

Explaining the Behaviour of Springs, Pendulums, and Cognizers

Michael Wheeler

CSRP 284, June 1993

Cognitive Science Research Paper

Serial No. CSRP 284

The University of Sussex
School of Cognitive and Computing Sciences
Falmer
BRIGHTON BN1 9QH
England, U.K.

Explaining the Behaviour of Springs, Pendulums, and Cognizers

Michael Wheeler

School of Cognitive and Computing Sciences
University of Sussex, BRIGHTON BN1 9QH, U.K.
E-Mail:michaelw@cogs.susx.ac.uk

Abstract

The received opinion in mainstream cognitive science is that cognition is the manipulation of representations by computational information processing mechanisms. However, in this paper, I argue that cognition ought to be conceptualized not as computation, but as state-space evolution in certain classes of non-computational dynamical systems. Following a review of van Gelder's arguments for this dynamical cognition hypothesis, I suggest that the most compelling reasons for adopting a dynamical perspective come from research into situated (world-embedded) autonomous agents. We should think of cognitive architectures as control systems for the situated activity of sufficiently complex embodied agents. I draw on key work in the simulation of adaptive behaviour which indicates that situatedness is best explained using the concept of dynamical coupling. I then present a preliminary sketch of other ways in which the conceptual framework of dynamics provides a powerful basis for thinking about and studying cognition.

1 Introduction

Dynamics is a field emerging somewhere between mathematics and the sciences. In our view, it is the most exciting event on the concept horizon for many years. The new concepts appearing in dynamics extend the conceptual power of our civilization and provide new understanding in many fields. (Abraham & Shaw, 1992)

Until now, cognitive science has attempted to explain cognition in terms of the construction and manipulation of representations by computational information processing mechanisms. This empirical hypothesis, which is accepted by both classical symbolists and the overwhelming majority of connectionists alike, has been transformed into an unchallenged article of faith. But the tide is turning. An interdisciplinary wave of discontent threatens to soak the cognitive science establishment. The notion that cognition is computation no longer defines the entire research programme called cognitive science.

But it is all very well calling for the overthrow of the current orthodoxy; throwing eggs is not the hardest part. The real challenge is to find something more powerful to put in place of the supposedly outmoded framework. With this aim in mind, there is evidence that the embryonic insurrection is beginning to cohere around the idea that cognition is best thought of in dynamical terms, i.e., in the language of dynamical systems theory. In short, one current, revolutionary hypothesis is that "cognition is state-space evolution in certain kinds of non-computational dynamical system" (van Gelder, 1992, p.1)

In essence, a dynamical system is any system for which, potentially at least, we have a rigorous analysis of the way it evolves over time. Formally, this is any system for which we can, in principle, provide a state space evolution equation, an equation which tells us how the values of the state variables of the system change with time. (What the particular state variables of interest are will be decided in context.) At various points in this paper, I shall employ the technical language of dynamical systems theory. The terms and ideas which I use will be explained as and when they are introduced. I hope this will make the overall

argument accessible to those readers who are not conversant with the field of dynamics. For a thorough introduction to the discipline, see (Abraham & Shaw, 1992) ¹

It has to be said that my title is somewhat misleading as to the exact orientation of the paper. The behaviour of springs and pendulums has been extensively studied by dynamicists (see, for example, (Abraham & Shaw, 1992) and (Baker & Gollub, 1990)). Consequently, I shall concentrate on the claim that essentially the same conceptual language appropriate to the study of springs and pendulums carries over to the study of cognizers. At the risk of sacrificing dramatic tension, I should begin by openly admitting that my personal sympathies lie with the dynamical rebels. However, that does not commit me to the position that every statement published by the emerging cadres should be treated as a gospel truth, with disagreement punishable by a ten year contract to recreate human-level intelligence using only production rules and a heuristically guided depth-first search algorithm.

2 On The Importance of James Watt to Cognitive Science

The name of James Watt is *not* synonymous with issues in the foundations of cognitive science. However, this situation may be only temporary, because an invention usually attributed to this famous engineer plays a key role in the arguments advanced by van Gelder in support of the hypothesis that cognition is fundamentally dynamical (van Gelder, 1992).

The Watt governor is a highly successful solution to an engineering problem which first arose during the industrial revolution. Consider the following situation:

1. a steam engine is being used to run spinning machinery via a connecting flywheel;
2. the speed of that flywheel needs to be kept constant in the face of fluctuations in engine workload and boiler pressure;
3. the speed of the flywheel is best controlled by adjusting the amount of steam entering the pistons from the boiler via a throttle valve.

What is required is a device (a ‘governor’) to turn the throttle valve so as to preserve a constant flywheel speed in spite of the destabilizing fluctuations. Watt’s elegant solution was to attach a vertical spindle to the flywheel in such a way that the speed of rotation of the spindle was directly dependent upon the speed of rotation of the flywheel. Two arms, each with a metal ball at its end, were connected to the spindle by hinges so that, as the spindle turned with the flywheel, the balls, affected by centrifugal force, would move either outwards and upwards (as the speed of the flywheel increased) or inwards and downwards (as the speed of the flywheel decreased). This motion on behalf of the arms was linked to the throttle valve by an arrangement which resulted in the valve being closed as the flywheel speed increased (restricting the flow of steam), and opened as the flywheel speed decreased (allowing more steam to flow). The upshot was a very effective method for maintaining constant engine speed. All very clever of course, but what has it got to do with cognition?

Van Gelder argues that there is a crucial lesson to be learnt by cognitive science from a careful consideration of the centrifugal (Watt) governor. In broad terms, this is because the behaviour of the centrifugal governor cannot be analysed satisfactorily using the concept of representation. And this is not because the governor is too simple a device to warrant such an explanation; it is precisely because the notion of

¹During the course of this paper, I knowingly elide a distinction between dynamical systems as those things in the world which can be described by state space evolution equations, and dynamical systems as idealized, mathematical models of real-world systems. The latter of these perspectives is almost certainly the more scientifically rigorous (Abraham & Shaw, 1992). All sorts of foundational questions of realism versus instrumentalism in science are lurking just around the corner. But, as important as it must be to recognize and maintain the distinction in certain contexts, I feel that it is of little or no consequence for the argument of this paper.

representation is too impoverished to capture the true nature of the relationship to be explained. The relevance of this message to the study of cognition depends on the observation that both the governing task and human behaviour essentially involve subtle interactions with constantly changing environments. According to van Gelder, it is on the basis of this parallel that essentially the same considerations brought to bear in the case of the centrifugal governor can be applied with equal force in the seemingly disparate domain of cognitive science.

