Proceedings of the 2007 IEEE Symposium on
Artificial Life (CI-ALife 2007)

A Critical Reappraisal of the Dynamical Approach to Cognition

S. McGregor, N. Virgo and J. Dyke

***Abstract*— Approaches to cognitive science have been socially divided into *dynamical* and *computational* camps. We break down the dynamical approach into finer components, suggesting a new taxonomy of dynamical approaches to cognition and questioning the logical unity of the dynamical school. We dispel some confusions surrounding the concepts of dynamical systems, computation, and the relation between the two. We introduce and argue for the notion of "cognition as it could be" and show its value in analysing the dynamicists' account of time.**

I. Introduction

Many researchers in Artificial Life identify themselves as being part of a "Dynamical Approach" to cognition, usually contrasted with a "Computational Approach". The dynamical approach conceives of cognitive processes in terms of dynamical systems; the computational approach in terms of digital computers. Unfortunately not just one but *both* of these terms ("dynamical system" and "digital computer") are used with multiple meanings by different authors, which inevitably leads to confusion.

Our position is that the dynamical approach is a "cluster concept" with overlapping semi-independent characteristics none of which is necessary or sufficient to define the dynamical approach. Depending on context, the phrase "dynamical approach" can informally mean

- "I consider a process occurring over time"
- "I consider a process occurring in continuous time"
- "I use a model with floating-point parameters"
- "I evolved a Continuous-Time Recurrent Neural Network (CTRNN)"
- "I don't subscribe to the physical symbol system hypothesis"
- "I use a situated and/or embodied approach"

The lack of a single defining feature of the dynamical approach is most obvious when one tries to define it positively rather than negatively - that is, to characterise specific properties which are shared by studies described as having a dynamical approach, rather than contrasting them with computational ones. This paper attempts to clarify the relationships between different aspects of the dynamical approach, expose some common unwarranted assumptions and clear up some confusion about the relationship between the dynamical approach and the use of computer simulation.

We begin with a brief tongue-in-cheek overview of the dynamicist and computationalist research traditions. Then we catalogue current usages of the terms "dynamical system" and "digital computer" and suggest a finer-grained vocabulary. The important notion of *situatedness* is discussed next. After that we consider what sort of environment cognitive agents are assumed to operate in, and the notion of "cognition as it could be" is contrasted with cognition in real physical systems.

A longer section deals with the properties of the specific class of dynamical system which dynamicists typically advocate. We examine different possible reasons for favouring these systems and by introducing the concept of "continuous-like" behaviour in discrete systems we hopefully clear up some of the relationship between continuous systems, discrete models of continuous systems and discrete systems.

Time has been advanced as playing a defining role in the dynamical approach, but the dynamicist conception of time turns out to be problematic. Our final section addresses the difference between continuous time, discrete time, and "mere sequence" in a way which highlights the central importance of situatedness to time.

II. Computationalists and Dynamicists

The following folk tale is a broad but not wholly inaccurate caricature of how self-proclaimed dynamicists see the difference between themselves and the people they call computationalists.

> *Once upon a time, there were some scientists who thought the mind was a computer. They tried to build robots which would reason before they moved around. They forgot what the brain was made of and what sort of body it was in. They didn't notice that the world itself was helping people and animals to be smart.*
>
> *Then some of them had a clever idea. What if agents were a part of the physical environments they lived in, just like the planets, and chemical reactions, and electrical circuits, and pendulums? The scientists would be able to use ordinary science to understand them.*
>
> *So they tried building robots which just moved around without planning anything. They tried making models which used ordinary sensible numbers. And guess what? It worked!*

Like most folk tales, this narrative is more about culture than it is about history. The computationalists and dynamicists presently occupy more or less socially distinct research communities in cognitive science and artificial intelligence. They use different research tools, have different research ideologies and concentrate on broadly different problems. So far, in respect to engineering intelligent systems, the two approaches have both been successful but in rather different domains. Computationalist systems can beat the human world champion at chess [1]; dynamicist systems provide the state of the art in robust robot control [2].

A. The "Dynamical Hypothesis"

Although its author is no longer active within the dynamicist community, perhaps the most influential attempt to define and argue for the dynamical approach is [3], in which two separate claims are identified. One, the "Nature Hypothesis," is a metaphysical claim that relies on some peculiar terminology to claim that cognitive agents in some sense "are" the abstract dynamical systems that could be used to model them. We will not discuss this here. The other, the "Knowledge Hypothesis," is the pragmatic claim that dynamical systems models can and should be used in cognitive science. This is not strictly a hypothesis but a sort of scientific meta-theory: a theory about what general form a successful scientific theory of cognition will take [4], [5].

