"Time and Mind" JOURNAL OF PHILOSOPHY: XCV: 7 :1998: p. 354-376

TIME AND MIND*

Andy Clark

Philosophy/Neuroscience/Psychology Program
Washington University in St. Louis
St. Louis, MO
USA
andy@twinearth.wustl.edu

ABSTRACT.

Mind, it has recently been argued[1], is a thoroughly temporal phenomenon: so temporal, indeed, as to defy description and analysis using the traditional computational tools of cognitive scientific understanding. The proper explanatory tools, so the suggestion goes, are instead the geometric constructs and differential equations of Dynamical Systems Theory. I consider various aspects of the putative temporal challenge to computational understanding, and show that the root problem turns on the presence of a certain kind of causal web: a web that involves multiple components (both inner and outer) linked by chains of continuous and reciprocal causal influence. There is, however, no compelling route from such facts about causal and temporal complexity to the radical anti-computationalist conclusion. This is because, interactive complexities notwithstanding, the computational approach provides a *kind* of explanatory understanding that cannot (I suggest) be recreated using the alternative resources of pure Dynamical Systems Theory. In particular, it provides a means of mapping information flow onto causal structure -- a mapping that is crucial to understanding the distinctive kinds of flexibility and control characteristic of truly mindful engagements with the world. Where we confront especially complex interactive causal webs, however, it does indeed become harder to isolate the syntactic vehicles required by the computational approach. Dynamical Systems Theory, I conclude, may play a vital role in recovering such vehicles from the burgeoning mass of real-time interactive complexity.

1. TIME, CAUSE, AND COMPLEXITY.

> The heart of the problem is *time. Cognitive processes and their context unfold continuously and simultaneously in real time*. Computational models specify a discrete sequence of static internal states in arbitrary "step" time (t1, t2, etc.). Imposing the latter onto the former is like wearing shoes on your hands. You can do it, but gloves fit a whole lot better
>
> Van Gelder, T. & Port, R. It's About Time: An Overview of the Dynamical Approach to Cognition. In R. Port & T. Van Gelder (Eds.), Mind as Motion: Explorations in the Dynamics of Cognition. (Cambridge, MA: MIT Press. p.1-44) p. 2. Italics in original

Time, it seems, is the skeleton in the Cognitive Scientific closet. A recent spate of criticisms (see note 1) of the bedrock explanatory apparatus of cognitive science all focus, in varying degrees, on the need to factor timing into acceptable explanations of cognitive phenomena.

The basic argument[2] is quite simple. It consists of three observations and a conclusion. The observations are:

1. That cognitive processes all take place in real time;
2. That the timing really matters, and
3. That computational models ignore real timing in favor of the artificial notion of a simple sequence of states.

Observation (1) is certainly true, at least as a matter of fact. Observation (2) looks plausible enough. As van Gelder & Port point out, 'the timing of any particular operation must respect the rate at which other cognitive, bodily and environmental processes are taking place' (op. cit., p. 19). Nor can there be much doubt about observation (3). Certainly, the formal theory of computation makes no concessions to real timing. It concentrates entirely on the sequence of (discrete) states[3].

Yet for all that, the basic argument still looks too quick to be convincing. One immediate cause for concern is that it would seem to apply equally well to real-world computer systems! Such systems depend on processes that occur in real-time, and the timing of events really matters. Programs that must interact with other programs and peripherals (printers, scanners and the like) must likewise respect the rate at which various processes are taking place. Yet it would seem premature to conclude that a dynamical understanding is preferable to a standard computational one in such everyday cases.

The real argument, then, must be a little more complex. And so it is. It depends on two additional factors. The first concerns the continuous nature of the relevant processes. The second concerns the complexity of the causal webs in which such states participate.

A continuous process is one in which the time-series of explanatorily relevant sub-states cannot be reduced to a sequence of discrete states with jumps in between, but instead requires a genuine continuum of states. This is not the case in conventional computers, which are designed carefully so that their proper functioning is explicable in terms of sequences of discrete states. Here then, is one potential difference between the case of ordinary computation and biological cognition. Perhaps biological processing makes real use of a continuum of possible states. If so, then any formal computational model must leave out something of theoretical importance, whereas a dynamical model (as we will see) need not.

More important than this, however, is a nexus of considerations concerning the *type* of causal web in which various states and processes may participate. For it is here that we discern the most potent challenge that temporal considerations pose to traditional forms of explanations and analysis. The challenge is to account for cases of what I elsewhere[4] term continuous reciprocal causation, or CRC for short.

Continuous reciprocal causation (CRC) occurs when some system S is both continuously affecting and simultaneously being affected by, activity in some other system O. Internally, we may well confront such causal complexity in the brain, since many neural areas are linked by both feedback and feedforward pathways[5]. Worse still, we may find processes of CRC that criss-cross brain, body and local environment. Think of a dancer, whose bodily orientation is continuously affecting and being affected by his neural states, and whose movements are also influencing those of his partner, to whom he is continuously responding! Or imagine playing improvised jazz in a small combo. Each musician's playing is influencing and being influenced by everyone else. CRC looks, in fact, to pervade the field of natural adaptive intelligence. The delicate dance of predator and prey, or of mating animals, exhibits the same complex causal structure. Historically, CRC looms large in cybernetics[6] and is also wonderfully described by the phenomenologist Maurice Merleau-Ponty:

> When my hand follows each effort of a struggling animal while holding an instrument for capturing it, it is clear that each of my movements responds to an external stimulation; but it is also clear that these stimulations could not be received without the movements by which I expose my receptors to their influence ... the properties of the object and the intentions of the subject are not only intermingled; they also constitute a new whole
>
> Merleau-Ponty, M. La Structure du Comportment (A. Fisher, Trans.). France: Presses Universitaites de France 1942, p. 13.

Here, the motions of my hands are continuously responsive to those of the struggling animal, while the animal's wrigglings are continuously shaped by the motions of my hand[7].

Van Gelder (1995- see note 1 above) makes this sort of interactive complexity the centerpiece of an extended argument aimed to reveal the shortcomings of computational models of mind. The key move in the argument is the suggestion that agent-environment interactions may involve such a complex and continuous process of 'give and take' that the idea of information entering the agent through perception, giving rise to an internal representation and issuing in an action (and that cycle repeating over and over) is just too simple to do it justice. The kinds of agent-environment coupling characteristic of our cognitive engagements with the world are thus said to '*transcend* representation' (*ibid.,* p. 381). To illustrate this central theme, van Gelder offers a simple, non-cognitive illustration: the operation of the Watt (or centrifugal) Governor.

The task of this device (invented by James Watt in the late 18th century) is to keep constant the speed of a flywheel to which some machinery is connected. There is a tendency for the speed to fluctuate due to varying steam pressures and workloads. To smooth things out, the amount of steam entering the pistons is controlled by a throttle valve. How might such control be achieved?