Thus, when the dynamicist suggests that it may be possible to account for cognition without evoking the notion of representation, it will be because the concept of representation is too *simple* to be capable of describing the real relationships involved (p.19).²

So the traditionalist has nothing to fear unless van Gelder can establish that the centrifugal governor does not instantiate a representational solution to the governing problem. (If the centrifugal governor is, after all, a representational device, then its practical success does not pose a threat to orthodox cognitive science.³) To meet this demand, van Gelder presents us with an account of what a computational (and, hence, representational) governor would be like. The argument is then that the centrifugal governor is not, in any important theoretical sense, akin to this hypothetical device.

If you are in the business of writing software to handle complex problems, a useful principle is to take your overall task, and divide it up into its simpler sub-tasks. In this spirit, van Gelder supplies us with the following sequence of sub-tasks for our computational governor.

1. Measure the speed of the flywheel.
2. Compare the actual speed with the desired speed.
3. If there is a discrepancy then
 - a) measure the current steam pressure
 - b) calculate the desired alteration in steam pressure
 - c) calculate the necessary throttle valve adjustment
 - d) make adjustment
4. Go To 1.

On the basis of this breakdown, van Gelder concludes that the computational governor would require these component sub-systems:

1. A device to measure the speed of the flywheel.
2. A device to calculate the speed discrepancy.
3. A steam pressure meter.
4. A device for calculating the throttle valve adjustment.
5. A throttle valve adjuster.
6. An executive handling the sequencing of operations.

Now what are the differences between the two governors? Van Gelder's analysis identifies five areas of fundamental divergence. These can be summarized as follows (adopting van Gelder's usage of the various technical terms — this will be discussed later):

²All page numbers cited in sections 2 to 5 inclusive refer to (van Gelder, 1992) unless otherwise indicated.

³When I refer to 'traditionalists' or 'orthodox cognitive science', I do not mean to pick out merely those researchers advocating von Neumann style computing architectures in contrast to connectionist set-ups. For the purposes of this essay, the defining feature of traditionalism is adherence to the hypothesis that cognition is the construction and manipulation of representations by computational information processing mechanisms, and, as I stated at the start of this paper, such a commitment is part and parcel of most connectionism as well as the guiding principle of language-of-thought-style theories.

A) The computational governor functions by building, manipulating, and internally transferring representations. The centrifugal governor features no representations.

B) The computational governor respects the principle of homuncular decomposition. There are specialized sub-systems which solve particular sub-tasks. These sub-systems communicate with each other via the passing of representations which have some identifiable semantic relevance at the level of our description of the task as a whole (e.g., speed or pressure). For example, the homunculus responsible for measuring the speed of the flywheel sends a representation of that speed to the homunculus responsible for calculating the speed discrepancy. Although the centrifugal governor does submit to some form of decomposition, it cannot be analyzed into compartmentalized mini-agents which communicate with each other in the specified manner.

C) The computational governor clearly follows a measure-compute-act cycle as well as completing several distinct, sequential operations involving representations within the compute stage of the cycle. The centrifugal governor does not carry out internal computations; i.e., there are no temporally distinct stages at which the functions of measurement, computation and action are individually executed, and no sequential manipulation of representations.

D) In van Gelder's words, "the two kinds of governor differ fundamentally in their temporality, and the temporality of the centrifugal governor is essentially that of the engine itself" (p.9). To my mind, this is the most difficult to grasp of the five points of difference. Its significance is unpacked later in this paper.

E) The rigorous analysis of the centrifugal governor requires the weaponry of differential equations and dynamical systems theory. The operation of the computational governor is most appropriately specified by an organized hierarchy of distinct instructions (an algorithm). The idea is that a device should be described in the explanatory vocabulary that best fits its fundamental nature.

Given the way that van Gelder specifies the workings of the computational governor, points B and C (homuncularity and computation) follow from point A (representations) without strict need for supplementary arguments. If there are no representations in the centrifugal governor then, by van Gelder's notions of homuncularity and computation, neither of these terms can apply. (Van Gelder, himself, does make this observation.) Point E (appropriate conceptual framework) can be conceded by the traditionalist as it depends on (rather than helps to establish) the claim that the centrifugal governor cannot be analysed in representational or computational terms. Consequently, my primary concern will be with A (representations) and D (temporality), although there will obviously be some spillage in my treatment as all the points are deeply interconnected.

It seems uncontroversial that the computational governor is making full use of representations. Each sub-system individually produces or manipulates representations, and also both receives and/or outputs representations in communication with other sub-systems. Van Gelder's claim is that, by contrast, the centrifugal governor

...makes no use of representation at all. It adjusts the throttle valve on the basis of engine speed, but never *represents* that speed or the throttle valve adjustment (p.4).

Of course, van Gelder anticipates the obvious objection. Doesn't the arm angle represent the speed of the engine by virtue of the fact that two factors are reliably correlated, or, rather, with the correct emphasis, because there exists a *causal* correlation, the arm angle being an *effect* of the engine speed? But

this challenge fails. As van Gelder points out, the proposed criterion of causal correlation is insufficient as a characterization of the interactive subtleties of the relationship at issue: the arm angle controls the speed of the engine via the throttle valve; so the arm angle is simultaneously a cause and effect *in the very same interaction*. This prevents it from qualifying as a representation of the engine speed by *the criterion of one-way causal correlation*. In addition, van Gelder stresses that any such correlation between arm angle and engine speed exists only whilst the entire system is at equilibrium. When fluctuations in load or pressure occur, the neat correlation between the two variables no longer exists, yet the governor continues to function admirably; and, moreover, these periods of non-equilibrium are the most important. Therefore, the temporary correlation between arm angle and engine speed cannot play any *key* explanatory role.

If the complex interaction under investigation is beyond the explanatory reach of the representationalist, how should it be described? Van Gelder's answer comes from the language of dynamics. The relationship between the centrifugal governor and the engine (and, thus, between arm angle and engine speed) is one of dynamical coupling. Such a relation obtains when two separable dynamical systems are bound together in a mathematically describable way, such that, at any particular moment, the state of either system fixes the dynamics of the other system. Earlier, I defined a dynamical system as any system for which we can, in principle, provide a state space evolution equation describing how the values of the state variables of the system change with time. Other values in such an equation specify quantities which affect the behaviour of the system without being affected in turn; these are called the parameters of the system. In formal terms, the coupling of two dynamical systems means that some of the parameters of each system either become, or become functions of, some of the state variables of the other. In the current example, a parameter for the governor is effectively a state variable for the engine, and vice versa (see table i).

TABLE i.

DYNAMICAL COUPLING AND THE WATT GOVERNOR		
	Parameter	State Variable
Watt Governor	Engine Speed	Arm Angle
Steam Engine	Setting of the Throttle Valve: Directly dependent upon arm angle	Engine Speed

Given two coupled dynamical systems, X and Y, as X changes state, the dynamics of Y will change. This change in the dynamics of Y will, in turn, feed back into the dynamics of X, and so on. Effectively, what coupling means is that two systems evolve together through time in a continuous process of feedback and mutual interaction. Later in this paper, I shall specifically discuss the role of dynamical coupling in explaining the behaviour of cognizers.