Not all dynamicists would support the Dynamical Hypothesis as stated in [3], and it has been widely criticised on a number of grounds [6], [5] which we will not attempt to summarise or address here. We are concerned with classifying the different elements of the dynamical approach, rather than offering any new criticisms or defences of its ideology.

III. Dynamical Systems

A. "Pure" Dynamical Systems

The technical term *dynamical system* has several meanings in current usage, and it is worth differentiating between them. To a mathematician it has a precise definition which is much broader than any normally found in cognitive science since it is designed to capture the notion of *anything* that changes over (discrete or continuous) time. In the branch of pure mathematics known as *dynamical systems theory* a dynamical system is simply a *phase space* or *state space*, which can be any sort of set, a *time space*, which has to support an addition operator – typically the real numbers or the integers – and an *evolution function* which describes how the state of the system varies with time. There are just two axioms, which essentially say that the system doesn't change over zero time and the evolution function is uniformly applied at all points in time.

> The most characteristic feature of dynamical theories, which distinguishes them from other areas of mathematics dealing with groups of automorphisms of various mathematical structures, is the emphasis on... properties related to the behaviour as time goes to infinity. [7] (p2)

There is not all that much which can be said about mathematical dynamical systems in their most general form. [7] identify four main subdomains in dynamical systems theory which make different assumptions about the structure of the phase space and the evolution operator: *ergodic theory*, *topological dynamics*, *smooth dynamical systems*, and *Hamiltonian dynamics*.

In this article we use the term "dynamical systems theory" in the broadest mathematical sense, including but not limited to its various "pure" or "applied" subdisciplines.

B. "Applied" Dynamical Systems: Numerical Systems

Many of the notions of mathematical dynamical systems theory are still too abstract for most scientists and engineers, whose objects of study can usually be defined using real number variables and modelled using differential (or difference) equations. Systems of this sort are typically easier for scientists to visualise, model and analyse than abstract topological or measurable spaces. They are also the main focus of any introductory course in dynamical systems theory. See [8] for a well-regarded text which deals mainly with this class of dynamical system. When people refer to dynamical systems in the context of the physical and life sciences they very often mean these sorts of system, which from a mathematical point of view constitute only a tiny subset of all dynamical systems.

C. "Applied" Dynamical Systems: Non-Numerical Systems

The field of Artificial Life is somewhat unusual in that dynamical systems whose states aren't interpreted numerically are fairly common. We are likely to be interested in analysing the dynamics of discrete systems such as cellular automata or random Boolean networks as well as continuous systems. An example of discrete-state research in the dynamical cognitive science tradition is [9]. However, discrete dynamical systems are also relevant to those studying continuous nonlinear dynamical systems in the physical sciences: it is sometimes useful to analyse the behaviour of such systems by discretising them, even when the trajectories of the system do not resemble discrete jumps. This trick, invented by Poincare, was the original motivation for the dynamical systems theory subdiscipline of *Symbolic Dynamics* [10].

D. Dynamical Systems in the Dynamical Approach

Cognitive studies which might be broadly described as dynamicist differ greatly in their use of the tools or language of dynamical systems theory[1]. Although authors such as Beer [12] and van Gelder [3] are generally careful to use the term "dynamical" when talking about the conceptual ideology, reserving the term "dynamical systems" for the mathematical theory, this usage is not universally followed and some authors blur the distinction. For instance, the Dynamical Hypothesis is referred to as the "Dynamical Systems Hypothesis" in [13] whilst the "Dynamical Systems Approach" in [14] is defined primarily in terms of how it differs from computationalism. We believe this usage to be a misnomer and suggest the following finer distinctions between types of study:

1) Work which emphasises the role of timing of interactions with the environment. We suggest the term *dynamical* or preferably *time-critical* for work of this sort.
2) Work which simply uses a numerical differential-equation, or difference-equation, model. Models of this

[1]Many of the mathematical notions underlying dynamical systems theory, although not the modern terminology, were applied to cognition and behaviour by the cybernetics movement. For example, see [11]. Cybernetics is widely seen as a precursor to modern dynamicism.

sort are ubiquitous in quantitative science and are never described as dynamical systems models outside of cognitive science. For instance, Newtonian mechanics would never be referred to as a "dynamical systems approach" to physics. We suggest the term *numerical* or *quantitative* to describe such models.
3) Work which makes specific use of concepts or tools from dynamical systems theory, such as attractors, phase portraits, bifurcations, structural stability, metric entropy or conjugacy. We suggest that the term *dynamical systems* be used only in work of this sort. In fact, as observed in [15], dynamicist models tend to operate far outside dynamic equilibrium, which makes it difficult to apply the mathematics.
4) Work which makes use of a model which could formally be described as a dyamical system (i.e. in terms of a state space, time set and evolution function). This category is far too broad to deserve any particular name, covering all work in both computationalist and dynamicist camps.