One solution (which van Gelder describes as the computational solution) would involve a sequence of steps and measurements. For example, we might program a device to measure the speed of the flywheel, compare this to some desired speed, measure the steam pressure, calculate any change in pressure needed to maintain the desired speed, adjust the throttle valve accordingly, then begin the whole sequence anew (see van Gelder *op. cit.*, p. 348). What makes this kind of solution computational, van Gelder suggests, is a complex of familiar features. The most important one is representation: the device measures the speed of the flywheel, creates a token that stands for the speed and performs numerous operations (comparisons, etc.) on this and other representations. These operations are discrete and occur in a set sequence which then repeats itself. The sequence involves a perception (measurement)-computation-action cycle in which the environment is probed, internal representations created, computations performed, and an action selected. The overall device is, in addition, functionally articulated and homuncular in construction. For it involves a division of the problem into distinct parts each of which is dealt with independently, with the various parts linked by acts of communication (in which x tells y the value of z and so on). The features distinctive of the computational governor are thus 1) the use of internal representations and symbols, 2) the use of computational operations that alter and transform those representations, 3) the presence of a well-defined perception-computation-action cycle (what van Gelder calls 'sequential and cyclic operation'), and 4) the susceptibility to information-processing decomposition (what van Gelder calls 'homuncularity').

Now for the second solution, the one discovered by James Watt himself. Gear a vertical spindle into the flywheel and attach two hinged arms to the spindle. To the end of each arm, attach a metal ball. Link the arms to the throttle valve so that the higher the arms swing out, the less steam is allowed through. As the spindle turns, centrifugal force causes the arms to fly out. The faster it turns, the higher the arms fly out. But this now reduces steam flow, causing the engine to slow down and the arms to fall. This, of course, opens the valve and allows more steam to flow. By clever calibration this centrifugal governor can be set-up so as to maintain engine speed smoothly despite wide variations in pressure, workload and so on. (This story is condensed from van Gelder, *op. cit.*, p. 347-350).

Having presented the two governors, van Gelder devotes much energy to defending the claim that the centrifugal governor, unlike the computational one, does not rely on representation or computation. Rather than rehearse these considerations[8], I propose to focus on just one point. The deeper reason, we are told, for supposing that the Watt Governor is not representational is that the relationship between the angle of the arm and the engine speed exhibits a special complexity that the framework of representations, computations, and information-processing simply cannot capture. It is a complexity that arises out of the process of continuous reciprocal causation itself and that is better captured using the alternative apparatus of dynamical systems theory. The key idea is that although there is a relationship between arm angle and engine speed, it is one that is 'much more subtle and complex than the standard notion of representation can handle' (ibid., p. 353). It is more subtle and complex because the arm angle is continuously modulating the engine speed at the same time as the engine speed is modulating the arm angle. The two quantities are best seen as being co-determined and co-determining -- a relationship nicely captured using the dynamical apparatus (see section 2 following) of coupled differential equations. The Watt Governor then fails to constitute a *computational* device for two reasons. First, because, on van Gelder's account, computation requires the manipulation of representations (op.cit., p. 353). And second, because there are no discrete operations in the governing processes and hence no distinct sequence of manipulations to identify with the steps in a computational solution[9]. The Watt Governor thus fails to exhibit any of the features associated with the computational solution, and for a single deep reason: the continuous and simultaneous relations of causal influence that obtain among the various factors involved. It is this distinctive kind of causal profile that both invites treatment in terms of an alternative dynamical analysis and that causes problems for the alternative (computational and representational) approach.

Why does CRC cause such problems for standard cognitive scientific modes of analysis and understanding? The root trouble, it should be clear, involves the search for illuminating problem decompositions. For systems whose causal embedding in the world is characterized by processes of CRC are unusually resistant to the kind of divide and conquer [10] approach taken by computational cognitive science. In a divide and conquer analysis, a complex phenomenon is explained by displaying a certain kind of articulation in the mechanisms that underlie it. This articulation need not appear at the most fundamental physical level(s), but may instead exist only at the level of a virtual machine implemented in the physical hardware. For example, word processing software creates a virtual machine whose significant features are quite different from those of the host computer. If we want to understand the distinctive phenomena (operations upon sentences, paragraphs, etc.) supported by such software, we need to display the organization and structure of the higher level virtual machine, for it is only at this level that we find the groupings and chainings of lower-level, basic-machine states that power the characteristic word-processing profile. The object of divide and conquer analysis may thus be either a physical machine or a virtual machine running on top of one. Either way, the analysis proceeds by isolating a number of distinct sub-mechanisms, assigning them specific roles and plotting the chains of causal influence that flow between them. The sub-mechanisms, in the case of computational/information-processing devices, are usually understood as performing operations on semantically-interpretable syntactic items such as numbers, strings and lists that encode information about the processed text. The flow of communication between sub-mechanisms then makes sense insofar as it enables a variety of different operations to be performed in some useful sequence, e.g., first isolating part of the text (by marking it), then moving it, then rendering it in italics, and so on. The analytic strategy is thus one that turns heavily on a) the decomposition of the (real or virtual) machine into component states, processes and sub-mechanisms, and b) some interpretation of the states and structures such that the flow of causation in the machine can be seen to respect constraints on the flow and manipulation of information.

Continuous reciprocal causation poses a clear problem for such a strategy. For although CRC is compatible with basic structural decomposition (distinct parts can be linked by complex reciprocal causal chains), it sits ill with the kind of analysis that assigns specific information-processing roles to components and that interprets those components as effecting a specifiable sequence of operations that solves some problem. In short, it sits ill with the standard idea of an algorithmic depiction of systemic activity. For such descriptions posit a sequence of distinct computational steps or actions that can be seen to solve some problem. The contrast between such a depiction and cases involving CRC is quite marked. Consider, as another simple example, returning a tennis serve:

> The ball is approaching; you are perceiving its approach, are aware of the other player s movements, are considering the best strategy for the return, and are shifting into position to play the stroke. *All this is happening at the same time*. As you move into place, your perspective on the approaching ball is changing and hence so is activity on your retina and in your visual system ... the path of the approaching ball affects which strategy would be best and how you move. *Everything is simultaneously affecting everything else*.
>
> (van Gelder & Port, *op. cit.*, p. 23.). Emphasis in original

In this example, the relevant structure of events is only very coarsely sequential. To cast the phenomena as a step-wise progression, involving a sequence of discrete states, would clearly do extreme violence to the time-coordination of events and processes. The motion of the ball, the tracking (and anticipating) eye movements and the various gross bodily actions, are all unfolding in a kind of phased and inter-animated parallelism. To re-cast such a story in terms of a standard sequential algorithmic description would clearly be to obscure much of importance.

The standard divide and conquer computationalist approach, as van Gelder & Port go on to point out, tends to do just that, and as such, is fundamentally ill-equipped for the analysis of such complex phenomena. Parallel computational models (such as connectionist and neural network approaches[11]) are superficially better bets. But even these threaten to fall short in some important respects. For one thing, they lack intrinsic temporality.