3 The Internal State of Representationalism

Naturally, the staunch representationalist won't throw in the towel without a fight. In this section, I consider some possible criticisms of van Gelder's arguments.⁴

The representationalist might accept van Gelder's general picture of the theoretical nature of a computational governor, but claim that his (van Gelder's) specific characterization of such a device greatly over-complicates the amount of representing and computing necessary to achieve a traditional-style solution to the governing problem. In effect, van Gelder is held to be guilty of attacking a 'straw representationalist'. Perhaps it is true that van Gelder's computational governor is somewhat over-complicated (although I wonder how many traditionalists would have settled on a high level algorithm not all that different from van Gelder's). But, even if the complaint of over-elaboration is justified, van Gelder's central point remains intact. The structural organization of a *reduced* computational governor must still display the key properties at issue; it must still feature representations, computations, and homunculi, even if there are less of each of them, and even if those which are there are considerably simpler than those postulated in van Gelder's original conceptualization. Otherwise it is hard to see how it could even count as a computational governor. What the traditionalist needs to show is that the centrifugal governor is a representational/computational device. It is incumbent upon the traditionalist to find a plausible analysis (of the centrifugal governor) which cites representations and computations, and accounts for the interactive subtleties that van Gelder explains via the notion of dynamical coupling. The point is not that there could not be a computational solution to the governing problem, but that the centrifugal governor is not a computational system.⁵

Another style of argument, and one which van Gelder, himself, considers, would be to object that unnecessarily restrictive criteria are being imposed for what count as representations and computations. Doesn't the centrifugal governor instantiate an analog computation by virtue of the fact that it implements a particular mathematical function from engine speed to valve angles? The operation of the centrifugal governor could indeed be described as analog computation in this sense but, as van Gelder observes, he "never intended to deny that the centrifugal governor actually solves the governing problem" (p.7). It is not *that* sense of computation to which he takes exception.⁶

This highlights an important fact about the dynamical critique of orthodox cognitive science. The thoroughgoing dynamicist bets against the truth of any hypothesis which claims that the crucial aspects of cognition involve the processing of semantically interpretable, internal states which function to encode objective states of an external world (i.e., representations). This representational theory of the mind is easily linked to the thought that the mind is also fundamentally computational because, among other things, the latter idea seems to provide some hope of an explanation as to how a physical system, whose processing mechanisms have access to only the structural features of that system's representational states, could manipulate those states so as to respect semantic constraints; after all, that is what computers are so good at.⁷ Given this link, the relevant sense of computation is one which essentially binds the notions of computation and representation. As van Gelder puts it, "there is computation when there are semantically appropriate manipulations of representations" (p.8). The dynamicist holds that this portrait captures the theoretical commitments of traditional cognitive science (including most connectionism), and that when these commitments are followed through to their natural conclusions, the traditionalist is just stuck with

⁴During section 3, I make various moves on behalf of the computational/representational camp in response to van Gelder's arguments. Many of these moves are based on points made to me both in private correspondence, and in conversation, by Matthew Elton.

⁵A liberal-minded traditionalist could allow that look-up-tables qualify as computational systems, and then argue that the centrifugal governor implements, or is equivalent to, some look-up-table arrangement. At the very least, this proposed equivalence fails to hold because a look-up-table still implements a measure-compute-act cycle in which each of the functionally individuated steps is sequentially discrete. This is not true of the centrifugal governor.

⁶In any case, the *mere* fact of physically implementing a particular mathematical function must be far too weak a notion of computation for the traditionalist to use as her explanation of cognition.

⁷For the full story of the link, see (Sterelny, 1990), especially chapter 2.

something like the senses of representation and computation made explicit in van Gelder's paper. It is always open to the traditionalist to respond first by explaining that *those* senses are not what she meant at all, and then by providing 'accurate' explications of the two terms. The ball would then be back in the dynamicist's court.

At this stage, I feel compelled to make an observation which does not have the status of a philosophical claim, but which I find highly suggestive. Of late, I have been involved in many discussions in which the issue of what counts as a representation has been a central question. In these interchanges, it is my experience that, if a traditionalist rejects the notion of representation as identified by the dynamicist, it seems very difficult indeed for her to avoid sliding all the way into the impossibly weak position that 'possessing a representation' equals 'having internal state.' Some traditionalists feel the need to summon up representations when confronted by any behaviour at all which requires some form of synchronization between one system and another. For instance, there are plants which track the sun as it moves across the diurnal sky. But there is one species of plant whose behaviour is locked in to the general movements of the sun in such a way that, around an hour before the dawn, the plant swings into position so as to be in the correct orientation when the sun appears. Apparently, for some traditionalists, this 'anticipatory' behaviour means that the plant must have a representation of the sun. (Occasionally, the proposed, causally efficacious internal state is referred to as a proto-representation, whatever one of those is!) Unsurprisingly, the plant's adjustments can be fully accounted for without the notion of representation (Simons, 1992). This just goes to show how unconstrained (and, thus, potentially vacuous) the use of the term 'representation' can become; and, as far as this brand of traditionalist is concerned, it makes a mockery of any contention that the Watt Governor is too simple to qualify as a representational system.⁸

No one (except, perhaps, some Skinnerian zealot) would deny that, in the case of cognitive systems, there exists behaviour which arises independently of the system's immediate sensory input, or that that system's internal state is of explanatory interest with regard to such behaviour. So any worthwhile account of cognitive systems will include some notion of internal state in its explanatory vocabulary.⁹ Purely reactive behaviours, in which the sense-act transition is unmediated by any distinctive, behaviour-organizing internal states, cannot be the whole story. But the notions of representation and internal state cannot be equivalent: a system without internal states presumably cannot be said to have representations; but, surely, a system can have internal states without possessing representations. Consider Beer's example of the concentrations of reactants in an industrial fractionation column. These concentrations result in the column displaying time-dependent input-output behaviour which will affect the company's performance, but they *represent* nothing about the company (Beer, 1992).

The dynamicist can embrace a rich concept of internal state. A dynamical system follows a trajectory dependent upon its own current state, and the laws of state space evolution which specify its dynamics. As Beer observes, the "very fact that an agent is modelled as a dynamical system rather than a function implies that it has internal state" (Beer, 1992, p.13). Those internal states will have a critical explanatory interest for the scientist interested in explaining the entire range of cognitively-driven behaviour in dynamical terms. This counters a possible representationalist ploy in which it is conceded that there are certain behaviours for which it may well be correct to insist on a dynamical-systems-style analysis, but which attempts to restrict such explanations to behaviours such as 'knee-jerk' reactions which could, in principle, be explained by low-level sensory and motor capacities.

Cognitive science would benefit from a fully worked out theory of representation. At present, it seems reasonable to conclude, with van Gelder, that the traditionalist is committed to concepts of representation

⁸Inman Harvey brought the existence of the sun-tracking plant to my attention. The 'sun-tracking plants have representations too' claim has been made to me (in conversation) by both Ron Chrisley and Chris Thornton (individually and as a team).

