Appears in *Philosophy of Science*, 2000

Anti-Representationalism and the Dynamical Stance

Anthony Chemero
Scientific and Philosophical Studies of Mind Program
Franklin and Marshall College
Lancaster, PA 17604-3003, USA
tony.chemero@acad.fandm.edu

Abstract: Arguments in favor of anti-representationalism in cognitive science often suffer from a lack of attention to detail. The purpose of this paper is to fill in the gaps in these arguments, and in so doing show that at least one form of anti-representationalism is potentially viable. After giving a teleological definition of representation and applying it to a few models that have inspired anti-representationalist claims, I argue that anti-representationalism must be divided into two distinct theses, one ontological, one epistemological. Given the assumptions that define the debate, I give reason to think that the ontological thesis is false. I then argue that the epistemological thesis might, in the end, turn out to be true, despite a potentially serious difficulty. Along the way, there will be a brief detour to discuss a controversy from early twentieth century physics.

Appears in *Philosophy of Science*, 2000

0. Introduction

Classical cognitive science (Fodor and Pylyshyn 1988) is founded on the idea that mind is a digital computer and that thinking is computation. Since computation is usually understood as the rule-governed manipulation of representations (Haugeland 1985), this foundational idea requires the assumption that the mind contains representations of aspects of the environment. Yet, due to advances in modeling within cognitive science, many have been denying this necessary condition. That is, many have embraced anti-representationalism--very roughly, the idea that cognition does not involve representations at all. A well-known example of such an advance is the insect-like robots built by Rodney Brooks (1991, 1999). These artificial insects get about using what Brooks calls a *subsumption architecture*. A behavior control module is designed to impart a particular skill (say, obstacle avoidance) when utilized in a legged robot. That module is then re-deployed in a robot with several other modules each of which bestows a complementary skill as part of an overall control system (e.g. the avoidance module is wired up to modules for wandering and exploring). The "higher" modules are able to suppress the activities of "lower ones" (*explore* can suppress *wander*, and so on), but this is the only communication that goes on between the modules; and none of the modules builds or manipulates representations of the environment. With these simple modules appropriately connected, intelligent situated activity emerges, apparently without the building or maintenance of representations of the environment. To whatever extent real animals are like his robots, Brooks argues, they also do not build and employ representations.

Many other cognitive scientists, influenced by Brooks, have developed similar models of intelligent behavior, and employed a similar argumentative strategy in favor of anti-representationalism.[1] These arguments tend to follow the following schema:

> Here is a model of some cognitive phenomenon. There are no representations in this model. If cognition in general really works like this model does, then there are no representations in cognition either.

If these arguments are sound, cognitive science must be re-built from the ground up, re-conceptualized without reference to representation or computation. But unfortunately, it is difficult to know how seriously to take such arguments. Although, anti-representationalists usually carefully characterize the cognitive models in question, they rarely say in any detail what *exactly* they take representations to be (van Gelder

---

[1] Some examples: Skarda and Freeman 1987; Brooks 1991; Thelen and Smith 1993; van Gelder 1995; van Gelder and Port 1995; Beer 1995; Wheeler 1996; Hendriks-Jansen 1996; Ramsey 1997; Keijzer 1997.

1995 and Wheeler 1996 give rough ideas); they are also often not clear what *exactly* anti-representationalism amounts to. The purpose here will be to be more clear. In so doing, I will show that at least one form of anti-representationalism is potentially viable.

Here is how I will proceed. In section 1, I give a teleological definition of representation. In section 2, I distinguish two versions of anti-representationalism, one ontological, one epistemological. In section 3, I suggest that the ontological thesis is almost certainly false. This means that anti-representationalists should hold only the epistemological thesis. In the remainder of the essay, I argue that the epistemological thesis *might* turn out to be true. To do so, I draw parallels between debates in cognitive science over representation and the debates between Mach and Boltzmann over the status of (what we would now call) realist and instrumentalist interpretations of atomism in physics.

1. Defining Representation

I will start by providing a definition of representation, with which to evaluate purportedly anti-representationalist models. The definition will be *more restrictive* than those offered by van Gelder 1995 and Wheeler 1996. I will be rebutting their anti-representationalist arguments specifically, so it would be unfair to have a definition of representation which called more things representations than their definitions do. If that were the case, it could be said that these systems seem representational only because my definition lets too many things be representations. The definition here is more restrictive is that it adds a teleological requirement that is absent in both van Gelder and Wheeler. (See van Gelder 1995 and Wheeler 1996).[2]

I will take it that to show that something is a representation, it will suffice to show that it meets the following three conditions.

A feature R_0 of a system S will be counted as a *Representation for S* if and only if:

(R1) R_0 stands between a representation producer P and a representation consumer C that have been standardized to fit one another.

(R2) R_0 has as its proper function to adapt the representation consumer C to some aspect A_0 of the environment, in particular by leading S to behave appropriately with respect to A_0, even when A_0 is not the case.

(R3) There are (in addition to R_0) transformations of R_0, $R_1 \ldots R_n$, that have as their function to adapt the representation consumer C to corresponding

---

[2] There are definitions of representation more restrictive than the one I outline here. See for example Grush 1997. For a comparison of definitions of representation, see Chemero and Eck 1999.

transformations of A_0, $A_1...A_n$.

This definition is a version of Ruth Millikan's teleological theory of content (Millikan 1984 see also Bechtel 1998). Though, I will not defend this definition here (see Chemero 1998, Chemero and Eck 1999, Millikan 1984), I will very briefly point out a few of its features. First, as mentioned above, since it requires a representation to have functions, it is *teleological* (R2). Second, it requires that the representation serve as a representation in the context of producing and consuming devices (R1). Third, it requires that a representation be part of a system of representations (R3). Agents must be able to represent more than one thing, else they should not be thought of as representing anything at all Fourth, it requires that we follow Millikan (1984) in focusing on the representation consumer in determining the content of a representation--the content is the way the world would need to be for the behavior caused by the representation consumer to be adaptive (R2).

An important subset of the things that meet the criteria of this theory of representation are what Andy Clark (1997) calls *action-oriented representations*.[3] Action-oriented representations (AORs) are representations that *both* describe a situation *and* suggest an appropriate reaction to it; as Clark puts it, they are maps that are also controllers. AORs are more primitive than other representations in that they can lead to effective behavior without requiring separate representations of the state of the world and the cognitive system's goals. That is, the perceptual systems of agent need not build an action-neutral representation of the world, which can then be used by the action-producing parts of the agent to guide behavior; instead, the agent produces representations that are geared toward the actions it performs from the beginning . Gibsonian affordances (Gibson 1979) are examples of AORs; a rabbit simply sees something like "opportunity-for-cover-over-there" behind a tree and need not represent each of "there is a tree," "one can hide behind trees," and "I want to hide now".[4]

2. Defining Anti-Representationalism

In the most well-known and philsophically sophisticated essay recommending anti-representationalism, van Gelder (1995) suggests that the Watt governor ought to be considered as a prototypical model for cognition, a replacement for the digital computer. He argues in detail that the Watt governor does not work by representing its environment (see below). If van Gelder is correct that dynamical systems like the

---

[3] Clark's action-oriented representations are a descendent of Millikan's *pushmi-pullyu representations* (1996).