Nothing in a standard connectionist model fixes the real-time profile of events. Such real-time profiles are, for the connectionist as much as the classicist, relegated to the level of 'implementation detail' -- a result that is profoundly inimical to the dynamicist's insistence that timing is *cognitively* crucial. For another, both connectionist and classical computational models alike are challenged by cases in which processes of continuous causal influence forge analytic unities that criss-cross brain/body/world boundaries. In the tennis example, for instance, we confront inner cognitive processes that must at some point interface with external (bodily and environmental) events whose temporal structure is crucial to the interaction. (Think of tracking the approaching ball, and factoring in leg motions and bodily reorientations). If we insist on a computationalist, hence fundamentally atemporal, model of all cognitive processes, then such interfaces present a problem; they must be analyzed in a double-sided fashion, by being shown both to successfully engage the real-time phenomenon and to figure in the more temporally-neutral computational specifications that are supposedly constitutive of the cognitive phenomenon. A common upshot of the pressure to preserve real-time profiles at the interface, however, is a kind of creeping temporal infection. The interface mechanism is depicted in dynamical (real-time involving) terms and this “drives dynamical conceptualizations inward, into the cognitive system itself “ (van Gelder & Port *ibid.*, p. 29). This process repeats as the interface is itself required to make contact with more central processes and the real-time dynamics seep further and further into the system. The final result according to van Gelder & Port is a fully temporalized model in which all processes, from input all the way to output, are conceived in dynamical terms. One advantage of the fully dynamical approach is thus that it does not need to deploy different tools to model bodily and environmental processes on one hand and more traditionally cognitive ones on the other. Where inner and outer processes are linked in cycles of continuous reciprocal causation, there seems to be no need to draw special boundaries around the inner aspects, nor to use a different vocabulary to describe them.

It is this same deep feature (CRC) that thus inclines the radical dynamicist to question even the traditional divisions between cognitive system, body, and world. For if similarly complex causal webs sometimes include neural, bodily and environmental factors (as in the tennis example), then we seem driven to the use of total state analyses that include all these factors on an equal footing. This, I believe, must be the reasoning behind bold[12] assertions such as:

> Cognitive processes span the brain, the body and the environment: to understand cognition is to understand the interplay of all three. Inner reasoning processes are *no more essentially cognitive than the skillful execution of coordinated movement or the nature of the environment in which cognition takes place*. The interaction between inner processes and outer world is not peripheral to cognition, it is the very stuff of which cognition is made
>
> van Gelder & Port *op. cit.*, p. ix. (My emphasis).

All the truly radical consequences (skepticism concerning the notions of computation and representation and concerning the cognitive/non-cognitive boundary itself) thus flow from a single deep source: the idea that intelligent action implicates complex webs of continuous reciprocal causal influence. The problem, in short, is *not* (simply) that time matters and that traditional computational approaches leave it out. Instead, the real culprit is a very specific kind of causal complexity: one that impedes the step-wise problem de-composition characteristic (so the argument goes) of a computational approach, and that even fudges the boundary lines between the cognitive system and the environment.

To sum up, there are three time-related problems that look to motivate recent anti-computationalist conjectures. The first is the simple observation that real-timing matters deeply to adaptive success and practical intelligence. The second is that such timing often involves continuous processes rather than discrete steps. The third is that inner and outer events are linked in complex causal webs that defy decomposition into temporal (and non-cognitive) and cognitive (and non-temporal) components. What we need, it seems, is some framework that builds the temporal facts into the very heart of the analysis, that uses a common vocabulary to describe inner and outer processes, and that is able to capture patterns involving continuous reciprocal causal influences among multiple (inner and outer) parts. Such, it transpires, is the profile (and hence the attraction) of Dynamical Systems Theory, as we shall now see.

2. TOTAL STATE EXPLANATIONS.

Dynamical Systems Theory is a well-established framework[13] for describing and understanding the temporal evolution of complex systems. In a typical explanation, the theorist specifies a set of parameters whose collective evolution is governed by a set of differential equations. Such equations always involve a temporal element, and in this way timing is factored into the heart of the approach. Moreover, such explanations are easily able to span organism and environment. In such cases the two components are treated as a coupled system in a specific technical sense, viz, the equation describing the evolution of each component contains a term that factors in the other systems current state (technically, the state variables of the first system are the parameters of the second and vice versa).

Thus consider two wall-mounted pendulums placed in close proximity on a single wall. The two pendulums will tend (courtesy of vibrations running long the wall) to become swing-synchronized over time. This process admits of an elegant dynamical explanation in which the two pendulums are analyzed as a single coupled system with the motion equation for each one including a term representing the influence of the other's current state[14].

A useful way to think of this is by imagining two co-evolving state spaces. Each pendulum traces a course through a space of spatial and temporal configurations. But the shape of this space is determined, in part, by the ongoing activity of the other pendulum, which is itself behaving in ways continuously modified by the action of its neighbor! This kind of complex, constant, mutual interaction is, van Gelder and others[15] claim, much closer to the true profile of agent-environment interactions than is the traditional vision of a simple perception-computation-action sequence.

The crucial upshot of the emphasis on constant mutual interaction is a corresponding emphasis on what van Gelder & Port (op.cit., p. 14) usefully term *total state*.

Because it is assumed that there is widespread and complex interanimation between multiple systemic factors (x influences y and z, and x is itself influenced by y, which also influences z and so on), the dynamicist chooses to focus on changes in *total system state* over time. The various geometric devices used to put intuitive flesh on the models (trajectories through state spaces populated by attractors, repellors and so on-see note 13) thus reflect motion in a space of possible *overall* system states, with routes and distances defined relative to points *each of which* assigns a value to *all* the systemic variables and parameters. This emphasis on total state marks one of the deepest contrasts between dynamical and standard-computationalist approaches, and it is both a boon and a burden, as we shall see.

A natural question, at about this point, is in what sense a dynamical/total state story counts as a real explanation of some phenomenon? Aren't such stories closer to descriptions than explanations? This is a big topic[16], but one we need not dwell upon here. Suffice it to say that there is no doubt but that a good dynamical systems story is explanatory in at least this sense: it has the power to illuminate counterfactuals -- to predict how the system would behave if perturbed in various ways. This is unsurprising, since a dynamical model aims to display the constraints that determine how overall system-state changes over time. What we don't find, however, is anything that looks much like an account of systemic *mechanisms*. Instead we are shown sets of interlinked parameters, state variables, etc. The discontent which some cognitive scientists feel when confronted with a dynamical story is thus due to the failure to provide something closer to an engineering blueprint. Dynamical models do not speak directly to the project of how to build a device that would exhibit the target features and in this they differ from standard computational models.

For present purposes, however, let us put such reservations aside. Let us use the term total state explanation for any dynamical model that emphasizes the fact that all aspects of a system are changing simultaneously (van Gelder & Port (op.cit.), p. 14) and that therefore invites us to understand the behavior of the system in terms of the possible sequences of changes in total state over time. It is this explanatory strategy that speaks most directly to the needs and worries displayed in section 1, since it is this which allows the dynamicist to respect the burgeoning complexity of causal webs in which everything (both inner and outer) is continuously influencing everything else. Such a strategy, for example, allows us to give a satisfying account of the action of the Watt Governor using a system of coupled differential equations in which the arm angle (and hence the setting of the throttle valve) appears as a parameter in the description of the engine and in which the engine speed appears as a parameter in the description of the governor. The mathematics of such a system of interlocking differential equations perfectly captures the way in which the two systems engage in their continuous, real-time and effectively instantaneous dance of mutual co-determining interaction[17].