Note that in work done so far, these categories tend to overlap in a hierarchical manner. Dynamical systems models (in our sense) tend to be quantitative, and quantitative models tend to be time-critical. But the conceptual categories are all in fact logically independent. One could use dynamical systems analysis of a model which was time-critical but not quantitative. For instance, one might want to use dynamical systems theory in understanding the behaviour of a robot controlled by a finite state machine (FSM) in a fast-changing grid world, or to consider convergence in offline neural network learning (quantitative but not time-critical).

E. What Is Dynamical Systems Theory a Theory Of?

The occurrence of the word 'theory' in the phrase 'dynamical systems theory' suggests to some casual readers that dynamical systems theory is, or could be, a scientific theory of cognition. In fact, it is a theory only in the mathematical sense of a set of axioms, proofs, methods and tools. (The major exponents of the dynamical approach never make this confusion.)

IV. The Computational Hypothesis

Perhaps much of the confusion about the meaning of the phrase "dynamical systems approach" can be attributed to a lack of clarity in defining the computationalist approach that it is formulated in opposition to. The dynamical approach is pitted in [3] against the "Computational Hypothesis" that "cognitive agents are basically digital computers." ([3], p615.), echoing the statement in [16] (p169) that "[t]he popular metaphor of calling the brain an 'information-processing device' is... patently wrong." The phrase "information-processing device" or "digital computer" here is a somewhat overloaded one. It appears to have slightly different meanings to different people (within both the dynamical and computationalist traditions), and can more or less strongly imply several distinct features. It is also worth noting that, perhaps somewhat counterinuitively, "digital computer" is usually meant to imply a more or less abstract entity rather than an actual computational device such as a desktop computer. Use of the term can imply at least the following four aspects:

1) Having distinct input and output phases, as in the traditional sense-model-plan-act (SMPA) architecture [17]. Van Gelder ascribes this meaning to the term "computer," so that anything which gives outputs that are a function of its inputs can be considered a computer. Note that this does not necessarily apply to practical computers in the real world such as in industrial control applications, or even to a laptop interacting with its user.
2) Having discrete rather than continuous input, output and internal state sets. That is, the possible inputs to and outputs from a digital computer come from a discrete set, such as the natural numbers or the set of strings in a given character set. This does not apply to physical analogue computers.
3) Being symbolically representational. That is, a digital computer is considered to operate in a logical fashion on discrete symbols which represent well-defined things. It is not enough that the sets of possible inputs and outputs be discrete: by "symbolically representational" we mean that some meaning has been ascribed to the system's input, output and internal states, as part of their definition. The operation of a digital computer is taken to consist of the manipulation of symbols in such a way that the system's mechanism respects the pre-defined syntax and semantics of the symbols. (This is the aspect of computation that Maturana and Varela [16], [18] are referring to when they state that the nervous system does not operate like a computer.)
4) Computability: the idea that a digital computer can only calculate functions that could be calculated by some Turing machine, regardless of any meaning that may be ascribed to that machine's internal states. It is often not clear whether this is intended as part of the definition of a digital computer when talking about cognition but it is fundamental to the meaning of the word "computer" in computer science.

Note that we are using a very strict definition of symbolic representation in criterion 3. There are some approaches to symbols and representation from within the dynamical school, e.g. [19], [20], [21]. These differ from the symbolic representation that we are describing here because the meaning of the symbols is not pre-defined by the experimenter and must instead be acquired by the agent.

It is important to note the logical independence of criteria 3 and 4. The phrase "digital computer" or even "computational model" can imply either or both, and confusion over this dogs explanations of the dynamical approach, particularly when addressing the use of computers in dynamical research such as in evolutionary robotics. A computer which is simulating a dynamical system in an evolutionary robotics context is operating on a symbolic representation of the dynamical system — that is, it is generally operating on floating point numbers which act as symbols representing an agent's internal state and that of its environment — but this is not the *agent's* representation of anything, because in evolutionary robotics tasks are not specified in such a way that the agent's behaviour has to respect any defined semantics. These floating point numbers change over time according to dynamical rules: the agent's state is just changing over time and it cannot necessarily be said to be doing logical operations on symbols, even though the computer simulating it is.

When the dynamical approach is contrasted with computational approaches, the word "computational" can take on various combinations of the above meanings. Most dynamical models do not have distinct input and output phases, being time-critical (see section III-D above), and sometimes this is all that is meant. Sometimes symbolic representation is also implied in the use of the word "computational." Dynamicists' models are not symbolically representational[2], but they share this quality with non-dynamical connectionist models such as feed-forward neural networks [22].