⁹'Internal state' cannot be defined as merely a physical state which occurs within some crude physical boundary such as an organism's skin. For the dynamicist, internal states will be individuated by reference to the state variables and parameters relevant to some time-dependent behaviour of interest. In section 7, I introduce the notion of a context-dependent control-boundary which suggests one way in which a useful dividing line between internal and external may be drawn.

and computation which just cannot be applied to the centrifugal governor.

So van Gelder's claim, that the orthodox framework is inappropriate for analysing the interactive subtleties present in the example of the Watt governor, seems to be safe. The crucial step of transferring the same conclusion to the case of cognition rests on the insight that the behaviours of the different devices (governors and cognizers) are equivalent in the sense that both the governing task and cognitively driven human behaviour inherently involve subtle interactions with constantly changing environments. Van Gelder's polemic against the computational analysis of the Watt governor amounts to the charge that traditionalists cannot explain that system's situatedness (the way in which it is embedded in its environment). In this context, the relative simplicity of the governor works in the dynamicist's favour. Assuming that both of the situated behaviours have the same fundamental character, if the traditional style of explanation fails for the Watt governor, it's just bound to collapse in the face of the complexity displayed by cognitive systems and, therefore, will fail to explain their situatedness too.¹⁰

4 The Status of Time in a Cognitive Theory

Van Gelder claims that there is "an important sense in which time does not matter in the operation of the computational governor" (p.8). What does this mean? In this section, I shall endeavour to unpack van Gelder's temporality argument.

Obviously, there is a practical constraint on the computational governor; i.e., it must be successful in its functional niche. Consequently, whatever the necessary sub-tasks turn out to be, they must occur in the right order and happen sufficiently fast. This requires decisions regarding the choice of algorithms and hardware. But, according to van Gelder, these are pragmatic implementation-details, beyond which

...there is nothing which dictates *when* each internal operation takes place, *how long* it takes to carry out, and *how long* the temporal interval between each operation is...The timing of the internal operations is thus essentially *arbitrary* relative to that of any wider course of events (pp.8-9).

The idea is that, in the case of the centrifugal governor, significant temporality is embedded in the *theoretical* language most appropriate for the analysis of the device. Such theoretical constraints do not appear in the conceptual framework which best describes the fundamental character of the computational governor. It is in this sense that we are supposedly *beyond* considerations of mere implementation. This needs further elaboration.

From the general tone of this paper so far, the reader would be forgiven for thinking that the dynamicist about cognition holds that computational systems are not dynamical systems. But, in fact, the dynamicist says no such thing because computational systems are special cases of dynamical systems (Giunti, 1991). So why do we typically think of computational systems in terms of, for example, algorithms rather than state space evolution equations?

The traditional conceptual anchor for computational theory is the Turing machine. A Turing machine, in its most general form, consists of a tape of infinite length on which are cells, each containing one symbol from some finite alphabet. The tape is viewed by a reader which is, itself, always in some state. At each step, the reader looks at the cell over which it is stationed, and carries out some action determined by its own current state and the symbol in view. This action is limited to some combination of the following: the reader writes a symbol onto the tape in the cell over which it is stationed; it moves one cell to the left or right; it changes its own state. Particular Turing machines are described by means of tables which

¹⁰As far as physical systems are concerned, the character of their situatedness is ultimately dependent upon the ways in which those systems are embodied. For animals — including human beings — this will be cashed out primarily in terms of sensorimotor capacities. For the remainder of this paper, wherever the situatedness of physical systems is the issue, the dependence of situatedness upon embodiment should be assumed.

specify all the possible pairings of symbols and states, together with the accompanying action. Despite the apparent simplicity of the device, Turing was able to establish that any procedure that is computable can be computed by a Turing machine (Turing, 1936).

Usually, the state of a Turing machine is identified as the state of the reader. However, if we think of a Turing machine as a dynamical system, we need to take into account *all* facets of the system that change over time: i.e., not only the current reader-state, but the current contents of every cell on the tape. A tape of unbounded length therefore requires an infinite dimensional state space in order for us to describe, in dynamical terms, the way the system evolves through time. Any individual symbol-state-action square in a Turing machine table will specify the change in three state space variables (the symbol in the current cell in view, the reader-position, and the reader-value). The change in all the other variables is implicitly specified to be null at this time step. Thus it is possible to think of a Turing machine table as a close relation of a state space evolution equation. (This result of Giunti's is summarized by van Gelder pp.27-30.)

Following van Gelder (p.31), we can see that the operation of a Turing machine is effectively a series of discrete events which are temporally (although not sequentially) arbitrary. Certainly, state transitions take place from one time step to the next. But there is no *theoretical* sense in which

1. a state transition takes any time to complete,
2. there is a time interval between states, or
3. the states themselves persist for any length of time.

In addition, a Turing machine is digital and deterministic, and displays only localized state-space transitions and low inter-variable interdependency. (The latter two properties are just consequences of the fact that, at most, only three state space variables are changed in each discrete state transition.) Van Gelder argues that it is these various features of a Turing machine which make it both *possible* and *preferable* to think of it in computational and algorithmic terms. So it is precisely the reconceptualization of a Turing machine as a dynamical system that brings the character of that system's computational nature into focus. Of course, because a computational system *is* a dynamical system, the hypothesis that cognition is state-space evolution in *some sort* of dynamical system is accepted by the orthodox and the rebellious alike, and the arguments over the correct theoretical framework for cognitive science can be restated as being over which sub-classes of dynamical systems are also cognitive systems.

We are now in a position to sharpen up our picture of what is going on in van Gelder's temporality argument. As defined earlier, a dynamical system is any system for which we have a rigorous analysis of the way that that system evolves through time. The conceptual anchor of computational theory (and, therefore, of the computational approach to cognition) is the Turing machine. A Turing machine is a dynamical system, but one which, from a dynamical systems perspective, has some peculiar characteristics which mean that it is most appropriate to describe it in the language of computations and algorithms. One of these properties is arbitrary temporality. And if the system which provides the theoretical foundation of computationalism has this quality, it seems fair to claim that temporality is not theoretically significant for computational systems in general. For computational systems, temporal constraints arise only as implementation considerations. So this must be the case for the computational governor.¹¹

By contrast, when we look at the relationship between temporality constraints and the centrifugal governor, we find a different story. There is no sequence of discrete, temporally distinct operations, and the governor's corrective behaviour is inextricably part of the very same real time frame as the change

¹¹It could be argued that 'real-time' computing systems employed in control-engineering constitute counter-examples to this claim. In such systems, highly complicated strategies are employed in the algorithms to take account of temporal constraints. However, the enormous increase in algorithmic complexity demonstrates just how 'unhappily' fundamental temporality fits with digital computation, whereas it is the essence of dynamical systems theory. I am indebted to Phil Husbands for bringing this point to my attention.

in the speed of the flywheel caused by some fluctuation in load or pressure. Such observations support the hypothesis that dynamical systems theory is the most appropriate vocabulary in which to analyse the behaviour of the centrifugal governor. In dynamical coupling, remember, the linked systems evolve *together* through time. Hence van Gelder's conclusion that "the temporality of the centrifugal governor is that of the engine itself" (p.9).