The question before us is therefore this: should the undoubted power and success of dynamical explanations in such cases, combined with the demonstration of similarly complex causal webs criss-crossing agents and environments, lead us to reject the computational model of mind? I next scout some reasons to suppose it should not.

3. DYNAMICS & THE FLOW OF INFORMATION.

The deepest problem with the dynamical alternative lies precisely in its treatment of the brain as *just one more factor* in the complex overall web of causal influences. In one sense this is obviously true. Inner and outer factors do conspire to support many kinds of adaptive success. But in another sense it is either false, or our world-view will have to change in some very dramatic fashions indeed. For we do suppose that it is the staggering structural complexity and variability of the brain that is the key to understanding the specifically intelligence-based route to evolutionary success. And we do suppose that that route involves the ability, courtesy of complex neural events, to become appraised of *information* concerning our surroundings, and to use that information as a guide to present and future action. If these are not truisms, they are very close to being so. But as soon as we embrace the notion of the brain as the principal seat of information-processing activity, we are already seeing it as fundamentally different from, say, the flow of a river or the activity of a volcano. And this is a difference which needs to be reflected in our scientific analysis:

a difference which typically *is* reflected when we pursue the kind of information-processing model associated with computational approaches, but which looks to be lost if we treat the brain in exactly the same terms as, say the Watt Governor, or the beating of a heart, or the unfolding of a basic chemical reaction[18].

The question, in short, is how to do justice to the idea that there is a principled distinction between knowledge-based and merely physical-causal systems. It does not seem likely that the dynamicist will deny that there is a difference (though hints of such a denial[19] are sometimes to be found). But rather than responding by embracing a different vocabulary for the understanding and analysis of brain events (at least as they pertain to cognition), the dynamicist re-casts the issue as the explanation of distinctive kinds of behavioral flexibility and hopes to explain that flexibility using the very same apparatus that works for other physical systems, such as the Watt Governor.

Such apparatus, however, may not be intrinsically well-suited to explaining the particular way neural processes contribute to behavioral flexibility. This is because 1) it is unclear how it can do justice to the fundamental ideas of agency and of information-guided choice, and 2) the emphasis on total state may obscure the kinds of inner structural variation especially characteristic of information-guided control systems.

The first point is fairly obvious and has already been alluded to above. There seems to be a (morally and scientifically) crucial distinction between systems that select actions for reasons and on the basis of acquired knowledge, and other (often highly complex) systems that do not display such goal-oriented behaviors. The image of brain, body and world as a single, densely coupled system threatens to eliminate the idea of purposive agency unless it is combined with some recognition of the special way goals and knowledge figure in the origination of some of our bodily motions[20]. The computational/information-processing approach provides such recognition by embracing a kind of dual-aspect account in which certain inner states and processes act as the vehicles of specific kinds of knowledge and information. The purely dynamical approach, by contrast, seems committed (at best) to a kind of behavior-based story in which the purposive/non-purposive distinction is unpacked in terms of such factors as resistance to environmental perturbation.

The second point builds on the first by noting that total state explanations do not seem to fare well as a means of understanding systems in which complex information flow plays a key role. This is because such systems, as Aaron Sloman has usefully pointed out[21], typically depend upon multiple, 'independently variable, causally interacting sub-states' (*op. cit.*, p. 80). That is to say, the systems support great behavioral flexibility by being able cheaply to alter the inner flow of information in a wide variety of ways. In a standard computer, for example, we find multiple databases, procedures and operations. The real power of the device consists in its ability to rapidly and cheaply reconfigure the way these components interact. For systems such as these the total state model seems curiously unexplanatory. Sloman (op.cit. p.81) notes that:

> a typical modern computer can be thought of as having a [total] state represented by a vector giving the bit-values of all the locations in its memory and in its registers, and all processes in the computer can be thought of in terms of the machine's state space. However, in practice, this [Total State Explanation] has not proved a useful way for software engineers to think ... Rather, it is generally more useful to think of various persisting sub-components (strings, arrays, trees, networks, databases, stored programs) as having their own changing states which interact with one another

The dynamicist may suggest that this is an unfair example, since *of course* a standard computer will reward a standard computational analysis. This, however, is to miss the real point, which is that information-based control systems tend to exhibit a kind of complex articulation in which what matters most is the extent to which component processes may be rapidly de-coupled and re-organized. This kind of articulation has recently been suggested as a pervasive and powerful feature of real neural processing[22]. The fundamental idea is that large amounts of neural machinery are devoted not to the direct control of action but to the trafficking and routing of information within the brain. The point, for present purposes, is that to the extent that neural control systems exhibit such complex and information-based articulation (into multiple independently variable subsystems) the use of total state explanations will tend to obscure the important details, such as the various ways in which sub-state x may vary independently of sub-state y and so on.

The dynamicist may then reply that the dynamical framework really leaves plenty of room for the understanding of such variability. After all, the location in state space can be specified as a vector comprising multiple elements and we may then observe how some elements change while others remain fixed and so on. This is true. But notice the difference between this kind of dynamical approach and the radical, total state vision pursued in section 2. If, as I suspect, the dynamicist is forced to a) give an information-based reading of various systemic substates and processes and b) to attend as much to the details of the inner flow of information as to the evolution of total state over time, then it is unclear that we still confront a real *alternative* to the computational story. Instead, what we seem to end up with is a (very interesting) hybrid: a kind of dynamical computationalism in which the details of the flow of information are *every bit as important* as the larger scale dynamics, and in which some *local* dynamical features lead a double life as elements in an information-processing economy.

This kind of dynamical computationalism is surely attractive. Indeed, it is the norm in many recent treatments that combine the use of dynamical tools with a straightforward acceptance of the notions of internal representation and of neural computation[23]. Nonetheless, such an accommodation is clearly rejected by those, who like van Gelder, depict the dynamical approach as in some deep sense non-computational. It is now time to face this characterization head-on.

4. COMPUTATION & CONTENT.

Three problems make it hard (but not impossible) to map computational ideas on to the dynamical framework. The problems, isolated in section 1, are:

1). The presence of continuous variations in key physical quantities;

2). The presence of complex loops of reciprocal causal influence; and

3). The difficulty of mating rich external influence, described in dynamical and non-computational terms, with a computational account of inner activity.

The mere presence of continuous variation does not, I think, pose any real problem. We may concede that the formal theory of computation (as exemplified in Turing Machine Computationalism) is defined only for discrete state machines, and that the feature of digitality is crucial to many classical results in the theory of computability. But there is also a less formal notion of computation (with an equally impressive historical pedigree in early work on so-called analog computation) which is tied to the much more general idea of automated information processing and semantically sensible transitions between representational elements[24]. A common worry about the proposal to thus 'allow in' analog computation is that such liberality broadens the notion of computation so greatly as to rob it of explanatory force. The worry (see e.g., Sloman *op. cit.*, p. 71) is that everything then turns out to be *some* kind of computer and hence the thesis that the brain computes becomes trivial and uninteresting. This result, however, does not follow. For one thing, what has genuine theoretical bite is never the bald claim that such-and-such a physical device computes. Rather, it is the much more substantive claim that such- and-such (a device or subsystem) computes *some specific function*. Thus[25] even if every object computes *something*, our interest is focused on that special class of computations whose implementations support flexible behavior and reason-guided action. Moreover, given the liberal but intuitive characterization of computation as semantically sensible state-transitions between representations, physical systems that do not traffic in representations (e.g. the solar system, on any intuitive reading) won't count as computational in any case. The threat of trivialization is thus more apparent than real[26].