It is often also felt that operating on continuous input/output or state sets is important for the dynamical approach. This is certainly true given the flavour of the models that tend to be used but it is important to note that it is not necessary for a system to operate on continuous states in order to be time-critical or non-symbolically representational. A connectionist example of a system that operates on a discrete state space but is not symbolically representational is the Hopfield network [23], a recurrent neural network whose nodes can be either on or off and which updates in discrete time steps. Indeed [24] is an example of a self-described "Dynamical Systems" approach which incorporates a Hopfield network in its model.

The one aspect of the word "computational" which is not generally implied when contrasting it with a dynamical approach is the notion of computability. [5] observes that some continuous systems can perform tasks that cannot be emulated by a Turing-equivalent system but that no researcher has this as a motivation for advocating the dynamical approach.[3]

V. Situated cognition

It has been shown that for a variety of interesting behaviours, both simulated and real-world agents take advantage of ongoing interaction with their environment. The traditional computationalist view, exemplified by the sense-model-plan-act model, was that cognitive processes essentially had distinct input, processing and output stages.

As rightly pointed out by dynamicists, this is an extremely poor basis for explaining real-time real-world behaviours such as obstacle avoidance or walking. In [26], which is claimed by the dynamicists as part of their tradition, the anthropologist Lucy Suchman's term *situated* [27] is used to describe agents which are in ongoing interaction with a relevant environment, and van Gelder identifies situatedness as an important part of the dynamical approach. There is still debate over to what extent situatedness is involved in "higher level" cognition. Certainly human beings can cope in some situations where they are artificially constrained to operate in distinct input, processing and output stages; and plausibly, some interesting natural situations are approximately separated into such stages. Authors such as Kirsh [28] have argued for a range of natural human behaviours which are not "situation determined". On the other hand, humans in the real world do make extensive use of their situatedness even for very abstract tasks (e.g. using pen and paper to do mathematics).

We are not convinced that situatedness and the mathematics of dynamical systems theory logically have to go together. Brooks's mobots and his subsumption architecture were not described using tools or language from dynamical systems theory [2], and one could in principle model and analyse non-situated (i.e. input-processing-output) systems using differential equations and dynamical systems tools. Indeed, the section on "short-term memory" Continuous-Time Recurrent Neural Network (CTRNN) agents in [29] represents what might be called a dynamical approach to non-situated cognition. In this study a simulated robot controlled by an evolved continuous-time neural network is able to follow a moving stimulus even though the robot is deprived of sensory information when the robot starts moving. This agent is specifically constrained to solve the task without using a sensory-motor loop mediated by interaction with its environment. In other words, it is forced to operate in distinct input and output stages, like the paradigmatic computationalist agent.

VI. Abstract and real cognition

When trying to formulate theories about cognition, we should bear in mind that we do not know what we mean by "cognition." A frequently asserted but rarely defended assumption in dynamicist literature is that cognition is synonymous with cognition in the real world. That is to say, cognition could not in principle occur in a non-physical entity such as a simulated agent. When researchers who make this assumption use computational or abstract mathematical models they are intended on some level as models of physical, real-world phenomena in living systems. A typical example is due to van Gelder and Port:

> [C]ognitive processes are ultimately physical processes taking place in real biological hardware. [30] (p19)

This does not square with A-Life's tradition of considering "life as it could be" - that is, trying to identify and understand properties of complex living systems which are not restricted by the contingent physical history of our planet or even the physical laws of our universe. General, abstract properties

[2] This is not to say that they cannot use representations. The important point is that explicit symbols with defined meaning are not included in the model.

[3] Some researchers from the symbolic tradition (e.g. [25]) have suggested for completely different reasons that humans are able to perform uncomputable tasks, but this view is not widely held.

may or may not be of interest to field biologists depending on to what extent the phenomena they study are particular to our unique environment. Likewise, maximally abstract notions of cognition are not necessarily useful to cognitive scientists whose domain of study is biological wetware. But these are not the only sort of cognitive scientist around. The study of "cognition as it could be" cannot begin by assuming that cognition is something that occurs in physical systems in continuous time.

Consider researchers who are interested in intelligent behaviour in machines, in contexts which do not closely resemble the natural physical environment of humans or other animals. For example, machines to perform mathematical reasoning, automatically classify written material, or handle real-time resource allocation in computational networks, do not *obviously* have to have robot bodies. Although non-robot machines still have to be physically instantiated in the real world, their intelligent behaviour would occur in a radically different, virtual, world. If it is accepted that agents operating in such domains can in principle be cognitive, we are justified in wanting *abstract*, not *real-world-specific* cognitive science principles which we can apply to them.