Van Gelder is surely right both that cognitive science needs to rethink its attitude towards temporality, and that the likely result will favour the dynamicist. No cognitive agent can be understood without some notion of real-time interaction with an environment, and dynamics, as the science of change, has intrinsic temporality embedded in its theoretical language. This insight can be developed further. Arbitrary temporality is symptomatic of a process of abstraction away from situatedness and embodiment, a process which has been a defining characteristic of most cognitive science (cf. the criticisms levelled at traditional A.I. by Brooks (1991a)).

The enabling assumption of traditional cognitive science is that the admittedly messy and complicated business of achieving real-time interaction with an environment is an implementation headache to be overcome by the hardware department. On this view, a cognitive agent is, first and foremost, a disembodied intellect which is then hooked up to the world (shades of old Descartes). The dynamicist suspects that this attitude woefully underplays the importance of situatedness and embodiment, and, thus, temporality. So this worry won't just disappear for the traditionalist who maintains the same conceptual prejudices, but becomes more implementation-minded.¹²

5 Towards A New Cognitive Science

Van Gelder asks us to imagine a dynamical system which behaves just the opposite to a Turing machine: i.e., a system which exhibits

1. continuous operation in real time and space,
2. stochastic state transitions,
3. dynamics such that change in any one variable depends on the current state of a great number of other variables in the system, and
4. the possibility of very large rates of change in any given state.

Using the criteria outlined above, we can conclude that such a dynamical system is a non-computational system. It is also just the kind of dynamical system realized by biological neural networks and, of course, neurobiologically speaking, that is what the most complex cognizers we know of, i.e., human beings, are (p.33). Keeping this thought in mind, we are ready to introduce van Gelder's Dynamical Cognition Hypothesis which states that

1. cognitive systems are dynamical systems found relatively close in the space of possible systems to the Watt Governor;
2. cognition is state space evolution in such systems;
3. the most appropriate framework for the study of cognition is dynamical modelling and dynamical systems theory.

¹²It seems to me no accident that a philosopher who took situatedness very seriously indeed concentrated much of his effort on exploring the conceptual links between world-embeddedness and temporality. I refer, of course, to Heidegger (1926). I admit that, at present, the mention of this potentially powerful link between the type of dynamicism I wish to develop and Heideggerian philosophy smacks of hand-waving and philosophical name-dropping. That is why I have relegated it to a speculative note. Defending and expanding on the rapprochement is a job for another day.

The first sub-hypothesis requires some clarification. On van Gelder's picture, the Watt governor usurps the Turing Machine as the conceptual anchor for cognitive science. It does so because it fixes the class of non-computational dynamical systems, and *not* because the level of dynamical complexity is equivalent for both governors and cognizers. Van Gelder is perfectly clear on this point (pp.36-37). The dynamical properties of cognitive systems are, of course, far more complex than those of the centrifugal governor. (In section 7 of this paper, I shall explore the nature of cognitive dynamics, arguing, along the way, that the increase in dynamical complexity buys the cognizer a certain sort of independence from her environment. The centrifugal governor does not enjoy such independence.)

The traditionalist is not finished yet. "Look," she says, "I grant you dynamicists that the nervous system *is* (probably) a dynamical system of the sort you describe. If we are talking neurophysiology you have a point. At *that* level of description, you are most likely correct; human beings are dynamical systems. But that doesn't prove to me that *cognition* is best described as state space evolution in dynamical systems."

How might the dynamicist respond? Two possible responses take the form of appeals to current empirical research in cognitive science. However, despite the tactical similarity, the two appeals should be kept apart. The first cites a growing body of empirical research which is both concerned with explanation at an obviously cognitive level, and employs dynamics as its fundamental conceptual terminology. This includes research into decision making (Townsend & Busemeyer, 1989), development (Thelen, 1990), and motor control (Turvey, 1990). In addition, *some* connectionists who claim to be doing cognitive modelling use concepts from dynamical systems theory in order to characterize and explain the behaviour of their networks, e.g., (Smolensky, 1986), (McAuley *et al.*, 1992). Smolensky has also provided a theoretical treatment of connectionism which is steeped in the language of dynamical systems theory (Smolensky, 1988).¹³ Surely this continuing research shows that concepts from dynamical systems theory are already being used to formulate explanations of cognitive capacities.

Van Gelder does not explicitly consider the objection to the dynamical cognition hypothesis with which we are currently concerned, but, in the course of his paper, he does cite (with approval) all of the key empirical work mentioned in the last paragraph. Perhaps, therefore, citing the success of this work in providing powerful explanatory models of distinctively cognitive abilities would be part of his response. But some care is needed here. Replacing talk of representations and computations with talk of state variables, evolution equations, and coupling is only a pigeon-step in the right direction if it is unaccompanied by at least one other realization, namely that non-situated approaches in cognitive science, which concentrate on abstracted sub-domains of cognition, will always run the risk of making too many assumptions about the way in which the particular capacity of immediate interest is integrated with perception and action. (This is another form of the criticism of orthodox cognitive science made at the end of the last section, i.e., that it mistakenly abstracts away from the key issues raised by situatedness and embodiment.) In this respect, why should abstract dynamical models fare any better than abstract representational models? In the case of the dynamical connectionists at least, this problem still has to be overcome.

Consequently, I prefer to focus on a different area of current research; and it is this which supplies the ammunition for our second response to the traditionalist. In this paper, I have indicated that it is the nature of situatedness which, in the end, provides the most powerful arguments for dynamicism. If this is right, then the most rewarding progress for the dynamical systems movement will be in the study of world-embedded agents.

¹³Unfortunately, many commentators on Smolensky's paper have seemingly been oblivious to this fact, concentrating, instead, on issues such as the symbolic/subsymbolic distinction.

6 The Dynamics of Situatedness

Autonomous agents are fully integrated, self-controlling, active systems which, while in continuous long-term interaction with their environments, behave so as to achieve certain goals. Physical autonomous agents — such as people, animals, insects and autonomous robots — are necessarily embodied, and inextricably situated. Cognitive architectures can be thought of as the control systems for sufficiently complex situated agents (cf. (Harvey *et al.*, 1993)). There now exists a breed of autonomous agent researcher which refuses to take refuge in disembodied, isolated sub-domains of cognition. Instead the hunt is on for robots which exhibit adaptive behaviour (“behaviour which is selected to increase the chances that a situated agent can survive in an environment which is noisy, dynamic, hostile, and uncertain” (Cliff *et al.*, 1993)). The direct result of this dedication to the building of complete agents is

a methodological commitment to a biologically informed, simple-systems-first policy.