[4] Assuming they use representations at all, it is likely that many non-human animals and some human sub-personal mechanisms use only action-oriented representations.

paradigmatic models of cognition are not representational, cognitive science would need to be utterly transformed. The ultimate purpose of such arguments is to claim that the governor is not a computational system: if one takes computation to be the rule-governed manipulation of representations (Haugeland 1985), then there can be no computation without representations. Arguing against treating the Watt governor as a representational system is merely a means to this end.[5] There are two distinct anti-representationalist claims made in that essay: one ontological, one epistemological.[6] First, ontologically, van Gelder claims that the existence of a (partial) correlation between parts of the Governor and states of its environment (the steam engine whose speed it controls) is not sufficient for states of the Governor *to be representations*. Second, epistemologically, van Gelder claims that the correct conceptual tools for explaining the behavior of the Watt Governor are dynamical and that ascribing representational states to the Governor *has no explanatory utility*. These are separate claims. Thus we should distinguish two distinct formulations of anti-representationalism in our account:

> **1) The Nature Hypothesis**: Natural cognitive systems do not traffic in representations. That is, nothing inside any cognitive agent meets the standards of the teleological theory of representation; and
>
> **2) The Knowledge Hypothesis**: The best models and/or explanations of natural cognitive systems will not invoke *explicit* representations. That is, there will be nothing in our best models and/or explanations of cognitive agents that is a representation according to the teleological theory of representation.

The nature and knowledge hypotheses are distinct. It is easy to imagine, for example, that the nature hypothesis is true and that humans really are just complex dynamical systems, but they are so complex that the best way for us (with our limited intellects) to explain them is by metaphorically or instrumentally ascribing them mental representations. This is roughly Dennett's intentional stance (1987), under one interpretation. Nonetheless, the nature and knowledge hypotheses are related at least in that the truth of the knowledge hypothesis might be evidence for the truth of the nature hypothesis. Furthermore, the truth of the nature hypothesis seems to imply the truth of a close relative of the knowledge hypothesis, the knoweldge$_{\text{in principle}}$ hypothesis: if cognitive systems really have no representations, then there should be some

---

This is borne out by van Gelder 1998 in which it is argued that the distinction between traditional cognitive science and dynamical systems theory is over computation, not over representation.

[6] van Gelder does not acknowledge making more than one claim in van Gelder 1995 (personal communication), but distinguishing between two varieties of anti-representationalism is his suggestion.

explanation or model of them that does not refer to internal, mental representations, whether or not we can find it or understand that explanation. That is, if natural cognitive systems really are not representational, God at least should be able to understand them that way.

An anti-representationalist, then, might defend either (or both) of two distinct claims**.** The main difference between the nature and knowledge hypotheses can be put as follows: the knowledge hypothesis is to a much greater extent a *(meta)scientific* hypothesis. That is, the knowledge hypothesis concerns how we ought to do cognitive science, whatever the mind is really like. The nature hypothesis, on the other hand, is to a much greater extent a *philosophical* hypothesis; it concerns what the some region of the world (cognitive agents) is really like, however that region is best explained scientifically. In what follows, I will discuss the prospects for these two hypotheses, focusing on two dynamical systems models that have been cited by anti-representationalists.

3. The Watt Governor and the Dynamical Stance

As noted above, van Gelder (1995) describes the operation of the Watt governor of steam engines, which he intends as a benchmark dynamical system, and argues that it is not representational, and hence not computational. His arguments are based on a contrast between two cases: the actual Watt governor and a fictional computational steam governor. Here I will compare three cases: (1) van Gelder's dynamically-modeled Watt governor, (2) van Gelder's account of a fictional computational governor, and (3) a representational, though non-computational, account of the governor as described in (1). Doing so will show that the Watt governor does in fact use representations, so it cannot support the nature hypothesis. This leaves open whether it can be cited as evidence for the knowledge hypothesis.

Start with the actual Watt Governor, and its dynamical systems model. The Watt governor controls the speed of a steam engine as follows.

> It consisted of a vertical spindle geared into the main flywheel so that it rotated at a speed directly dependent on that of the flywheel itself. Attached to the spindle by hinges were two arms, and on the end of each arm was a metal ball. As the spindle turned, centrifugal force drove the balls outward and hence upwards. By a clever arrangement, this arm motion was linked directly to the throttle valve. The result was that as the speed of the main wheel increased, the arms raised, closing the valve and restricting the flow of steam; as the speed decreased, the arms fell, opening the valve and allowing more steam to flow. The engine adopted a constant speed, maintained with extraordinary swiftness and smoothness in the presence of large fluctuations in pressure and load. (van

Gelder, 1995, p. 349)

In van Gelder's proposed dynamical explanation of the governor, its operation is described mathematically. Just as Newton did in his descriptions of the physical world, the behavior of the system of interest is observed, and mathematical equations that describe that behavior are found. In the case of the Watt governor, the instantaneous change of the arm angle when the steam engine is disconnected from the throttle valve is described by the following equation:

$$\frac{d^2\theta}{dt^2} = n\omega^2\cos\theta\sin\theta - \frac{g}{l}\sin\theta - r\frac{d\theta}{dt}$$

where θ is the angle of the arms, n is a gearing constant, ω is the speed of the engine, g is the gravitational constant, l is the length of the arm, and r is a friction constant (see van Gelder 1995; this paragraph and the next follow van Gelder's discussion of the Watt governor closely). This equation describes the instantaneous acceleration of the arm angle, given the instantaneous arm angle. Only θ, the arm angle, is a variable in this equation; n, ω, g, l and r are parameters, hence remain constant and fix the dynamics of the system. This equation is completely general in that it gives the acceleration for any arm angle. Solutions to this equation specify a state space, and trajectories through this space can be used to predict future instantaneous accelerations and arm angles, given the current values of these variables.

The governor's behavior when connected to the throttle valve can be described by the following, more complicated equation:

$$\frac{d^n\omega}{dt^n} = F(\omega, \ldots, \tau, \ldots,)$$

where τ is the setting of the throttle valve. This equation, also perfectly general, describes the instantaneous change of the speed of the engine ω as a function of the throttle setting, which is itself a function of the arm angle θ. Just as ω is a parameter in the former equation, θ is a parameter in this equation, so we should think of these two dynamical systems as *coupled*. Any change in the arm angle θ changes the total dynamics of the system that describes the speed of the engine ω, in which it is a parameter; and any change in the engine speed ω changes the total dynamics of the system that describes the change of the arm angle θ, in which it is a parameter. Indeed, given the closeness of the relationship between the two systems, for many purposes, it is best to just think of them as just one system (governor plus engine).