Even further removed from the physical world, some researchers might be interested in totally abstract agents. They would exist only on paper, in the form of equations, thought experiments, or rational arguments. The abstract worlds inhabited by these agents - if they inhabited any at all - could be stripped of everything but the bare conceptual necessities for cognitive-like processes; alternatively, they might be explorations of exotic ideas like multiple time dimensions.

Abstraction often provides useful insights into the workings of the particular, but this is by no means its only purpose. A theory of cognition in the abstract would be in part about the question "What do we mean by 'cognition'?". We may observe here that the lack of a generally agreed answer to the question "What do we mean by 'life'?" has not hindered Artificial Life research; rather, it has stimulated it.

There is no obvious reason why the idea of situatedness (or related ideas like embodiment) cannot usefully be applied in the case of abstract agents. We might consider what properties an abstract situated agent would need to have depending on the abstract environment it was situated in. In fact, we will see later on that doing so helps to illuminate the role of time in the dynamical approach.

VII. THE DYNAMICAL EMPHASIS ON NUMERICAL MODELS

A common feature in much dynamicist research is that the system under consideration only has real-number variables and the dynamics are defined in terms of differential or difference equations. This is not a consequence of mathematical dynamical systems theory: Beer, following the same textbook mathematical definition we do, observes that

> The state space S may be numerical or symbolic, continuous or discrete or a hybrid of the two, and it may be finite- or infinite-dimensional depending on the number of variables required to fully describe the state of the system. [12] (p92)

However, there is a strong emphasis on continuous numerical models in the dynamicist literature. For instance, Van Gelder is quite explicit [31] that these models are the sort which should be used to understand cognition. Indeed, he uses the term "dynamical systems" exclusively for this limited subset of dynamical systems. That might be acceptable as a shorthand for philosophers who want to contrast dynamicist with computationalist approaches to cognition, but to those of us who might want to use the tools or language of mathematical dynamical systems theory to study *discrete-space* models of cognition, it is frustrating. We have suggested above that the specific term *quantitative*, which is also used by van Gelder, should be used for this subclass of dynamical system.

Mathematically speaking, continuous spaces do not have to be interpreted using numbers. In practice, the only current way for scientists to get a handle on continuous spaces is to use numerical coordinates, so in science continuous spaces are effectively synonymous with real number spaces. Typically such models in A-Life have a fairly small number of variables (although as pointed out in [12], continuous-field models such as the Smith & Thelen [32] model of the A-not-B task in human infants do not have a finite number of real variables).

Most numerical dynamicist models, such as CTRNNs, are mathematically continuous-time as well as continuous-space. In fact there is an important mathematical distinction between systems based on differential equations (which operate in both continuous time and continuous state) and systems based on difference equations (which are discrete in time but not necessarily in state). For instance, discrete-time systems with one variable can have chaotic dynamics, but at least three variables are required for chaotic behaviour in smooth systems.

A. Properties of systems based on differential equations

Well-behaved continuous systems do have various specific dynamical properties which are not shared by all discrete systems. For instance, they are guaranteed to be *reversible* in a specific technical sense. Moreover, as observed in [3], real-number systems come with a naturally associated *metric* (notion of distance between points in phase space) without which many of the tools of mathematical dynamical systems theory are inapplicable; discrete systems do not always have any readily computable natural metric. Another difference is that in continuous systems, phase-space trajectories do not cross one another, which is either meaningless or false for many discrete systems.

The degree to which discrete systems respect these properties is the degree to which they work as acceptable models of continuous systems. For instance, in digital computer simulations of continuous (in the abstract) CTRNN systems, if the Euler integration step size is set sufficiently small then (over the range of phase space considered) the discrete system will behave qualitatively like the abstract continuous system. For instance, it will respect reversibility and non-intersection of trajectories. However, if the integration step size is set too high, these dynamical properties will break down, producing discretisation artifacts. As long as the system is operating "like a continuous system" dynamically, its macroscopic behaviour will be indistinguishable from that of a continuous system (albeit a slightly different continuous system from the analytical one being modelled, due to numerical errors).

There are a number of distinct possible reasons why someone might argue for the study of numerical systems over discrete non-numerical systems. In order to further the debate on what sorts of models cognitive scientists should use, we believe they should be separated.

1) The researcher is interested in biological systems, and holds biological systems to be continuous, or at least approximately so on the appropriate scale. For many tasks, people and animals have to respond to stimuli in (effectively) continuous time and space using (effectively) continuous low-level system variables such as membrane potentials or chemical concentration. Similar considerations apply to mechanical robots.
2) The researcher believes that the dynamical properties of continuous systems are essential to cognition even in the abstract (whilst allowing that some discrete systems, such as computational models with floating-point numbers, can also have those properties). This could be an interesting hypothesis, but would need to be advanced and argued explicitly. We are not aware of anyone who has done so.
3) The researcher holds a weaker version of 2, namely that dynamical systems with continuous-like properties are the easiest ones to construct interesting cognitive models in (regardless, for instance, of whether they are operating in a discrete or continuous environment).
4) The researcher wants to apply dynamical systems theory tools which require a real-vector phase space to the analysis of his or her model. Note that none of cases 1-3 necessarily imply that the analytical tools of dynamical systems theory are relevant to understanding particular cognitive processes; rather, they suggest that certain basic properties are desirable (or necessary) for cognitive models.