The analog-friendly notion of computation thus shifts much of the burden to the presence or absence of internal representations. This seems sound enough, since even van Gelder asserts that "manipulable representations lie at the heart of the computational picture" (*ibid.*, p. 351). But there is, alas, a large and complex current debate concerning exactly how we should unpack the notion of internal representation itself[27]. One very plausible constraint, however, is that an internal representation must be a reliably identifiable yet in some sense arbitrary configuration[28] that plays a specifiable semantic role as an intermediary in some problem-solving process, i.e., it must be an item, state or process that we can pick out by physical means but whose functional role is that of *carrying information*.

The presence of continuous variability is clearly compatible with meeting this minimal constraint. But what (moving now to problem number 2) about the presence of complex loops of reciprocal causal influence? Let us assume (for now) that the problematic loop is fully internal and involves some relation of continuous reciprocal causal influence binding the activity of two elements. Would it follow that we could not assign representational roles to the elements? Not at all. It might be, for example, that the coupled activity of the elements is itself the vehicle of some specifiable content. Thus two oscillators, each of which has a distinct, individual frequency of oscillation, can be joined (by continuous reciprocal coupling) so as to constitute a new higher level element with a new and uniform frequency of oscillation. Technically the phase of each oscillator continuously influences the periodic behavior of the other [resulting in] an oscillation frequency for the coupled system[29]. Such regular coupled behavior could *itself* be used as a systemic stand-in for some state of affairs.(For example, a somewhat augmented circuit, known as an adaptive oscillator, can entrain its activity so as to replicate the frequency of an input pulse and to continue the rhythm in the absence of the input[30].) The point, in any case, is just that there is no reason why entire inner loops, involving multiple components in continuous reciprocal causal exchange, should not, at times, themselves constitute the internal vehicles of specifiable contents. Such loops would stand revealed as high-level syntactic items whose functional role is to carry, communicate and transform specific bodies of information; in short, as internal representations implemented in complex dynamics.

Here, then, is a new and fascinating way in which dynamical and computational analyses may proceed hand-in-hand. The dynamical analysis may help identify the complex and temporally extended physical exchanges that act as the *vehicles* of representational content. Traditional computationalism may have been too narrow-minded in its vision of the likely syntactic form of the inner bearers of information and content. Our fascination with the static characters and strings of natural language may have led us to expect simple, local, spatially extended states to function as inner content-bearers. Connectionist approaches helped us see beyond that vision, by identifying the content-bearers as distributed patterns of activity. But it may take the full firepower of dynamical systems theory to reveal the rich space of possible content-bearers: a space that includes very complex temporally and spatially extended processes and interactions.

What, though, of those loops of continuous reciprocal influence that look to criss-cross brain/body/world boundaries? These loops are especially problematic, as we saw in section 1, because they threaten to create tightly coupled systems that incorporate elements (neural, bodily, and environmental) that traditionally require very different kinds of description and explanation. The problem is especially acute when the external influences are described in richly temporal terms whilst the inner economy is depicted as purely computational[31]. This is the 'mating' problem identified as number 3 on our original list.

Once again, the situation looks a lot worse then it really is. First of all, it is simply not the case that temporality must somehow fade away at the point where a computational story takes over. If, as we just conjectured, temporal features of inner processes can play representational roles (and why not? -- temporal duration is a perfectly good non-semantic feature that nature may employ for information-encoding purposes), then the dynamical characterization may be married[32] to the informational and computational accounts. In this way the temporally-characterized inner states and processes are *also* given an information-processing gloss, and this opens the door to viewing them as participants in processes of (perhaps analog) computation.

If it is asked why we should offer this further gloss on the inner events and not on the outer ones, the answer is clear enough. It is that *typically*, the inner events are playing a very different functional role to the outer ones. The difference is two-fold. First, the inner events are acting to control the behavior of a well-defined individual organism whereas the outer events (again only typically) constitute the sources of variance with which the controller must cope. And second, the control in question is special insofar as it is information-and-goal based. Regarding the first point, we can see that even the Watt governor plays a different functional role to the engine, for it is the *job* of the governor to keep the engine speed constant, whereas there is no plausible teleological story in which the reverse is true. Governors evolved to control engines, and this fact alone marks an important difference that is unaffected by the dense continuous and reciprocal nature of the causal exchanges that bind the governor to the engine[33]. The Watt governor, however, is the kind of control system that does not in addition require us to invoke an information-based account of its control strategies. Instead, we gain a full and sufficient understanding of its operation by focusing on physical laws relating physical quantities. This is not the case in control systems that, for example, create and manipulate models of the domain with which they deal, or ones that maintain hierarchies of goals and sub-goals, reason about what action to perform next, and so on.

To simply assert that human intelligence involves the use of inner models and representations would, of course, be to beg the question at issue. The point is rather that even if the leading skeptical consideration (concerning relations of continuous reciprocal causal influence that criss-cross brain, body, world boundaries) is correct, this is *in itself* no reason to suppose that we cannot motivate talk of inner models, representations, computations and all the rest. In so doing we opt for a distinctive vocabulary and explanatory style that highlights the special functional role of certain events as control structures and that reveals the control function as based on the extraction, preservation and use of specific types of information[34].

Next recall the argument[35] that the complexities of certain organism/environment interactions may (courtesy of our old friend continuous reciprocal causation) be 'so subtle and complex as to defy description in representational terms' (Van Gelder, *op. cit.,* p. 381). One immediate reaction is that even here the case remains unproven. Even in the example of the on-going tennis game, it is by no means obvious that there are not distinct inner states playing specific representational roles that may be identified by future neuroscientific investigations. Moreover, even if the temporal complexities of the interactions cannot be directly captured using talk of representations and information-flow, it does not follow that we should reject such descriptions outright. It is quite possible, for example, that the pure dynamical description is *equally* inadequate. If it highlights physical exchanges at the expense of revealing the fine structure of goals, strategies and desires that help modulate the play, it too will leave out much that is of interest. The point about the possible super-representational status of the agent/environment relation gains its force from an implicit assumption to the effect that one kind of story (either dynamical or representational) ought to tell us *all* that we want to know about the agent-environment relation in question. I suggest, however, that this assumption is unwarranted. It is perhaps fueled by the use of the Watt governor illustration in which the dynamical account really does seem to tell us all we need to know. But this is not a demonstration we can trust since the governor does not produce the kinds of highly flexible, rapidly re-configurable, multiple goal and task oriented behaviors so characteristic of cognitive agents. The governing task is thus not 'hungry' for the kind of information-and-representation based control strategies that seems especially illuminating in cases where the behavior is much more highly flexible and depends on the maintenance of many bodies of information and complex goal structures. In such cases it is surely implausible to expect any model that eschews talk of information and representation to satisfy us. Instead we should pursue a hybrid approach that both tracks the complex interactive dynamics and makes some attempt to reveal the information processing roles of component states and processes.