In this, as in all dynamical explanations, once we have found equations such as these for the Watt governor, it is agreed that we have explained the Watt governor's behavior: we have a perfectly general, counter-factual supporting description its behavior, as is provided in Netwonian physics. Note also that, again as in physical explanation, there is no reference to representation, computation, or teleology. If cognitive systems are dynamical systems like the Watt governor, cognition can be explained just as any other complicated physical system is explained.

van Gelder contrasts the actual Watt governor with a *computational governor*, a fictional machine that is meant to embody a computationalist approach to the problem of smoothly controlling the speed of a steam engine. To design such a governor, one would find a description of the task to be performed, then implement that task description in a finite number of simple steps. The computational solution van Gelder imagines to the problem, which would in fact be easily implemented on a digital computer, consists in running the following program:

```
(1)     Begin:
        (i)     Measure the speed of the flywheel;
        (ii)    Compare the actual speed against the desired speed.
(2)     If there is no discrepancy, return to step 1; otherwise:
        (i)     Measure the current steam pressure;
        (ii)    Calculate the desired alteration in steam pressure;
        (iii)   Calculate the necessary throttle-valve adjustment;
        (iv)    Make the necessary throttle-valve adjustment.
(3)     Return to step 1. (van Gelder, 1995, p.423)
```

Note that this governor is much different from the one which Watt actually built. Suppose one were to use this computational description to empirically investigate Watt's actual governor. One would observe the governor carefully, searching for the devices that implement the computations it is assumed to perform, perhaps precisely measuring the time it takes the governor to change the engine's speed to the desired value in an attempt to determine how many steps its computation uses. But here, the computational description is completely misleading. One could look forever and not find the way that these computations are implemented because the Watt governor does not implement the task as described. For one thing there are no digital representations in the Watt governor, and for computation (as standardly understood) to occur, there must be digital representations. Indeed, it is far from obvious that there are any representations at all in the Watt governor.

The computational and dynamical governors that van Gelder describes, however, do not exhaust the space of possibilities: a representational, but non-computational, description of the Watt governor is possible. The Watt governor is designed so that the speed of a flywheel controlling the flow of steam into the engine is in turn controlled by the angle of the rotating arms. In the functioning of the Watt governor, the spindle spins, causing changes in the arm angle, in turn causing the valve to open and close. To see that the arm angles are representations, we start by considering that the angle of the arm is used by the valve to control the engine speed: the higher the arm, the slower the valve makes the engine run; the lower the arm, the faster the valve makes the engine run. It is the function of particular arm angles to change the state of the valve (the representation consumer[7]), and so adapt it to the need to speed up or slow down. For consider that the governor was *designed* so that the arm angle would play this role; that is, arm angle tokens are parts of the functioning of the governor *because* they lead to appropriate control of the engine speed (satisfying R2). So the function of arm angles is to control the speed of the engine, and since each arm angle indicates both a speed and the appropriate response to that speed, is both map and controller, it is an action-oriented representation, standing for the current need to increase or decrease the speed. Since different arm angles are appropriate for different engine speeds, there is a system of representations (satisfying R3). Furthermore, the arm angle can "be fooled", causing behavior for a non-actual engine speed: imagine what would happen if we used a flat surface to hold the arm at an artificially high angle. If we held the arm up in this way, the speed of the engine would decrease and finally halt altogether because the representation used to control the engine speed is of a situation (and its corresponding action) that does not obtain. Thus the arm angles of the Watt governor are action-oriented representations.

Since van Gelder offers the Watt governor as a prototypical dynamical system and a new paradigm for the modeling of cognition, the fact that it is representational is significant, and it implies that other dynamical systems models of cognition are also representational.

The non-computational, representational explanation of the Watt governor begins with and adds significantly to the dynamical story, but does not displace it. In particular, the representational story about the Watt governor adds a teleological dimension[8] to the description of the governor: it answers a "why" question. The representational explanation can be seen as telling us why the Watt governor works the way it does. The explanation begins by assuming that the governor was designed to perform a certain task, and then assigns content to its states based on the way it performs that task. In this case, since we know that the task is to control the speed of the steam engine, we look for parts of the system that are designed to adapt the

---

[7] The spindle is the representation producer in the Watt governor *qua* representational system. (This satisfies R1).

[8] It has been a tradition since Taylor (1964) to think of intentional explanations as a variety of teleological explanation. Representational explanations are a variety of intentional explanation.

governor to aspects of the environment relevant to controlling the engine's speed. We find that we can assign the roles of representation producer and consumer to the spindle and throttle valve, respectively. And the arm angles are action-oriented representations of situations relevant to controlling the engine's speed.

Since the Watt governor has representations according to our definition, it does not provide support for the nature hypothesis. To be sure, there is nothing *computational* (by the standard account) in these models: there are no rule-governed transformations of these representations. The representations in the Watt governor are produced and used, without being subject to rule-governed manipulations, and without necessarily taking part in any inferences.[9] So as a model of cognition, the Watt governor is much different from business-as-usual computational cognitive science. But it still contains entities that meet the standards of the teleological definition of representation.

Systems like the Watt governor might support the *knowledge hypothesis*, however. Indeed to many people, the representational explanation seems superfluous in the case of the Watt governor. The same will be true for many other models favored by dynamical systems theorists, for example Beer's robotic insects (1995) and the evolutionary robots of Harvey, Husbands and Cliff (1994) (see below). For these models, as for the Watt governor, the teleological, representational story doesn't seem any more informative than saying that the robots evolved or were designed for their tasks. Thus one can take up what we might call, with apologies to Dennett, *the dynamical stance* toward these models, explaining their behavior with the tools of dynamical systems theory and avoiding representational vocabulary, while remaining agnostic on the status of the nature hypothesis. In doing so, one admits that the representational story could be told, but maintains that the dynamical systems theory explanation tells us everything important about the system. In fact, this is a natural reaction to the argument that the Watt governor is a representational system. Why should we bother with representational explanations when we have precise, perfectly general, counter-factual supporting mathematical ones? Perhaps part of the reason for this feeling is that the representations in former set are *action-oriented*, so it is fairly difficult for us to say exactly *what* they represent. Another reason, which will be discussed below, is that one must have the complete dynamical story first, before one

---

[9] Bechtel (1998) has also argued that the Watt governor is representational. Bechtel argues that the Watt governor is computational as well. Bechtel and I disagree here because we have different views of what computation is. Bechtel makes a distinction between "processes operating on representations" and "representations figuring in processes". Either case, he thinks, counts as computation. I follow Haugeland (1985) in thinking that only the former of these is computation: having representations figure in processes is not enough.

can have the representational story. If one has the *complete* dynamical story, what is left to be explained? But note that the representational gloss is not positively misleading in the way the computational gloss on the Watt governor would be.