(Dave Cliff suggests that ‘A.I.’ ought to stand for ‘artificial insects’ (Cliff, 1990).)

If we focus on adaptive behaviour, as opposed to chess-playing, we are no longer blind to a suggestive realization. The capacity for situated activity exists throughout the animal world. On evolutionary grounds, it seems reasonable to hypothesize that linguistic competence and deliberative thought are overlays on the prior (and, in terms of survival, more fundamental) capacity for situated activity. In fact, it is entirely plausible that human-level, cognition-driven behaviour is ‘merely’ a form of adaptive behaviour in general. So perhaps it is with the study of relatively simple situated agents, rather than fancy heuristics designed to prune exponentially expanding search trees, that the search for an understanding of cognition should begin.

In this vein, Beer, who uses ‘computation’ and ‘representation’ with practically the same senses as employed by van Gelder, argues that it *may* be correct to hold out for a computational theory of deliberative reasoning (Beer, 1992). After all, he observes, our very notion of computation was abstracted from our own conscious, deliberative reasoning, and it seems that, on occasions, humans do ‘compute’, in the strict sense that we deliberately form and manipulate symbolic representations (e.g., when we do mathematics or formal logic). However, there is evidence to suggest that the notions of representation and computation are inappropriate to characterize the more fundamental situated activity of autonomous agents. The relevant evidence comes from recent research into situated agents which demonstrates, empirically, that relatively complex and robust behaviour can be generated by autonomous systems with control architectures which fail to reflect the traditional principles, e.g., (Beer, 1990), (Brooks, 1991b), (Harvey *et al.*, 1993), and (Maes, 1990). For example, those working within what has been dubbed ‘behaviour-based robotics’ (e.g., Brooks and Maes) reject the sense-model-plan-act framework embraced by the traditionalist (cf. the measure-compute-act cycle of the computational governor). These roboticists tend to advocate highly reactive architectures, with no central reasoning systems, no manipulable, symbolic representations, and radically decentralized processing.

Leaving behind the hypothesis that cognition is literally, in all essential respects, computation (in the relevant sense of ‘computation’) does not mean that computers have no role to play in the study of situated agents. The potential exists for the computer-modelling of agents and their environments. But, as Beer reminds us, “no computational fluid dynamicist would mistake a computer model of a fluid for the theoretical claim that fluids compute their behaviour” (Beer, 1992, p.2).¹⁴

What comes out of this line of argument is a demand for a theoretical vocabulary which does not assume computation and representation as primitive notions. This leads Beer to suggest that the most suitable theoretical framework for the study of autonomous agents is dynamical systems theory, a suggestion also to be found in (Smithers, 1992), (Cliff *et al.*, 1993), and (Husbands *et al.*, 1993). So a revolutionary shift in

¹⁴A simulated autonomous agent in an artificial environment would be situated but not embodied (Brooks, 1992). This raises important issues regarding the danger of working with simulations. Simulations allow the researcher to define, and, potentially, oversimplify the sensorimotor relationship between agent and world. See (Harvey *et al.*, 1993) for a methodology designed to keep this worry at bay.

the foundations of cognitive science may be in the air. But something has been missing from my discussion. Aside from a few remarks, details of the new path to the cognitive holy grail have been conspicuous by their absence. The traditionalist is owed at least a sketch of the dynamical alternative. It is time to discharge the debt.

7 Where We Go From Here

Along with van Gelder and Beer, I wish to conceptualize an agent and its environment as two coupled dynamical systems. How might we flesh this out? Consider Beer's dynamical framework for autonomous agents (Beer, 1992). A (the agent-system) and E (the environment-system) are coupled via two functions, S (environmental state variables to agent parameters) and M (agent state variables to environmental parameters). As an approximation, S can be thought of as sensory input, and M as motor behaviour, although Beer is clear that the scopes of these two functions are not necessarily restricted to what are normally thought of as sensory and actuation channels. They are supposed to capture *any* effects which one of the systems can have on the future trajectory (state space evolution) of the other. In the rest of this section, I shall investigate the character and consequences of this picture of autonomous agents and their environments, with particular emphasis on those most complex of autonomous agents, cognizers.

Earlier in this paper, I characterized coupling as a relationship in which two systems evolve together through time in a continuous process of feedback and mutual interaction. I propose that dynamical coupling is the fundamental mechanism of situatedness. As a result, agents and their environments become, in some sense, inseparable. In fact, it should be said that they stand in a relation of cospecification (cf. (Varela *et al.*, 1991)). This view is supported by Beer's observation that whilst it is useful, under certain circumstances, to think of the agent and its environment as separate but coupled dynamical systems, it is equally valid to redescribe the coupled agent-environment system as one larger dynamical system in which the observed patterns of interaction between the agent-system and the environment-system are properties of that larger system (Beer, 1992). So the decision to draw the line between agent and environment is inherently revisable. Furthermore, if a division is imposed, where the dividing line is drawn is, itself, a matter of choice. To my mind, this fits nicely with Gibson's point that a tool-in-use is plausibly better thought of as part of the agent-system rather than the environment-system (Gibson, 1979).

All of this should not be taken to imply that the dynamicist does not recognize any sense in which agents are 'independent' of their environment. In the following passage, Beer explains the nature of the interactions between the coupled agent and environment systems.

S [the environment-to-agent function] cannot necessarily be interpreted as "instructing" an agent as to the configuration of its environment. Generally speaking, one dynamical system cannot directly specify the state of another dynamical system. A dynamical system follows a trajectory specified by its own current state and dynamical laws [state space evolution equation]. By varying some of the parameters of those laws, a second dynamical system can certainly influence the subsequent trajectory of that system. However, influence is not the same as specify. In general, one dynamical system can at best bias the intrinsic "tendencies" of a second dynamical system to which it is coupled. Thus, it is perhaps most accurate to view an agent and its environment as mutual sources of perturbation, with each system continuously influencing the other's potential for subsequent interaction. (Beer, 1992, p.13)

Now, take two coupled dynamical systems, and consider, for a moment, a simplified abstraction in which the coupling is 'one way', i.e., we have a driving system and a driven system. Certain changes in the state of the driving system will perturb the dynamics of the driven system. If, for any small perturbation in the dynamics of the driven system, those dynamics remain *qualitatively* the same, then the driven system is said to be structurally stable. (This will be with respect to the parameter or parameters of interest.