But there is, in any case, one last reason why we should not be overly troubled by the spectre of complex causal chains that criss-cross the agent/environment boundary. It is that such complexity is not characteristic of most of the (intuitively) core cases of genuinely *cognitive* phenomena. For both historically and intuitively the very idea of cognition has been tied to (a) the idea of reasoned behaviors selected in the absence of any local guiding stimulus (as in advance planning, counterfactual reasoning etc.), and (b) to behaviors that seem to track non-nomic stimulus properties (such as 'being a movie ticket' - a property that does not figure in any natural laws but one to which we can reliably respond nevertheless (see e.g. Fodor, J. Psychosemantics (MIT Press, Cambridge, MA, 1986), p.14)). The idea of cognition has been tied, that is to say, to the idea of behaviors carried out in the absence of any constant, lawful and reliable signal from the local environment[36]. This is why a light-following robot does not strike us as a paradigmatically cognitive engine, whereas a robot that engages in off-line planning and abstract problem-solving does.

It would be unfair to make too much of this, since it is an admirable part of the agenda of the dynamicist to combat the 'passive, reflective' model of cognition that has so deeply permeated traditional philosophical and scientific thought. Nonetheless, it is surely fair to comment that it was the project of explaining both environmentally detached reasoning and consistent response to non-nomic stimulus properties that originally led theorists to posit inner models and manipulable representations. When, for example, an agent reasons about an event in advance, the theorist is led to ask what structural surrogates the system is using so as to select appropriate actions in the absence of actual local triggers and cues. And when an agent reasons about movie tickets, the theorist is led to ask how a physical system can possibly key it's responses to a feature of the world that does not fall (*qua* that feature) under the laws of physics.

In sum, the worries about temporality and complex causal exchanges look, on closer examination, unable to bear the kind of weight required by recent skeptical arguments. It is by no means clear that pure dynamical stories can even in principle do justice to the special requirements imposed by the project of explaining reason-guided action. Nor is it clear that the kinds of case to which the skeptical considerations most plausibly apply are really on a par with the cases for which the target notions (of internal representation and computation) were originally invoked.

5. CONCLUSIONS: ACCOMMODATION NOT ELIMINATION.

Suppose we accept that traditional computational models *have* tended to screen off a little too much of the real physical and temporal structure of events, and *have* tended to erect somewhat too firm a barrier between the cognitive system and the rest of the world. One response to these excesses is to reject outright the notion that cognition is a matter of computation, and with it the whole apparatus of internal representation and information-processing theory. But this, I have argued, is a distortive and even dangerous overreaction. It is distortive insofar as there is nothing in the detailed considerations and arguments put forward (concerning the importance of timing, and the possibility of continuous reciprocal causal influence) that is in any way inimical to a substantive computational and informational analysis. And it is dangerous insofar as it threatens to deprive us of the special kinds of tool and understanding needed to make sense of the special class of information-based control systems.

What is indicated is thus not the rejection of the computational/representational vision so much as the addition of an irreducibly dynamical dimension to the analysis. Such a dimension manifests itself in several ways, including: the use of dynamical tools to recover potential information-bearing vehicles from highly complex webs of causal exchange; the recognition that intrinsically temporal features may play representational and computational roles; and the extension of computational ideas to systems that change continuously in time and that exploit continuous state.

Accommodating this dynamical dimension clearly calls for a substantial broadening of the standard computationalist outlook. But as long as we maintain the distinctive and explanatorily potent physical/informational dual aspect that is at the heart of computational framework, we confront (I suggest) only refinement and progress, not outright elimination. The future of cognitive science may thus lie in a delicately negotiated union between the familiar framework of computational, representational and information-processing description, and the challenging and temporally charged project of dynamical analysis. For such a union to succeed, the dynamicist must look beyond total state explanations, attending as much to the details of mechanism and information-flow as to the shape of overall systemic unfolding. While the computationalist must recognize the large and baroque space of physical and temporal processes that nature may plunder for representational gain.

Andy Clark 1997

- Thanks to Tim Van Gelder, Randy Beer, Esther Thelen and Michael Wheeler for numerous discussions and exchanges on all these topics. Thanks also to friends, audiences and colleagues at the University of Sussex, UK, the Santa Fe Institute, and the Philosophy-Neuroscience-Psychology Program at Washington University in St. Louis.

Abraham, R., & Shaw, C. (1992). Dynamics -- The Geometry of Behavior. Redwood, CA: Addision-Wesley.

Ashby. (1956).

Beer, R. (1995). A Dynamical Systems Perspective on Environment Agent Interactions. Artificial Intelligence, 72, 173-215.

Beer, R., & Gallagher, J. C. (1992). Evolving dynamical neural networks for adaptive behavior. Adaptive Behavior, 1, 91-122.

Busemeyer, J., & Townsend, J. (to appear). Decision Field Theory. In R. Port & T. van Gelder (Eds.), Mind as Motion: Dynamics, Behavior and Cognition . Cambridge, MA: MIT Press.

Chalmers, D. (1994). On implementing a computation. Minds and Machines, 4, 391-402.

Churchland, P. M. (1989). The Neurocomputational Perspective. Cambridge: MIT/Bradford Books.

Clark, A. (1989). Microcognition: Philosophy, Cognitive Science and Parallel Distributed Processing. Cambridge: MIT Press.

Clark, A. (1993). Associative Engines: Connectionism, Concepts and Representational Change. Cambridge: MIT Press.

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

Clark, A. (in progress). Re-Thinking Innateness: A Critical Study.

Clark, A., & Chalmers, D. (1995). The Extended Mind (PNP Research Report). St. Louis: Washington University.

Clark, A., & Grush, R. (1997). Towards a Cognitive Robotics (PNP Technical Report). St. Louis: Washington University in St. Louis.

Clark, A., & Toribio, J. (1994). Doing Without Representing? Synthese.

Crutchfield, & Mitchell. (1995). .

Dennett, D. (1978). Brainstorms. Sussex: Harvester Press.

Elman (1996).

Giunti, M. (1995). Dynamical Models of Cognition. In R. Port & T. van Gelder (Eds.), Mind as Motion (pp. 549-571). Cambridge, MA: MIT Press.

Goodwin, B. (1994). How the Leopard Changed its Spots. London: Weidenfeld & Nicolson.

Hardcastle, V. (1995). .

Harnad, S. (1994). What is Computation? Minds and Machines, 4(4), 377-488.

Haugeland, J. (1995). Mind Embodied and Embedded. In Y.-H. Houng & J.-C. Ho (Eds.), Mind and Cognition (pp. 3-38). Taipei, Taiwan: Academia Sinica.

Hutchins, E. (1995). Cognition in the Wild. Cambridge, MA: MIT Press.