This sense of dissatisfaction with the representational explanation may be just what is necessary for the dynamically-inspired cognitive scientist to affirm the knowledge hypothesis. The dynamical systems theorist can argue for the knowledge hypothesis, via the dynamical stance, as long as (1) there is a large class of dynamical models for which representational glosses add little to the mathematical explanation, and (2) the best explanations of cognitive phenomena fall within this class. The second of these is an empirical matter: we will simply have to wait and see how much of cognition can be explained using dynamical systems models without representational glosses. In the next section, I make a case for the first claim.

4. Evolutionary Robots and the Knowledge Hypothesis

Recent work in robotics at the University of Sussex by Harvey, Husbands, and Cliff (henceforth HHC) (see Harvey, Husbands and Cliff 1994; Husbands, Harvey and Cliff 1995)[10] presents a case of dynamical cognitive science that, we will see, supports the knowledge hypothesis. In their work, control systems for robots are artificially evolved from a randomly generated initial population. Members of this initial group are selected to be "parents" (subject to mutations, etc) based upon their ability to complete particular tasks, for example finding and moving toward a (potentially-moving) target. The researchers purposefully take a hands-off approach to the architecture of the control systems; the only criterion used to determine which systems get to become parents is success at the particular task being selected-for. (In fact, it is their success in their *worst* trial.) Thus, the theorists have no bias for any particular cognitive architecture. Instead they are concerned with achieving skilled performance of the task. By focusing on evolution of skillful behavior, the robot-builders avoid building intricate models of the task domain themselves and coding them into the robot.

Wheeler (1996) claims that one would be hard pressed to produce a representational story for the control system of the successful HHC robots. It seems, he claims, that the robots get by not just without a set of representations constructed by their builders, but without any representations at all. But the fact that the Watt governor (advertised as a prototypical dynamical system) has representations suggests

---

[10] Everything I say about HHC here applies equally to the robots of Beer (1995) and Brooks (1991, 1999).

that *any* dynamical model of cognition will be a representational model. That is, it seems that any such model will posit entities that are representations according to the teleological definition of representation we are using.[11] To see that this is also true of the Sussex evolutionary robots, we must take a closer look at the model than Wheeler does when he claims that it is not representational.[12]

According to the definition of representation in section 1 above, to argue that the HHC's robots do in fact traffic in internal representations, one must find a set of states of the system that are produced by one part of the system, for use by some other part of the system in adapting the system to some aspect of the environment; then one must argue that these parts were designed to interact with one another in this way. This is easily done for Sussex robots, despite the large number of recurrent connections, which lead Wheeler to describe the control system as resembling a bowl of spaghetti (1996, p.220).

Consider a robot with a control system as depicted in Table 1. This robot, described in Harvey, Husbands and Cliff 1994, was the most successful of those that were evolved for target following. After the artificial evolution process, which standardizes the parts of the system to work together to produce the desired behavior, the system works as follows. Nodes 0 and 1 take input from separate visual fields (V1 and V2, respectively). When there is strong input in V1, whatever V2 is like, node 0 causes a pattern of activation that speeds up the left and right motors, causing motion straight ahead. When there is strong input to V2 and weak input to V1, node 1 causes a pattern of activation that leads to increased excitement in node 14, which excites itself and slows the left motor. This causes the robot to rotate in a circle, until a strong input is found in V1. When there is strong input to neither visual field, the noisy node 5, which has no connections from either input node (0 and 1), creates feedback loops that cause the robot to spin in place, until one of the visual fields has input. Thus activation of node 0, from V1, causes a pattern of activation that leads the motors-- via nodes 13, 14, and 15--to behave appropriately with respect the target being straight ahead of the robot. Activation of node 1, from V2, causes a pattern of activation that leads the motors--again, via nodes 13, 14, and 15--to behave appropriately with respect to the target being in sight, but not straight ahead. And, perhaps most interestingly, self-

---

[11] In Chemero 1998, I argue that *every* model of cognition that meets a few conditions will be representational. If that is true, no dynamical model will support the nature hypothesis

[12] Interestingly, Wheeler claims, with little argument, that the model has no representations even though his definition of representation is less restrictive than the one we are using; that is, more things count as representations according to his definition.

activation of node 5 (when there is strong visual input from neither field) causes the robot to behave appropriately with respect to the target being out of visual range. These patterns of activation of the robot's control system are representations according to the teleological definition of representation: the input nodes play the role of representation producers; nodes 13, 14 and 15 are representation producers; patterns of activation across intermediate nodes are representations.

[TABLE 1 ABOUT HERE]

But compare this description to the purely dynamical one that HHC prefer (see Husbands, Harvey and Cliff 1995), the one that comes from adopting the dynamical stance toward the system. In their preferred analysis, the robot and the environment are coupled dynamical systems. To explain their operation, HHC give mathematical descriptions of the structure of each separately, and then, based on those, a unified account of the robot-environment coupled dynamical system. (In what follows, I will provide only a sketch of their argument, leaving out the mathematical details. See Husbands, Harvey and Cliff 1995.) The control system of the robot is given a mathematical account--in fact, I exploited this account in giving my representational account of the robot. Consider the noisy node 5, which excites itself in the absence of visual input to the system: this node, which they call a *generator unit*, provides input to a feedback loop comprising nodes 1 (a visual input unit), 6, 9, 10, 12, and 14 (a motor output unit), which has connections to the other two output nodes, 13 and 15, as well. To see how such a complicated loop works, HHC begin by providing a mathematical analysis of the behavior of a simple, single self-exciting node. This analysis is extended to multiple unit feedback loops, where the behavior is equivalent, but with time delays depending on the number of additional units. Since two of the nodes that take give output to the motors (14 and 15) have connections back to the nodes that take input from the visual fields, the analysis in terms of feedback loops can be extended to cover the behavior of the whole network. With such an account, the behavior of the network, given any input to the visual fields, can be predicted.

The next step in HHC's analysis is a dynamical description of the robot's task environment. This description is done in a space of egocentric polar coordinates $r\theta$, where r is the distance from the robot to the center of the task environment and θ is the clockwise angle from the front of the robot to the center of the task environment. It is possible to determine what the robot's visual input will be for every possible coordinate in the $r\theta$ space, providing a full description of the robot's visual environment. Finally, HHC combine these two dynamical accounts into a combined system that captures the properties of the agent-environment coupling. Once one knows the visual input at every point in the task environment and the behavior given every type of visual input, one can construct a phase portrait that predicts the robot's behavior *no matter where it is* in its environment. For the robot whose control system we have been discussing, this phase portrait has just one attractor, corresponding to the location of the target. Furthermore, every point in the phase portrait is in the basin of attraction for the target's attractor. So one can predict that the robot will succeed at its task every time, and when HHC tested the robot's performance it did succeed every time.