¹⁵) Structural instability ensues when a small perturbation results in a qualitative change in the total dynamics of the system. By their very nature, structurally unstable dynamics will tend to be ephemeral. As observers, we may choose to distinguish any number of dynamical cognitive systems and sub-systems at varying levels of abstraction, depending upon the properties (state variables and parameters) we decide to be of current interest. On grounds of overall behavioural coherence, it seems likely that the vast majority of these systems will tend to exhibit structural stability with respect to most of their changing parameters. Now and then, of course, it may be very useful for the dynamics of a cognitive agent to undergo fundamental change (e.g., in the face of immediate physical danger). In cognitive dynamics, occasional structural instability might not always be a bad thing.

To develop our portrait of dynamical cognition, recall that, in section 3, I argued that the dynamicist can recognize a rich sense of internal state, which allows an agent to engage in non-reactive behaviour (behaviour not dependent upon immediate sensory circumstances). In some cases, it is no doubt a mistake to attribute large amounts of internal complexity to an agent when, in fact, the observed complexity of behaviour is more a matter of environmental interaction (cf. Simon's ant (Simon, 1969)). This important appeal for parsimony notwithstanding, there seems little doubt that the sophisticated behaviour demonstrated by cognitive agents in interaction with their environments requires enormous complexity in the internal dynamics of the cognitive agent itself. To glimpse how rich the dynamicist's conception of cognitive dynamics can be, just adopt the sort of changed perspective hinted at in the last paragraph, and think of a cognitive agent as being made up of many coupled, dynamical sub-systems, interacting and evolving together through time. Given the point about arbitrary divisions, it may seem that we need a specification of what 'internal' is here being taken to mean; in this context, the organismic boundary will do. In general, however, such a crude dividing line will be inadequate. My own inclination is to favour distinctions in terms of some context-dependent control-boundary. As an example of such a boundary, consider the following example due to Dawkins (1982). One parasite of the fresh water shrimp has a particularly adventurous life cycle in which it not only invades a shrimp, but passes through a muskrat as well! After the parasite has carried out its shrimp-invasion, the shrimp switches from a mud-dwelling to a surface-dwelling existence, thus increasing the likelihood that it will be eaten by a muskrat. So the parasite achieves its muskrat-period by acting directly on the shrimp's nervous system, thereby extending its (the parasite's) control-boundary to include the shrimp's organismic boundary. ¹⁶

We could decide to think of, say, the agent as an uncoupled system (i.e., using Beer's scheme, none of the parameters of A are changing via S, and none of its state variable changes are reflected in changes in E via M). In principle, there will still be state space evolution of A through time as a result of purely internal dynamics. ¹⁷ But, of course, in general, the dynamics of two coupled dynamical systems will be vastly different from the dynamics of those systems when uncoupled. For this reason, no matter how complex the purely internal dynamics of an agent may be, its situated behaviour must be seen as a consequence of the interaction between it and its environment, and not explained in terms of a causal effect of its internal state alone. Properly speaking, an agent's behaviour is a property of the single

Despite this notion of behaviour, there will still be an important role for the sort of idealization that focuses on a dynamical model of an agent for some fixed values of the environmentally determined parameters. As we have seen, in coupling, the entire dynamics of a driven system can change from one moment to the next as a result of changes in the driving system. (In general, of course, both systems are driving and driven.) Naturally, this makes the empirical investigation of interesting coupled systems a difficult business. One way to approach the problem is to extract a number of temporally distinct, instantaneous

¹⁵Formally, qualitative similarity of dynamical behaviour is cashed out in terms of the topological equivalence of phase portraits. Whilst this concept of topological equivalence is not difficult to grasp, an explication would take several pages of background information and terminological explanations. In any case, an intuitive sense of qualitative similarity is both adequate for an appreciation of the point I wish to make, and faithful to the formal notion. Those interested in the full story, see (Abraham & Shaw, 1992).

¹⁶The notion of a control-boundary is also applicable to the Gibsonian point about tool-use which I mentioned earlier.

¹⁷This seems to be at least part of what Varela et al. mean by the 'operational closure' of a system (Varela *et al.*, 1991).

time slices of the system under investigation (change in parameters = 0), and gradually piece together an overall picture of the way change in the system is changing (Beer, 1992).¹⁸

So I'm prepared to bet that cognitive agents viewed as dynamical systems display a strong tendency towards structural stability, and exhibit highly complex internal dynamics. On the basis of these hypotheses, and accepting Beer's picture of the way in which two coupled systems affect one another's future trajectories, it seems that 'being coupled' does not mean 'being at the mercy of every environmental change'. However, this recognition is perfectly consistent with the view that the behaviour of a cognitive system cannot, in any final sense, be explained in isolation from the dynamics of the environment in which that agent is situated. Behaviour is a feature of a system in which an environmentally-embedded agent and an agent-embedding environment evolve together through time, in a process of mutual and continuous feedback.

It is time to remind ourselves of the charge against which the last two sections of this paper have been a defence. In section 5, our hypothetical traditionalist accused the dynamicist of doing a fine job of theoretical neuroscience, but missing cognition altogether. Some traditionalists might argue that my preferred defence, based, as it is, on a research programme which currently studies relatively simple systems, supports, rather than disposes of, that criticism. This would be justified if my defence of dynamicism rested on the levels of behavioural complexity demonstrated at present by robots (real or simulated) which have been developed by the adaptive behaviour community. But this is not the case. Here's a reconstructed summary of my argument:

1. Real cognitive agents are necessarily both situated and embodied. (In the case of animals, 'embodiment' is understood as referring to sensorimotor capacities.)
2. Cognitive architectures are the control systems for the situated activity of sufficiently complex embodied agents.
3. Thus, cognition cannot be studied effectively in isolation from situatedness and embodiment. (Orthodox cognitive science fails to recognize this fact.)
4. Recent empirical research into autonomous agents concentrates on situated activity, and, therefore, recognizes the fundamental importance of situatedness and embodiment. Situated activity is ubiquitous in the animal kingdom. So the way forward for cognitive science is to begin by developing and studying the control architectures required by simpler, but complete, systems in adaptive interaction with their environments. We may expect that the methodologies and analytical techniques developed at the simpler level will suggest profitable ways in which to approach more complex cases.
5. Dynamical systems theory provides the most suitable theoretical framework for autonomous agent research.
6. Dynamical systems theory provides the most suitable theoretical framework for cognitive science.

Thus, the fact that current autonomous agent research is concerned with relatively simple systems is seen as a necessary step along the road to understanding cognition.

In Beer's presentation of the dynamical perspective, the possibility remains that *some* aspects of cognition will be most profitably described as representational and computational ((Beer, 1992), and cf. (van Gelder, 1992)). This respects Giunti's result (described above) that computational systems are special cases of dynamical systems. How many of our cognitive capacities are best explained using the traditional

¹⁸In his example of the synthesis and analysis of a simulated walking agent, Beer takes a predefined network architecture, and then employs a genetic algorithm (Goldberg, 1989) to find a set of network parameters (and, thus, a network dynamics) which enables the system to walk. He also presents a dynamical analysis of the notion of adaptive fit, a concept which plays a key role in genetic-algorithm-theory and adaptationist evolutionary theory in general.

framework is ultimately a matter for empirical research to decide, and Beer is often rather measured in his attacks on traditionalism. But I, for one, bet in favour of the view that traditional styles of explanation will become increasingly marginalized as the field of adaptive behaviour reaches maturity.