B. A note on biology and continuous models

There is no current consensus on whether continuous models are the best ones for relevant biological systems. For instance, many interesting biological phenomena such as cell replication involve intrinsically discrete processes [33]. It is argued in [34] that various biological variables with upper and lower bounds essentially operate as digital systems due to the regulative dynamics.

On the other hand, there are already successful biological models which use continuous variables to model discrete phenomena. Examples are population dynamics (modelling the size of a population as a continuous variable), neural field models (modelling a large collection of individual neurons as a continuous field) and rate coding models of neurons (modelling a series of separate action potentials as a continuous firing rate).

C. A note on metric spaces

Occasionally dynamicist authors will make claims about the *a priori* suitability of continuous numerical models based on the metric properties of the state space. For instance: -

> [The time set in a continuous dynamical system] is a metric space, such that amounts of change in state are systematically related to amounts of change in time as measured by that metric. [35]

or: -

> The major drawback of hybrid systems is that... analogue processes and... symbolic processes cannot interact with each other intimately since the two pathways are defined in different metric spaces. [24] (p5)

We consider that these claims are slightly misleading. One of the features of complex dynamical systems, even ones characterised by sets of differential equations, is that their behaviour is often chaotic; over more than a very short temporal period, amounts of change in state are not systematically related to amounts of change in time in any way which is useful for an analyst. In other words, the state jumps about over the metric space much as a discrete system's state does (Beer makes this point in [4]).

VIII. Time and the Dynamical Approach

In this section we will illustrate the value of the "cognition as it could be" concept by examining the role of time in the dynamical approach to cognition. Contrary to what some previous commentators have claimed, we do not believe that continuous-time models are *a priori* a prerequisite for studying cognition in the abstract. Rather, they are appropriate when the agent under consideration is situated in a continuous-time world responding to changes in continuous time.

Timing has been put forward by as a central plank of the dynamical approach [30], [35]. This may be reasonable in terms of contrasting their focus with that of computationalist approaches, but Van Gelder also makes a distinction between what he calls "ersatz" time and genuine time: genuine time is continuous and comes in amounts, whereas ersatz time is a "mere order". On this view, computational processes operate in ersatz time as opposed to real or quasi-real time:

> However, none of the properties of the integers, over and above their constituting an ordered set, have any relevance to the Turing machine. [35]

In other words, labelling the time steps of a computational process using integers is held to be essentially a convenience - we might as well use the letters of the alphabet. Van Gelder is not wrong to insist that time in dynamical systems comes in amounts as well as being ordered; this is part of the mathematical definition of time in dynamical systems [7], where the evolution operator ψ satisfies the condition $\psi(t_1+t_2, x) = \psi(t_2, \psi(t_1, x))$, i.e. evolving for t_1 time units and then t_2 time units is the same as evolving for $(t_1 + t_2)$ time units.

However, as pointed out in [36], time in Turing Machines, or other computationally universal discrete-time discrete-space dynamical systems, does indeed come in quantities. The fact that each discrete step takes the same amount of time, and that four steps take twice as much time as two steps - in other words, the fact that one can apply integer arithmetic to time steps in computation - is essential to the field of algorithmic time complexity.

This observation is not a mere formal detail; both computationalist cognitive scientists and applied computer scientists typically assume that real-world implementations of computational processes take a consistent number of seconds per time step. This assumption underlies computationalist research which tries to infer cognitive "algorithms" from psychological response time data. For instance, the famous psychological study by Shepard and Metzler [37] on mental rotation is cited in Johnson-Laird's seminal computationalist text "The Computer and the Mind" [38]. In the Shepard & Metzler study, participants were shown two 2-D pictures of three-dimensional objects and were asked to determine whether or not they were the same object seen from different angles. Response times varied linearly with the actual angle of rotation, interpreted by the authors as suggesting that the cognitive process could be understood in terms of constant-speed mental rotation. Other examples are given in [36].

We believe that where the typical dynamicist approach differs from the computationalist approach in respect to time is actually in its situatedness. A major time-related drawback of computationalist accounts of human cognition is that humans exist in, and respond to, an environment which changes in continuous time. This supports dynamicist demands for human (and animal) cognition to be modelled in terms of dynamical systems with a continuous (or locally continuous-like) time dimension. It does not in general support an insistence that cognitive processes, in the abstract, need to occur in continuous-like time. The relevant issue is the relation between the time scale of the cognitive process and the time scale of the environment in which it is situated; processes situated in a discrete-time environment - or processes which operate in an input-processing-output temporal regime - might get by perfectly well on time "in ticks".