There are reasons to prefer HHC's dynamical account to the representational one described above. First, the representational story depends upon the dynamical story about the control system. It was that mathematical description of the control system in terms of feedback loops that allowed me to predict what behavior would be produced when particular patterns of activation were produced in the system. To find out what those patterns of activation represented, I determined what environmental situations those activations would adapt the agent to, in particular, by determining what environmental situations the ensuing behavior would be appropriate to. So the representational description is dependent upon the dynamical one.

Perhaps more importantly than this dependence, though, the representational description of the system does not add much to our understanding of the system. Once we have the full dynamical story, we can predict the behavior of the robot in its environment completely, and we can do so without making reference to the representational content of any states of its control system. The same is true of the Watt governor. In both cases, the dynamical stance pays off: fully predictive mathematical descriptions of the systems are provided. And despite the fact that a representational gloss is possible, once one has the dynamical explanation, the representational gloss does not predict anything about the system's behavior that could not be predicted by the dynamical explanation alone.

The lack of a need for representational explanation to accompany dynamical models can be made more clear via a historical parallel. As van Gelder (1998) points out, dynamical cognitive science is an attempt to fulfill Hume's goal of a scientific psychology similar to Netwon's mechanics--a psychology in which cognition would be explained by mathematical laws. One of the most striking and important features of Newtonian physics is that the sort of covering law explanations[13] that were provided by

---

[13] A covering law explanation is one in which a fact is explained iff a statement of that fact is deduced from

Newton's mechanics obviated any need for teleology in physical explanation. In providing a fully general set of mathematical laws for physics, Newton sidestepped speculation about Aristotelian final (=teleological) causes, taking his laws of motion as axioms not in need of further explanation. Similarly, once one has mathematical covering laws for psychology, laws which like HHC's model can predict the behavior of agents in their environments with great accuracy, there may be no need for teleological explanations in psychology. And since representational explanations are a species of teleological explanations, a mature dynamical cognitive science might make them obsolete.

So dynamical models of cognition like HHC's robots do indeed provide support for the knowledge hypothesis. They show that it is possible that significant portions of cognition might be explained without referring to explicit, internal representations. But it must be noted that these models are capable of only very simple behaviors, behaviors which have not been traditionally considered "cognitive". This is not to say that dynamical cognitive science is incapable of explaining behaviors traditionally thought of as "cognitive" such as linguistic processes, syntax, and decision making. In fact, work by Ellman, Petitot, and Busemeyer and Townsend (collected in Port and van Gelder 1995) addresses these very topics. But these theorists invoke representations in their explanations. For dynamical cognitive science to vindicate the knowledge hypothesis, for the dynamical stance to pay off, its proponents must provide models--and covering law explanations--of more traditionally cognitive behavior that are not usefully viewed as representational. Only time will tell whether this will be possible.

5. A (Potential) Problem for Dynamical Accounts

Although it seems to be an empirical matter whether dynamical cognitive science will provide compelling, non-representational explanations of great stretches of cognition, there is a potentially serious methodological problem for dynamical accounts, one that might arise for any research program that provides only covering law explanations. And if there is a serious methodological problem for dynamical cognitive science, there is reason to think that dynamical cognitive science will not provide sufficient explanations to vindicate the knowledge hypothesis. In what follows, I will explain this potential methodological problem for dynamical cognitive science, and explore a

---

statements, at least one of which is a general scientific law. Note also that according to the covering law model of explanation (and dynamical models of cognition), prediction is the dual of explanation: a fact is explained iff could have been predicted. See Hempel 1965.

possible resolution to it. This will require a brief, but instructive, digression into the history of physics.

The debate between computationalists and dynamicists in cognitive science runs closely parallel to that between atomists (e.g., Boltzmann) and phenomenalists (e.g., Mach) in theoretical physics at the beginning of the 20th century. Because of this, the argument made by Boltzmann against Mach's physics can also be made by computationalists against dynamical cognitive science. This argument, which I will call the "guide to discovery" argument, was devastating to the phenomenalist picture of physics, and is potentially devastating to the dynamical hypothesis in cognitive science.

The guide to discovery argument was made in response to Mach's philosophy of science (see Mach 1886). Mach was a *phenomenalist*; that is, he believed that everything that there is is available to the senses. This constrains the sort of theories that might be admissible in the sciences; they too must posit no non-sensible entities or properties. Thus Mach argues for a strictly phenomenological physics, the purpose of which is to provide what we would now call covering law explanations for physical phenomena. He was therefore opposed to any theory that posited non-sensible entities of any kind, as atomists in physics did. This led to an ongoing, sometimes heated debate between phenomenological and atomist physicists.

In his "Recent Methods in Physics" (1900), Boltzmann makes the "guide to discovery" argument in favor of atomistic physics. Boltzmann begins the argument as follows:

> The question simply is whether there are not additional results which atomism only could have achieved, and of such results the atomistic theory has had many remarkable specimens to show, even long after the period of its greatest glory. (p.253)

As an answer to "the question" Boltzmann describes some recent triumphs of atomistic physics, such as Van der Waals' greatly improved formula to predict the behavior of the aggregate states of simple chemical substances, improvements to Avogadro's law, Gibb's theory of dissociation, and hydrodynamics. All these successes, Boltzmann claims, could not have been achieved without atomistic assumptions. He concludes the argument:

> If phenomenology deems it expedient, as it certainly must, constantly to institute new experiments for the purpose of discovering necessary corrections for its equations, atomism accomplishes much more in this respect, in that it enables us to point definitely to the experiments which are in most likelihood to lead to its correction. (p.254)

Atomism, then, is the best methodology for physics because it provides a *guide to* *discovering* new equations that describe the phenomena more accurately; by assuming that *there are atoms*, one is led to certain testable predictions of new phenomena.

The point of this argument is best seen by coopting Peter Clark's (1976) account of the ultimately failure of phenomenological physics. Clark characterizes phenomenological physics as *fact-dependent*. In other words, phenomenological physics proceeded by making empirical generalizations about substances and then altering the parameters of descriptive equations to fit anomalies in the experimental results. Thus the only way to improve phenomenological physics was by *ad hoc* additions to the theory, in light of new empirical facts. That is, phenomenological physics, because it refuses to postulate underlying, unobservable structure provides no guide to discovery. And, as Clark puts it, "[t]his quite marked limitation of the heuristic of thermodynamics meant that there was no way of systematically improving the theory." (1976, p.44)

Boltzmann's criticism of non-atomistic theories in physics boils down to the observation that they are fact-dependent, and so unlikely (because they avoid reference to underlying structure) to offer reasonable testable predictions that might improve abilities to explain. That is, unlike atomistic theories, phenomenological theories offer no guide to discovery, and can only proceed in an *ad hoc* manner. Atomistic physics, on the other hand, is not fact-dependent; its practitioners, therefore, can make substantial predictions and then test them. Put simply, phenomenological physicists must constantly alter their theories to fit *existing* empirical results after experimentation, while atomists can use their micro-theory to *predict* empirical results before experimentation. This is a significant methodological advantage for the atomists.