8 Summary

In this paper, I have defended the emerging dynamical approach to cognitive science. The space of possible dynamical systems is immense, and representational/computational systems fill one tiny corner of that space. There is mounting evidence that explanations couched in the vocabulary of representations and computations are far too restrictive to account for the behaviour of cognitive systems. The most compelling arguments for a more general, dynamical perspective come from the study of the adaptive behaviour of situated, autonomous agents. Once situatedness and embodiment are firmly on the agenda (as opposed to being tucked away as the sort of ‘any other business’ which always gets left until the next meeting), dynamical systems theory looks to be the most promising theoretical language in which to conduct cognitive science.

Acknowledgements

This work was supported by British Academy award no. 92/1693. Many thanks to Maggie Boden, Dave Cliff, Matthew Elton, Inman Harvey, Phil Husbands, Geoff Miller, and Chris Thornton for their helpful comments on an earlier version of this paper.

References

- Abraham, R. H., & Shaw, C. D. 1992. *Dynamics - The Geometry of Behaviour 2nd edition*. Redwood City, California: Addison-Wesley.
- Baker, G. J., & Gollub, J. P. 1990. *Chaotic Dynamics - An Introduction*. Cambridge: Cambridge University Press.
- Beer, R. 1990. *Intelligence as Adaptive Behaviour: An Experiment in Computational Neuroethology*. San Diego, California: Academic Press.
- Beer, R. 1992. *A Dynamical Systems Perspective on Autonomous Agents*. Tech. rept. 92-11. Case Western Reserve University, Cleveland, Ohio.
- Brooks, R. 1991a. Intelligence Without Reason. *Pages 569–95 of: Proceedings of the Twelfth International Joint Conference on Artificial Intelligence*. San Mateo, California: Morgan Kaufman.
- Brooks, R. 1991b. Intelligence Without Representation. *Artificial Intelligence*, **47**, 139–59.
- Brooks, R. 1992. Artificial Life and Real Robots. *Pages 3–10 of: Varela, F. J., & Bourgine, P. (eds), Toward a Practice of Autonomous Agents: Proceedings of the first European Conference on Artificial Life*. Cambridge, Massachusetts: MIT Press / A Bradford Book.
- Cliff, D. 1990 (May). *Computational Neuroethology: A Provisional Manifesto*. Cognitive Science Research Paper CSRP162. University Of Sussex.
- Cliff, D., Husbands, P., & Harvey, I. 1993. Analysis of Evolved Sensory Motor Controllers. *Pages 192–204 of: Proceedings of the 2nd European Conference on Artificial Life*.
- Dawkins, R. 1982. *The Extended Phenotype*. Oxford: Oxford University Press.

- Gibson, J. J. 1979. *The Ecological Approach to Visual Perception*. Boston: Houghton Mifflin.
- Giunti, M. 1991. *Computers, Dynamical Systems, Phenomena , and the Mind*. Ph.D. thesis, Department of History and Philosophy of Science, Indiana University.
- Goldberg, D. E. 1989. *Genetic Algorithms in Search, Optimization and Machine Learning*. Reading, Massachusetts: Addison-Wesley.
- Harvey, I., Husbands, P., & Cliff, D. 1993. Issues in Evolutionary Robotics. *Pages 364–73 of: Meyer, J. A., Roitblat, H. L., & Wilson, S. (eds), From Animals to Animats 2: Proc. second International Conference on simulation of Adaptive Behaviour*. Cambridge Massachusetts and London, England: MIT Press / A Bradford Book.
- Heidegger, M. 1926. *Sein Und Zeit (English translation by J. Macquarrie and E. Robinson. 1962. Being and Time)*. Oxford, England: Basil Blackwell.
- Husbands, P., Harvey, I., & Cliff, D. 1993. *Circle In The Round: State Space Attractors for Evolved Sighted Robots*. Unpublished Manuscript ,University of Sussex.
- Maes, P. (ed). 1990. *Designing Autonomous Agents: Theory and Practice from Biology to Engineering and Back*. Cambridge, Massachusetts: MIT Press.
- McAuley, J. D., Anderson, S., & Port, R. F. 1992. Sensory Discrimination in a Short-Term Dynamic Memory. *In: Proceedings of the 1992 Cognitive Science Society*.
- Simon, H. A. 1969. *The Sciences of the Artificial*. Cambridge, Massachusetts: MIT Press.
- Simons, P. 1992. *The Action Plant*. Cambridge, Massachusetts, and Oxford, England: Blackwell Books.
- Smithers, T. 1992. Taking Eliminative Materialism Seriously: a methodology for autonomous systems research. *Pages 31–40 of: Varela, F. J., & Bourgine, P. (eds), Toward a Practice of Autonomous Agents: Proc. 1st European Conference on Artificial Life*. Cambridge, Massachusetts: MIT Press / A Bradford Book.
- Smolensky, P. 1986. Information Processing in Dynamical Systems: Foundations of Harmony Theory. *Pages 194–281 of: Rumelhart, D. E., & McClelland, J. L. (eds), Parallel Distributed Processing: Explorations in the Microstructure of Cognition*, vol. 1. Cambridge, Massachusetts: MIT Press / A Bradford Book.
- Smolensky, P. 1988. On the Proper Treatment of Connectionism. *Behavioral and Brain Sciences*, **11**, 1–74.
- Sterelny, K. 1990. *The Representational Theory of Mind: an introduction*. Oxford, England and Cambridge, Massachusetts: Basil Blackwell.
- Thelen, E. 1990. Dynamical Systems and the Generation of Individual Differences. *Pages 19–43 of: Colombo, J., & Fagen, J. W. (eds), Individual Differences in Infancy: Reliability, Stability and Prediction*. Hillsdale, New Jersey: L. Erlbaum Associates.
- Townsend, J. T., & Busemeyer, J. R. 1989. Approach-Avoidance: Return to Dynamic Decision Behaviour. *In: Izawa, H. (ed), Current Issues in Cognitive Processes*. L. Erlbaum Associates.
- Turing, A. M. 1936. On Computable Numbers, with an Application to the Entscheidungsproblem. *Pages 230–65 of: Proceedings London Math. Society*.
- Turvey, M. T. 1990. Coordination. *American Psychologist*, **45**, 938–53.

- van Gelder, T. 1992. *What Might Cognition be if not Computation?* Tech. rept. 75. Indiana University Cognitive Sciences.
- Varela, F. J., Thompson, E., & Rosch, E. 1991. *The Embodied Mind: Cognitive Science and Human Experience*. Cambridge, Massachusetts, and London, England: MIT Press.