IX. Conclusion

We conclude that the terms "dynamical" and "computational" when applied to cognitive science are too loose in their meanings to be useful as technical terms, and should be understood as sociological labels for currently separate research communities. They are not solid enough to bear much philosophical weight, and if used as a theoretical basis their ambiguity risks merely reinforcing historical norms.

Our position is also that the phrase "dynamical system" is over-used. Its most general meaning in mathematics is simply any system whose state changes over time in a definable time-independent manner, including a Turing machine. Because mathematical dynamical systems theory is often (though not always) relevant to dynamicist research, casual use of its technical terminology is confusing. We strongly prefer the terms "dynamic" or "dynamical" (approach to cognition) over the term "dynamical systems". "Dynamical system" is best reserved for when technical concepts from dynamical systems theory are in use, such as attractors, phase portraits or Lyapunov exponents.

The theoretical emphasis put on continuous numerical models by some dynamicists is not sufficiently justified. In our opinion, even the claim that biological systems "have" continuous variables is contentious, and there is no *a priori* reason to use numerical differential equations when investigating "cognition as it could be" from a dynamicist perspective.

We recommend finer-grained technical terms to distinguish between different aspects of the dynamical and computational approaches. Some of these terms already exist: "situated", "quantitative", "sense-model-plan-act"; some may need to be coined anew, such as for what we have called "time-critical" behaviour or what we have termed "symbolic representation".

Acknowledgment

This paper would have been deeply impoverished without the contribution of Aaron Sloman and would never have been written without the benefit of discussions with Inman Harvey. We would like to thank our colleagues in the CCNR.

References

[1] M. Campbell, A. Hoane, and F. Hsu, "Deep blue," *Artificial Intelligence*, vol. 134, pp. 57–83, 2002.

[2] R. Brooks, "A robust layered control system for a mobile robot," *IEEE Journal of Robotics and Automation*, vol. 2, no. 1, pp. 14–23, 1986.

[3] T. van Gelder, "The dynamical hypothesis in cognitive science," *Behavioural and Brain Sciences*, vol. 21, no. 5, pp. 615–628, 1998.

[4] R. D. Beer, "Framing the debate between computational and dynamical approaches to cognitive science," *Behavioural and Brain Sciences*, vol. 21, no. 5, p. 630, 1998.

[5] T. van Gelder, "Disentangling dynamics, computation, and cognition," *Behavioural and Brain Sciences*, vol. 21, no. 5, pp. 654–661, 1998.

[6] Various, "Open peer commentary on van gelder: The dynamical hypothesis," *Behavioural and Brain Sciences*, vol. 21, no. 5, pp. 629–654, 1998.

[7] A. Katok and B. Hasselblatt, *Introduction to the Modern Theory of Dynamical Systems*. Cambridge, UK: Cambridge University Press, 1995.

[8] S. H. Strogatz, *Nonlinear Dynamics and Chaos: With Applications to Physics, Biology, Chemistry and Engineering*. Perseus Books Group, January 2001. [Online]. Available: http://www.amazon.co.uk/exec/obidos/ASIN/0738204536/citeulike-21

[9] R. D. Beer, "Autopoiesis and cognition in the game of life," *Artif. Life*, vol. 10, no. 3, pp. 309–326, 2004.

[10] D. Lind and B. Marcus, *An Introduction to Symbolic Dynamics and Coding*. Cambridge University Press, November 1995. [Online]. Available: http://www.amazon.co.uk/exec/obidos/ASIN/0521559006/citeulike-21

[11] W. R. Ashby, *Design for a Brain - The origin of adaptive behaviour*. John Wiley and Sons, 1960.

[12] R. D. Beer, "Dynamical approaches to cognitive science," *Trends in Cognitive Science*, vol. 4, no. 3, pp. 91–99, 2000.

[13] R. Port, *Dynamical Systems Hypothesis in Cognitive Science*. MacMillan, 2002, vol. 1.

[14] I. Harvey, E. Vaughan, and E. D. Paolo, "Time and motion studies: The dynamics of cognition, computation and humanoid walking." in *Proceedings of Fourth Intl. Symp. on Human and Artificial Intelligence Systems: From Control to Autonomy.(2004)*, University of Fukui, Fukui, Japan, December 2004.

[15] H. Jaeger, "Today's dynamical systems are too simple," *Behavioural and Brain Sciences*, vol. 21, no. 5, pp. 643–644, 1998.