We can substitute computational and dynamical cognitive science into this debate. Computationalism, like atomism, posits an underlying mechanism--computations performed upon representations. This mechanism can be used to predict new, as-yet-unobserved phenomena, and then perform tests in order to improve our understanding of the mind. There are literally thousands of results in cognitive science that, like the improvement to Avogadro's number that depended upon the assumption of atomism, would not have been achieved without the positing of internal mental representations. One obvious example is the results on *mental rotation* (See Shepard and Metzler 1971). In a well-known experiment, Shepard and Metzler posited picture-like mental representations and predicted that there would be temporal effects associated with the operations performed upon them. In particular, they predicted that to determine whether two similar three-dimensional shapes were the same shape at different orientations, subjects would mentally rotate one of them and the time it would take them to decide would be proportional to the degree of rotation. The experiments showed exactly the temporal effect that was predicted.[14]

On the other hand, dynamical cognitive science, at least dynamical cognitive science that does not posit mental representations, is like phenomenalism in that it posits no underlying mechanism for cognition. It is essentially a phenomenological psychology: it is successful when it provides equations that capture observed behavior. It is, therefore, fact-dependent. Dynamicist cognitive scientists must first make empirical observations and then alter their theory to fit them. This, then, is the "guide to discovery" criticism of the dynamical hypothesis: dynamical systems theory, because it posits no underlying mechanism for cognition, is a fact-dependent theory. It provides no guide to discovery, and therefore is bad scientific method. Indeed, given the Humean nature of dynamical cognitive science (van Gelder 1998) and Mach's obvious intellectual debts to Hume, it is not surprising that dynamical cognitive science has the same problems as Mach's phenomenological physics.

Notice that the guide to discovery argument can only be an argument for *instrumental theories*, such as methodological atomism or computationalism. If we conclude that atomism is better equipped to yield progress than fact-dependent theories, then we conclude only that we should do atomistic physics, as opposed to phenomenological physics. On the face of it, this is different from arguing for what we might call *realist* atomism. A realist atomism would claim not just that we should do atomistic physics, but also that the world really is composed of atoms, unobservable though they may be. The guide to discovery argument cannot lead to such a conclusion. It can only support a conclusion such as this: whatever the underlying nature of reality, atomism is more likely than fact-dependent theories to increase our ability to describe the world. Similarly, it is not possible to argue from the claim that dynamical cognitive science is fact-dependent, to the conclusion that cognition *really* is representational or computational. The argument is instead for the *practical necessity* of positing representations and computational operations upon them. But since we are considering the knowledge hypothesis--the claim that the best explanations of cognition will not invoke representations--here, any argument for the practical necessity of positing representations in our explanations of cognition is sufficient. That is, it might be that although cognitive systems really are dynamical systems, our understanding of

---

[14] Commenting on a draft of this essay, Bill Bechtel has suggested a second means of discovery available to computationalists, but unavailable to dynamicists. With computational explanation one might make predictions about the consequences of replacing component X of the system with component Y, and in so doing learn/confirm details about the roles of X and Y in the system. He is correct that dynamical models give no indication as to how one might learn modifying or re-designing a cognitive system in this way. But what matters is that there is *some* way for the dynamical stance to offer a guide to discovery. We shall see below that it can.

them requires that we treat them as having representations. The point is that representational stories might provide crucial leverage for understanding behavior, perhaps especially more complicated, more "cognitive" portions of behavior. If this were to be true, the dynamical stance would not pay off, and the knowledge hypothesis would be false.

The question, then, is whether dynamical cognitive science can avoid being purely fact-dependent, despite the fact that it posits no single underlying mechanism. There is reason to think that it cannot. Think about the way that one does dynamical cognitive science: first, one observes some cognitive behavior; then one tries to find the relevant parameters and variables that define the dynamical system that the behavior instantiates; finally, one finds equations that specify the trajectories through the state space defined by the dynamical system. All of this proceeds *ad hoc*, by adding another dynamical explanation for another observed phenomenon, just as in the case of phenomenological physics. To see this, suppose that all of the dynamical accounts in Port and van Gelder 1995 are correct, and combine them to form a full dynamical account of human cognition. The result would be a hodgepodge, with every covered aspect of cognition being given a different *sort* of dynamical explanation, related only in that they all use the multifarious tools of dynamical systems theory. To add Skarda and Freeman's account of olfaction (1987) to that account, one simply says that olfactory behavior will be explained by the covering laws of Skarda and Freeman's model of olfaction. One does this, despite the fact that Skarda and Freeman's model differs fundamentally from the models in Port and van Gelder 1995 in that it is a *chaotic* dynamical system. The same sort of *ad hoc* addition would accompany every new dynamical explanation of some aspect of cognition, whether or not the explanation was interestingly related to those that are already part of the overall account. Taking the dynamical stance toward cognitive systems might not provide enough predictive leverage to be scientifically fruitful.

6. A Possible Solution

Although dynamical cognitive science as it is currently practiced has fact-dependent tendencies, it may not have to be that way. Dynamical cognitive science could be such that it posited no underlying mechanism for cognition, but did provide a guide to discovery, if it were able to posit a generally applicable type of dynamical model that accounts for a wide range of cognitive phenomena. This would allow scientists to predict that other, similar behaviors would fall under the same covering laws, and then test that prediction. Such a generalizable dynamical model could provide a guide to discovery, putting dynamical cognitive science on equal methodological footing with computational and representational cognitive science. As it happens, one such generalizable dynamical model of cognition has proven widely applicable, and its range is being extended to more aspects of cognition: the well-known HKB model (Kelso 1995; Haken, Kelso and Bunz 1985).[15]

The HKB model is based upon a very simple, very robust experimental result. Subjects asked to wag their index fingers left-to-right can produce only two stable patterns of bimanual coordination. In one, called *in-phase* or *relative phase 0*, the fingers approach one another at the mid-line of the body; in the other, called *out-of-phase* or *relative phase .5*, the fingers move simultaneously to the left, then to the right, like the windshield wipers on most cars. As subjects were asked to wag their fingers out-of-phase at gradually increasing rates, they eventually were unable to do so, and slipped into in-phase wagging. The HKB model for this behavior applies a vector field to the relative phase of the fingers. At slower rates, this field has two attractors, one at relative phase .5, another at relative phase 0. This means that any finger wagging will tend to be stable only when one of these values for relative phase is maintained. But as the rate increases (and passes what HKB call the *critical point*), the attractor at .5 disappears, so the only remaining attractor is at relative phase 0. So finger-wagging at higher rates will tend to be stable only when it is in-phase.