Keijzer, F. & Bem, S. (1996). Behavioral Systems Interpreted as Autonomous Agents and as Coupled Dynamical Systems: A Criticism. Philosophical Psychology, 9, 323-46.

Kelso. (19xx). .

McClelland, J., Rumelhart, D., & Group, P. R. (1986). Parallel Distributed Processing: Explorations in the Microstructure of Cognition. (Vol. I & II). Cambridge: MIT Press.

Merleau-Ponty, M. (1942). La Structure du Comportment (A. Fisher, Trans.). France: Presses Universitaites de France.

Mitchell. (1994). .

Norton, A. (1995). Dynamics: An Introduction. In R. Port & T. Van Gelder (Eds.), Mind as Motion: Dynamics, Behavior, and Cognition . Cambridge, MA: MIT Press.

Port, R., Cummins, F., & McCauley, J. (1995). Naive Time, Temporal Patterns and Human Audition. In R. Port & T. Van Gelder (Eds.), Mind as Motion . Cambridge, MA: MIT Press.

Port, R. & van Gelder, T. (1995). Mind as Motion: Dynmaics, Behavior, and Cognition. Cambridge, MA: MIT Press.

Salzman, L. & Newsome, W. (1994). Neural Mechanisms for Forming a Perceptual Decision. Science, 264, 231-237.

Sloman, A. (1993). The Mind as a Control System. In C. Hookway & D. Peterson (Eds.), Philosophy and Cognitive Science: Royal Institute of Philosophy Supplement 34 (pp. 69-110). Cambridge, UK: Cambridge University Press.

Smith, B. C. (1996). On the Origin of Objects. Cambridge, MA: MIT Press.

Smolensky, P. (1988). On the Proper Treatment of Connectionism. Behavioral and Brain Sciences, 11, 1-74.

Sterelny, K. (1995). Understanding Life: Recent Work in Philosophy of Biology. British Journal for the Philosophy of Science, 46(2), 155-183.

Thelen, E. (1993).

Thelen, E., & Smith, L. (1994). A Dynamic Systems Approach to the Development of Cognition and Action. Cambridge, MA: MIT Press.

Turing. (1936). On Computable Numbers, with an Application to the Entscheidung Problem. Proceedings of the London Mathematical Society, 2(42), 230-265.

Van Essen, D., & Gallant, J. (1994). Neural Mechanisms of Form and Motion Processing in the Primate Visual System. Neuron, 13, 1-10.

van Gelder, T. (1995). What Might Cognition Be, If Not Computation? Journal of Philosophy, XCII(7), 345-381.

van Gelder, T., & Port, R. (1995). It's About Time: An Overview of the Dynamical Approach to Cognition. In R. Port & T. v. Gelder (Eds.), Mind as Motion: Explorations in the Dynamics of Cognition (pp. 1-44). Cambridge, MA: MIT Press.

Varela, F., Thompson, E., & Rosch, E. (1991). The Embodied Mind. Cambridge, MA: MIT Press.

Wheeler, M. (1994). From Activation to Activity. Artificial Intelligence and the Simulation of Behavior (AISB) Quarterly, 87, 36-42.

---

[1] See van Gelder, T. (1995). What Might Cognition Be, If Not Computation? Journal of Philosophy, XCII(7), 345-381, van Gelder, T., & Port, R. (1995). It's About Time: An Overview of the Dynamical Approach to Cognition. In R. Port & T. v. Gelder (Eds.), Mind as Motion: Explorations in the Dynamics of Cognition (pp. 1-44). Cambridge, MA: MIT Press. Related arguments are found in Thelen, E., & Smith, L. (1994). A Dynamic Systems Approach to the Development of Cognition and Action. Cambridge, MA: MIT Press, Kelso, S. (1995) Dynamic Patterns Cambridge, MA: MIT Press, Varela, F., Thompson, E., & Rosch, E. (1991). The Embodied Mind. Cambridge, MA: MIT Press and in Wheeler, M. (1994). From Activation to Activity. Artificial Intelligence and the Simulation of Behavior (AISB) Quarterly, 87, 36-42.

[2] van Gelder, T., & Port, R., *op. cit.* p. 18-23, van Gelder, T. *op. cit.* p. 354, 379.

---

[3] See, e.g.,Turing, A. (1936). On Computable Numbers, with an Application to the Entscheidung Problem. Proceedings of the London Mathematical Society, 2(42), 230-265, Sloman, A. (1993). The Mind as a Control System. In C. Hookway & D. Peterson (Eds.), Philosophy and Cognitive Science: Royal Institute of Philosophy Supplement 34 (pp. 69-110). Cambridge, UK: Cambridge University Press, and Giunti, M. (1995). Dynamical Models of Cognition. In R. Port & T. van Gelder (Eds.), Mind as Motion (pp. 549-571). Cambridge, MA: MIT Press.

[4] In Clark, A. (1997). Being There: Putting Brain, Body, and World Together Again. (Ch. 8) Cambridge, MA: MIT Press.

[5] E.g., Van Essen, D., & Gallant, J. (1994). Neural Mechanisms of Form and Motion Processing in the Primate Visual System. Neuron, 13, 1-10.

[6] Ashby, R (1956) Introduction to Cybernetics. Wiley, UK p. 54.

[7] This complex interactive dance is also stressed by Varela, Thompson,& Rosch *op. cit.* and is central to Varela's notion of cognition as 'enaction'.

[8] See Clark, A., & Toribio, J. (1995). Doing Without Representing? Synthese 101: 401-431 for discussion. I propose to grant that the Watt Governor is, in some fairly intuitive sense, a non-computational device insofar as we do not get our best explanatory grip on its operation by means of its information-processing profile.

[9] In addition, note that real-timing is crucial to these interactions in a way that it *isn t* crucial to computational operations. The Watt Governor, unlike the computational one, includes no components whose brief is simply to calculate x and output a result within so many milliseconds (see van Gelder, T. *op. cit.* p. 354). Instead, there is a constant temporal coupling in which change in various components is continuously coordinated in real-time. In this sense, real-timing matters more in the dynamical solution.

[10] See Dennett, D. (1978). Brainstorms. Sussex: Harvester Press.

[11] See McClelland, J., Rumelhart, D., & the PDP Research Group, (1986). Parallel Distributed Processing: Explorations in the Microstructure of Cognition. (Vol. I & II). Cambridge: MIT Press, Clark, A. (1989). Microcognition: Philosophy, Cognitive Science and Parallel Distributed Processing. Cambridge: MIT Press, Clark, A. (1993). Associative Engines: Connectionism, Concepts and Representational Change. Cambridge: MIT Press, Churchland, P. M. (1989). The Neurocomputational Perspective. Cambridge: MIT/Bradford Books.

[12] For similar comments, see Thelen, E., & Smith, L. *op. cit.* (p. 17) Thelen, E. (1993), p. 580, Varela, Thompson, & Rosch *op. cit.* p. 172-180.

[13] See Norton, A. (1995). Dynamics: An Introduction. In R. Port & T. Van Gelder (Eds.), Mind as Motion: Dynamics, Behavior, and Cognition, Kelso, S. *op. cit.*, and Abraham & Shaw, .*op. cit.* for useful introductions.