[16] H. Maturana and F. Varela, *The Tree of Knowledge: The Biological Roots of Human Understanding*. Shambhala, 1987.

[17] R. A. Brooks, "Intelligence without reason," in *Proceedings of the 12th International Joint Conference on Artificial Intelligence (IJCAI-91)*, J. Myopoulos and R. Reiter, Eds. Sydney, Australia: Morgan Kaufmann publishers Inc.: San Mateo, CA, USA, 1991, pp. 569–595. [Online]. Available: http://citeseer.ist.psu.edu/41420.html

[18] H. Maturana and F. Varela, *Autopoiesis and Cognition: The Realization of the Living*. D. Reidel, 1980.

[19] J. Tani, "Model based learning for mobile robot navigation from the dynamical systems perspective," *IEEE Trans. Syst. Man and Cybern. B*, vol. 26, no. 3, pp. 421–436, 1996. [Online]. Available: citeseer.ist.psu.edu/77989.html

[20] Y. Kuniyoshi, Y. Yorozu, Y. Ohmura, K. Terada, T. Otani, A. Nagakubo, and T. Yamamoto, *From Humanoid Embodiment to Theory of Mind*, ser. Lecture Notes in Computer Science. Springer, 2004, vol. 3139.

[21] T. Inamura, I. Toshima, H. Tanie, and Y. Nakamura, "Embodied symbol emergence based on mimesis theory," *The International Journal of Robotics Research*, vol. 23, no. 4–5, pp. 363–377, 2004.

[22] R. M. French and E. Thomas, "The dynamical hypothesis: One battle behind," *Behavioural and Brain Sciences*, vol. 21, no. 5, pp. 640–641, 1998.

[23] J. J. Hopfield, "Neural networks and physical systems with emergent collective computational abilities," *Proceedings of the National Academy of Sciences*, vol. 79, no. 8, pp. 2554–2558, 1982.

[24] J. Tani, "An interpretation of the 'self' from the dynamical systems perspective: A constructivist approach," *Journal of Consciousness Studies*, vol. 5, no. 5/6, pp. 516–542, 1998.

[25] R. Penrose, *Shadows of the Mind: A Search for the Missing Science of Consciousness*. Oxford University Press, 1994.

[26] I. D. Horswill and R. A. Brooks, "Situated vision in a dynamic world: Chasing objects," in *Proceedings of the Seventh National Conference on Artificial Intelligence (AAAI-88)*, (St. Paul, MN), 1988, pp. 796–800.

[27] L. Suchman, *Plans and Situated Actions*. Cambridge Universtiy Press, 1987.

[28] D. Kirsh, "Today the earwig, tomorrow man," *Artificial Intelligence*, vol. 47, pp. 161–184, 1991.

[29] A. C. Slocum, D. C. Downey, and R. D. Beer, "Further experiments in the evolution of minimally cognitive behavior: From perceiving affordances to selective attention," in *From Animals to Animats 6: Proceedings of the 6th International Conference on the Simulation of Adaptive Behaviour*, J.-A. Meyer, A. Berthoz, D. Floreano, H. Roitblat, and S. Wilson, Eds., France, Sept 2000.

[30] T. van Gelder and R. F. Port, *It's about time: an overview of the dynamical approach to cognition*. Cambridge, MA, USA: Massachusetts Institute of Technology, 1995, pp. 1–43.

[31] T. van Gelder, "Revisiting the dynamical hypothesis." [Online]. Available: citeseer.ist.psu.edu/315733.html

[32] E. Thelen, E. Schoner, C. Scheier, and L. Smith, "The dynamics of embodiment: A field theory of infant perseverative reaching," *Behavioral and Brain Sciences*, no. 1, pp. 1–34, 2000. [Online]. Available: citeseer.ist.psu.edu/thelen00dynamics.html

[33] A. Sloman, Personal Communication.

[34] R. Thomas and R. D'Ari, *Biological Feedback*. CRC, 1990.

[35] T. van Gelder, "Bait and switch? real time, ersatz time and dynamical models." in *Dynamical Cognitive Science: Proceedings of the Fourth Australasian Cognitive Science Conference*, R. Heath, B. Hayes, A. Heathcote, and C. Hooker, Eds., Newcastle, NSW: University of Newcastle, 1999.

[36] N. Chater and U. Hahn, "What *is* the dynamical hypothesis?" *Behavioural and Brain Sciences*, vol. 21, no. 5, pp. 633–634, 1998.

[37] R. Shepard and J. Metzler, "Mental rotation of three dimensional objects," *Science*, vol. 171, no. 972, pp. 701–703, 1971.

[38] P. N. Johnson-Laird, *The computer and the mind*. Cambridge, MA, USA: Harvard University Press, 1988.