The mathematical model of this behavior, the actual HKB model, is the following function

V = □acos□ □ bcos2□,

where □ is the relative phase and the ratio *b/a* is inversely related to rate. This function, it is worth noting, is the simplest that will accommodate all the data. The HKB model is an example of a general strategy for explaining behavior. First, observe patterns of macroscopic behavior; then seek collective variables (like relative phase) and control parameters (like rate) that govern the behavior; finally, search for the simplest mathematical function that accounts for the behavior. Because, HKB argue, complex systems (like the one involving the muscles, portions of the central nervous system, ears, and metronome in the finger-wagging task) have a tendency to behave like much simpler systems, one will often be able to model these systems in terms of extremely simple functions, with only a few easily observable parameters, which reflect the dynamic behavior. So it might be possible to develop models and covering laws for

---

[15] There are other possibilities of course. One such is Gibson's (1979) ecological psychology (see also Reed 1996). Ecological psychology is often seen as consistent with dynamical systems theory.

many complex behaviors using this method. Indeed, the HKB model itself can be applied to many facets of limb coordination (Kelso 1995). Furthermore, the model makes very specific predictions. First, it predicts that as rates increase, experimental subjects will be unable to maintain out-of-phase performance. Second, even at slow rates only relative phase of 0 and .5 will be stable. Third, the behavior should exhibit *critical fluctuations*: as the rate approaches the critical value, attempts to maintain out-of-phase performance will result in erratic fluctuations of relative phase. Fourth, the behavior should exhibit critical slowing down: at rates near the critical value, disruptions from out-of-phase performance should take longer to correct than at slower rates.

The HKB model is an interesting example of dynamical cognitive science. But as it stands, it is not sufficient to save dynamical cognitive science from fact dependence--it still provides no “guide to discovery”. Recent work by Robert Port and his students may, however, be changing this. Port has suggested that one might be able to develop a *general theory of meter*, based on the HKB model. First, Kaipainen and Port (unpublished) performed experiments showing that certain simple speech tasks conform to the predictions of the HKB model. Such results suggest, Port claims, that the HKB model captures something general which might underlie speech actions as well as limb motions. Furthermore, it suggests that a more general variant of the HKB model might explain large swaths of human behavior.

Port’s general model of meter derives from the fact that the speech task explored by Kaipainen and Port matches the predictions of the HKB model, but shows additional attractors. The speech task has attractors at relative phases of 0 and .5, as in HKB’s finger wagging task, and weaker attractors at .33 and .67. These two new attractors can be accounted for by adding a term to the HKB equation:

$$V = \square a\cos\square \; \square \; b\cos 2\square \; \square \; c\cos 3\square .$$

If c is large, attractors appear at relative phase .33 and .67, along with those at 0 and .5. The HKB model is a special case of this more general model, where c = 0. Such a model, Port says, might be fully general: we might be able to apply it to every repetitive motor pattern (such as walking, running, swinging a hammer, chanting), as well as to the perception and production of music and speech.

If Port is right, his model does provide a guide to discovery, providing numerous testable predictions that will allow non-*ad hoc* extensions to our understanding of cognition. It does so by potentially capturing *all of rhythmic behavior* with one dynamical model. A cognitive scientist interested in any rhythmic aspect of cognition can use the model to make predictions about phenomenon of interest. It should have all of the following characteristics: there should be attractors (stable performance) only at relative phases of 0, .33, .5, and .67; the attractors at .33 and .67 should be stable only at slower rates, while that at .5 should be stable at slightly higher rates, and at the highest rates only that at 0 should be stable; it should exhibit critical fluctuation; it should exhibit critical slowing down. The specificity of these predictions provides a principled way for dynamical cognitive scientists to extend their understanding.

So dynamical cognitive science may not be fact-dependent after all; the guide to discovery argument that was so damning to phenomenological physics might not apply. What is necessary to fully exonerate dynamical cognitive science of the charge of fact dependence is the development of more explanatory strategies as generally applicable as Port's extension of the HKB model. There is no principled reason to think that these will not be forthcoming. There is, then, no principled reason to claim that dynamical cognitive science is methodologically inferior to computational cognitive science.

7. Conclusion

Since dynamical cognitive science may be able to provide guidance to its practitioners, we have as yet no reason to think that it might not eventually vindicate the knowledge hypothesis. Its ability to do so depends not just upon the future production of many, many more satisfying accounts of cognition, but on the nature of those accounts: they must be like those provided by Harvey, Husbands and Cliff, in which the representational explanation does not explain anything that is not already explained by the dynamical account of the agent-environment system. Furthermore, they must be extendable to other cognitive phenomena, as Port has extended the HKB model. Whether this will happen is an empirical matter. This, despite the fact that we have seen reason to think that the nature hypothesis--the hypothesis that cognitive systems do not use representations--is false. This points to some advice for the would-be anti-representationalist. Refrain from arguing that cognitive systems *really* are not representational; instead, argue that the best way to understand cognition is with the tools of dynamical systems theory, by taking up what I have called the dynamical stance. The best way to argue for the fruitfulness of the dynamical stance is by example: get to work providing non-representational explanations of cognitive phenomena that are both convincing and sufficiently rich in their implications to guide further research.

Acknowledgements

Thanks to Bill Bechtel, Will Cordeiro, Mike Dunn, Michael Friedman, Ruth Millikan, Bob Port, Gregor Schöner, Brian Cantwell Smith, and Tim van Gelder for comments.

Works Cited

Bechtel, William. (1998). "Representations and cognitive explanations." *Cognitive Science* 22: 295-318.

Beer, Randy. (1995) "Computational and Dynamical Languages for Autonomous Agents." In Port and van Gelder 1995.

Boltzmann, Ludwig. (1900). "The Recent Development of Method in Theoretical Physics." *Monist* 11, 226-257.

Boltzmann, Ludwig. (1901). "On the Necessity of Atomic Theories in Physics." *Monist* 12, 65-79.

Brooks, Rodney. (1991). "Intelligence Without Reason." *Artificial Intelligence* 47: 139-159.

Brooks, Rodney. (1999). *Cambrian Intelligence*. Cambridge: MIT Press.

Busemeyer, Jerry and Townsend, James. (1995). "Decision Field Theory." In Port and van Gelder 1995.

Chemero, Anthony. (1998) *How to be an Anti-Representationalist*. Doctoral Thesis in Philosophy and Cognitive Science, Indiana University, Bloomington, IN.

Chemero, Anthony. and Eck, Douglas. (1999). "An Exploration of Representational Complexity via Coupled Oscillator Systems." In U. Priss (ed.) *Proceedings of the 1999 Midwest AI and Cognitive Science Conference*, Cambridge: AAAI Press.

Clark, Andy. (1997). *Being There*. Cambridge: MIT Press.

