Churchland and Sejnowski's account relies on general notions like 'causality' which allow most things to count as computing. Of course, Searle (1990) has long pointed out that standard definitions of computing (like the dynamicist one) also allow most things to count as computing.

³As odd as it is, this aspect of Churchland's position is not relevant to the point I'm making here.

References

- Amit, D.J. (1995), 'The hebbian paradigm reintegrated: Local reverberations as internal representation', *Behavioral and Brain Sciences* 18, pp. 617–657.
- Baer, W. Kock, C. (1994), *Precision and reliability of neocortical spike trains in the behaving monkey*. The Neurobiology of Computation: Proceedings of 3rd Computational and Neural Systems Conference, Kluwer.

- Bialek, W. and Rieke, F. (1992), 'Reliability and information transmission in spiking neurons', *Trends in Neurosciences* 15 (11), pp. 428–434.
- Bower, J.M. (ed.) (1998), *Computational neuroscience: Trends in research 1998*. Elsevier.
- Busmeyer, J.R. and Townsend, J.T. (1993), 'Decision field theory: A dynamic-cognitive approach to decision making in an uncertain environment', *Psychological Review* 50, pp. 432–459.
- Churchland, P. (1995), *The engine of reason, the seat of the soul: a philosophical journey into the brain*, Cambridge, MA: MIT Press.
- Churchland, P.S. and Sejnowski, T. (1992), *The computational brain*, Cambridge, MA: MIT Press.
- Cummins, R. (1980), Functional analysis. *Readings in philosophy of psychology*. N. Block, Cambridge, MA: Harvard University Press. 1, pp. 185–190.
- Eliasmith, C. and Thagard, P. (2001), 'Integrating structure and meaning: A distributed model of analogical mapping', *Cognitive Science* 25(2) pp. 245–286.
- Eliasmith, C. (1996), 'The third contender: a critical examination of the dynamicist theory of cognition', *Philosophical Psychology* 9 (4), pp. 441–463.
- Eliasmith, C. (1997), 'Computation and dynamical models of mind', *Minds and Machines* 7, pp. 531–541.
- Gallant, J.L., Conner, C.E. et al. (1994), 'Responses of visual cortex neurons in a monkey freely viewing natural scenes', *Society of Neuroscience Abstracts* 20, p. 1054.
- Globus, G.G. (1992), 'Toward a noncomputational cognitive neuroscience', *Journal of Cognitive Neuroscience* 4 (4), pp. 299–310.
- Henneman, E. and Mendell, L. (1981), Functional organization of motoneuron pool and its inputs. *Handbook of physiology: The nervous system*. V.B. Brooks. Bethesda, MD: American Physiological Society. 2.
- Lass, Y. and Abeles, M. (1975), 'Transmission of information by the axon. I: Noise and memory in the myelinated nerve fiber of the frog', *Biological Cybernetics* 19, pp. 61–67.
- Mumford, D. (1997), *Issues in the mathematical modeling of cortical functioning and thought*.
- Newell, A. (1990), *Unified theories of cognition*, Cambridge, MA: Harvard University Press.
- Reza, F.M. (1994), *An introduction to information theory*, New York: Dover.
- Searle, J.R. (1990), 'Is the brain a digital computer?', *Proceedings and Addresses of the American Philosophical Association* 64, pp. 21–37.
- Stevens, C.F. and Wang, Y. (1994), 'Changes in reliability of synaptic function as a mechanism for plasticity', *Nature* 371, pp. 704–707.
- Thelen, E. and Smith, L.B. (1994), *A dynamic systems approach to the development of cognition and action*, Cambridge: MIT Press.
- van Gelder, T. (1995), 'What might cognition be, if not computation?', *The Journal of Philosophy* DCI (7), pp. 345–381.
- van Gelder, T. (1998), 'The dynamical hypothesis in cognitive science', *Behavioral and Brain Sciences* 21 (5), pp. 615–665.
- van Gelder, T. and Port, R. (1995), It's about time: An overview of the dynamical approach to cognition. *Mind as motion: Explorations in the dynamics of cognition*, R. Port and T. van Gelder. Cambridge, MA: MIT Press.
- Yuille, A. and Geiger, D. (1995), Winner-take-all mechanisms, *Handbook of Brain Theory and Neural Networks*, M. Arbib. Cambridge, MA: MIT Press, p. 1056.
- Zeeman, E.C. (1965), Topology of the brain, *Mathematics and Computer Science in Biology and Medicine*. London, Medical Research Council: 277–292.