[14] See Salzman, L. & Newsome, W. (1994). Neural Mechanisms for Forming a Perceptual Decision. Science, 264, 231-237.

[15] See, e.g., van Gelder, T. *op. cit.* 345-381, Beer, R., & Gallagher, J. C. (1992). Evolving dynamical neural networks for adaptive behavior. Adaptive Behavior, 1, 91-122, Wheeler, M. *op. cit.* p.36-42. For discussion see Keijzer, F. & Bem, S. (1996). Behavioral Systems Interpreted as Autonomous Agents and as Coupled Dynamical Systems: A Criticism. Philosophical Psychology, 9, 323-46, Clark & Toribio *op. cit.* and Clark, A. & Grush, R. (To appear). Towards a Cognitive Robotics. Adaptive Behavior.

[16] For discussion, see Clark, A. (1997). Being There: Putting Brain, Body, and World Together Again. (Ch. 6) Cambridge, MA: MIT Press.

---

[17] For further mathematical detail, see van Gelder, T. *op. cit.*, p. 356-357.

[18] For the last two cases, see Goodwin, B. (1994). How the Leopard Changed its Spots. (p. 60) London: Weidenfeld & Nicolson.

[19] E.g., van Gelder's comments (*op. cit.*, p. 358) on tasks that may only initially appear to require 'that the system have knowledge of and reason about, its environment', and Thelen & Smith's (*op. cit.*, p. xix) stress on the brain as a thermodynamic system. By contrast, the dynamicist Scott Kelso (*op. cit.*, p. 288) sees the key problem as 'how *information* is to be conceived in living things, in general, and the brain in particular'.

[20] For a similar argument, see Keijzer & Bem *op. cit.* p. 339.

[21] See Sloman, A. (1993). The Mind as a Control System. In C. Hookway & D. Peterson (Eds.), Philosophy and Cognitive Science: Royal Institute of Philosophy Supplement 34 (pp. 69-110). Cambridge, UK: Cambridge University Press.

[22] E.g., the 'gating' hypothesis in Van Essen, D., & Gallant, J. (1994). Neural Mechanisms of Form and Motion Processing in the Primate Visual System. Neuron, 13, 1-10.

[23] See e.g., Smolensky, P. (1988). On the Proper Treatment of Connectionism. Behavioral and Brain Sciences, 11, 1-74 and Elman, J (1995) Language as a Dynamical System, in R. Port and T. van Gelder (eds) (*op. cit*., p. 195-226).

[24] For example, MacLennan, B. (1990) (Field Computation: A Theoretical Framework for Massively Parallel Analog Computation. University of Tennessee Technical Report CS-90-100)describes a continuous computational system which is in all other respects similar to a traditional computational system.

[25] This point is due to Chalmers, D. (1994). On implementing a computation. Minds and Machines, 4, 391-402.

[26] For further discussion, see Smith, B. C. (1996). On the Origin of Objects. Cambridge, MA: MIT Press, Clark, A. (1997). *op.cit.*, Chapter 8, Sloman, A. (1993). *op.cit.*, Chalmers, D. *op. cit.* p. 391-402, and papers in Harnad, S. (1994) (ed) What is Computation? Special Issue of Minds and Machines, 4(4), 377-488.

[27] See e.g., Clark, A *op. cit.* Ch. 8, van Gelder, T. *op. cit.* 345-381, Beer, R. (1995). A Dynamical Systems Perspective on Environment Agent Interactions. Artificial Intelligence, 72, 173-215, Wheeler, M. *op. cit.* pp. 36-42.

[28] This term is to be understood broadly so as to include physical states, virtual states and temporally extended processes; i.e., any identifiable feature or pattern capable of entering into causal relations.

[29] See Port, R., Cummins, F., & McCauley, J. (1995). Naive Time, Temporal Patterns and Human Audition p 361. In R. Port & T. Van Gelder (Eds.), *op. cit*., p. 339-372.

[30] See Port, R., Cummins, F., & McCauley, J. *Op. cit.* p 361-363

[31] See van Gelder, T., & Port, R. (1995). It's About Time: An Overview of the Dynamical Approach to Cognition. In R. Port & T. v. Gelder (Eds.), Mind as Motion: Explorations in the Dynamics of Cognition (pp. 1-44). Cambridge, MA: MIT Press.

[32] Just such a union is pursued in Crutchfield, J & Mitchell, M (1995) The Evolution of Emergent Computation Proceedings of The National Academy of Science 92: 10742-10746 and in Mitchell, M., Crutchfield, J., and Hraber, P. (1994) Evolving Cellular Automata to Perform Computations Physica D 75: 361-391. It also characterizes several papers in R. Port & T. Van Gelder (Eds.) (1995), *op. cit.*. Van Gelder s own notion of revisionary representationalism and his discussion of Decision Field Theory (van Gelder, T. op. cit. p. 359-363) show that he is open to the idea of such a union of dynamics and

---

representationalism, though troubled by the time-based considerations highlighted in the present treatment.

[33] A parallel argument aimed to distinguish genetic control from environmental backdrop is developed in Sterelny, K. (1995). Understanding Life: Recent Work in Philosophy of Biology. British Journal for the Philosophy of Science, 46(2), 155-183.

[34] For further argument to this effect, see Sloman, A. *op. cit.* (pp. 69-110), Clark, A., & Grush, R. *op. cit.*. Notice, in passing, an important wrinkle. By giving prominence to the notion of a distinctive kind of functional role (viz, control by information-and-representation) we allow that some *external* states and processes might indeed count as (perhaps temporary and transient) parts of the cognitive control system itself. What will determine whether specific external processes and events count as part of an extended cognitive system is not, however, the presence or absence of continuous reciprocal causal interactions. For this, we now see, is orthogonal to the real issue, which is whether the external events participate in the encoding, maintenance or transformation of information used to control intelligence action. Many external structures play --for humankind -- just such a role. Think of abacuses, computers, books, ships compasses, maps, etc., etc.. Such items are selected to function as elements in environmentally extended processes of computation and information-based control. In these cases, the inner and the outer do, I think, form extended computational systems that reward treatment as integrated, information-processing wholes. For further argument see Hutchins, E. (1995). Cognition in the Wild. Cambridge, MA: MIT Press, Clark, A. (1997). Being There: Putting Brain, Body, and World Together Again. (Ch. 9, 10) Cambridge, MA: MIT Press. It is, of course, a further and extremely vexing question when (if ever) we should count such external structures as *part of the agent* herself -- for some discussion, see Clark, A., & Chalmers, D. (In press). The Extended Mind. Analysis.

[35] See van Gelder, T. *op. cit.* p 345-381 and section 2 above.

[36] See, e.g., the definition of internal representation in Haugeland, J. (1995). Mind Embodied and Embedded. In Y.-H. Houng & J.-C. Ho (Eds.), Mind and Cognition (pp. 3-38). Taipei, Taiwan: Academia Sinica, and discussion in Smith, B. C. (1996). On the Origin of Objects. Cambridge, MA: MIT Press.