Clark, Peter. (1976). "Atomism versus Thermodynamics." In *Method and Appraisal in the Physical Sciences*, C. Howson, ed. New York: Cambridge University Press.

Dennett, Daniel. (1987). *The Intentional Stance*. Cambridge: MIT Press.

Fodor, Jerry and Pylyshyn, Zenon. (1988). "Connectionism and Cognitive Architecture." *Cognition* 28: 3-71.

Gibson, James J. (1979). *The Ecological Approach to Visual Perception*. NewJersey: Houghton Mifflin.

Grush, Rick. (1997). "The Architecture of Representation." *Philosophical Psychology.*

Haken, Hermann, Kelso, J.Scott, Bunz, H. (1985). "A Theoretical Model of Phase

Transitions in Human Hand Movements." *Biological Cybernetics* 51: 347-356.

Harvey, Inman, Husbands, Phil, and Cliff, David. (1994). "Seeing the Light: Artificial Evolution, Real Vision." In David Cliff (ed.) *From Animals to Animats 3*, Cambridge: MIT Press.

Haugeland, John. (1985). *Artificial Intelligence: The Very Idea*. Cambridge: MIT Press.

Hempel, Carl. (1966). *The Philosophy of Natural Science*. Englewood: Prentice-Hall.

Hendriks-Jansen, Horst. (1996). *Catching Ourselves in the Act*. Cambridge: MIT Press.

Husbands, Phil, Harvey, Inman, and Cliff, David. (1995). "Circle in the Round: State Space Attractor for Evolved Sight Robots. *Journal of Robotics and Autonomous Systems.*

Keizjer, Fred. (1997). *The Generation of Behavior*. Leiden University Ph.D. Dissertation.

Kelso, J. Scott. (1995) *Dynamic Patterns*. Cambridge: MIT Press.

Kaipainen, Mauri. and Port, Robert. (unpublished). "Speech Cycling."

Mach, Ernst. 1886. *The Analysis of Sensations*. Translated by Williams. New York: Dover.

Millikan, Ruth. (1984) *Language, Thought and Other Biological Categories*. Cambridge: MIT Press.

Millikan, Ruth. (1996) "Pushmi-pullyu Representations." In James Tomberlin, ed., *Philosophical Perspectives* vol. IX. Ridgeview Publishing.

Port, Robert. (unpublished). "A Dynamic Approcah to Language and Cognition."

Port, Robert. and van Gelder, Timothy (eds.). (1995). *Mind as Motion*. Cambridge: MIT Press.

Ramsey, William. (1997). "Do Connectionist Representations Earn their Explanatory Keep?" *Mind and Language*, 12: 34-66.

Reed, Edward. (1996). *Experiencing the World*. New York: Oxford University Press.

Shepard, Roger and Metzler, Jacqeline. (1971). "Mental Rotation of Three-Dimensional Objects." *Science* 171: 701-3.

Skarda, Christine and Freeman, Walter. (1987). "How the Brain Makes Chaos to Make Sense of the World." *Behavioral and Brain Sciences*., 10, 161-195.

Taylor, Charles. (1964). *The Explanation of Behavior*. London: Routledge and Kegan Paul

Thelen, Esther. and Smith, Linda B. (1994). *A Dynamic Systems Approach to the Development of Cognition and Action*. Cambridge: MIT Press.

van Gelder, Timothy. (1995). "What might Cognition be if not Computation?" *Journal of Philosophy*, 91, 345-381.

van Gelder, Timothy. (1998). "The Dynamical Systems Hypothesis." *Behavioral and Brain Sciences*, 21, 1-14.

van Gelder, Timothy. and Port, Robert. (1995). "It's about Time." In Port and van Gelder 1995.

Wheeler, Michael. (1996) "From Robots to Rothko." In Boden, M. (ed), *The Philosophy of Artificial Life.*

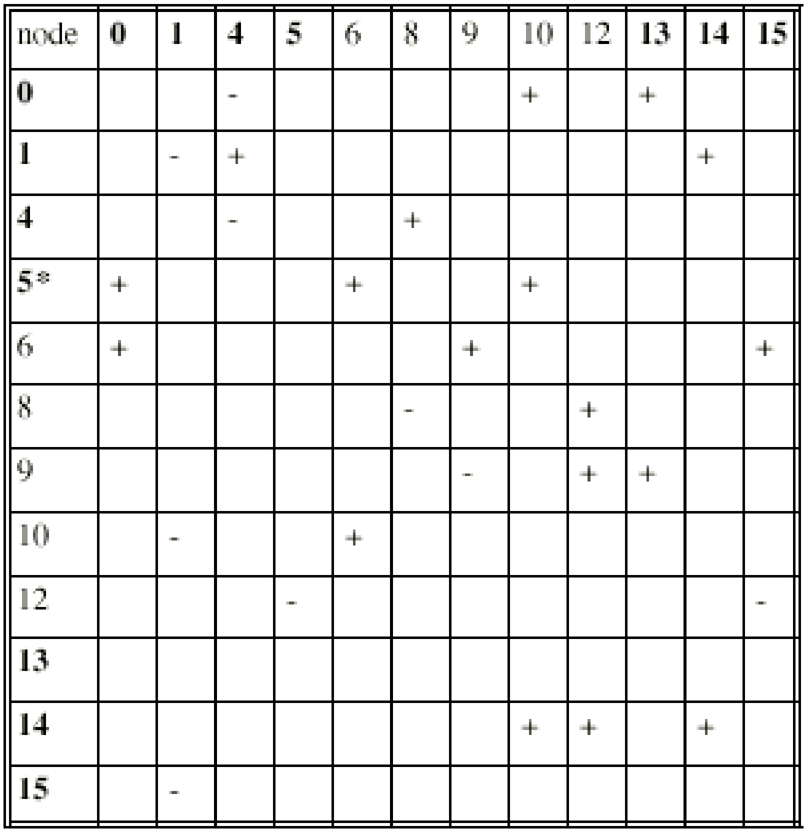

node	**0**	**1**	**4**	**5**	6	8	9	10	12	**13**	**14**	**15**
0			-					+		+		
1		-	+								+	
4			-			+						
5*	+				+			+				
6	+						+					+
8						-			+			
9							-		+	+		
10		-			+							
12				-								-
13												
14								+	+		+	
15		-										

Table 1

Table 1: Connections in a Sussex Robot. This table shows the connections in the most successful target tracking robot discussed in Harvey, Husbands and Cliff 1994. A mark in a square indicates that there is a connection from node with that row number to the node with that column number. A '+' indicates an excitatory connection; a '-' indicates an inhibitory connection. Nodes 0 and 1 take input from visual fields V1 and V2, respectively; node 13 increases voltage to the left motor; node 14 decreases voltage to the left motor; node 15 increases voltage to the right motor. Node 5 is an extra noisy unit that takes no input from either visual field.