The Place of Time in Cognition

ABSTRACT

Dynamical systems theorists (dynamicists) allege that symbolic models of cognition are essentially incomplete because they fail to capture the temporal properties of mental processing. I present two possible interpretations of the dynamicists' argument from time and show that neither one is successful. The disagreement between dynamicists and symbolic theorists rests not on temporal considerations *per se*, but on differences over the multiple realizability of cognitive states and the proper explanatory goals of psychology. The negative arguments of dynamicists against symbolic models fail, and it is doubtful whether pursuing dynamicists' explanatory goals will lead to a robust psychological theory.

1 Introduction

The symbolic theory of mind in cognitive science has been under pressure from a number of sources in recent years. Connectionists have long offered models of cognition that do not employ localized symbol tokens over which computations might operate. The problem of describing ordinary, situationally relevant deductive inference without succumbing to combinatorial explosion (roughly the Frame Problem) remains an open challenge for all extant symbolic models. Non-deductive forms of inference have proven still more recalcitrant to such descriptions (Fodor [2000]). And neuroanatomy and neurophysiology have suggested to some that mental symbols themselves, distinguished by their syntactic properties, may simply be found nowhere in the densely tangled networks of the brain (Garson [1997]).

All of these challenges are by now familiar. This paper is about a further challenge to the symbolic conception of the mind posed by dynamical systems theory. Advocates of dynamical systems theory, or dynamicists, challenge the empirical and conceptual adequacy of the symbolic conception of mind. They argue that symbolic models are inadequate because such models, based on discrete computational transitions between representational states, ignore or otherwise omit the obvious fact that '[c]ognitive processes always unfold in real time' (van Gelder and Port [1995], p. 18). While symbolic models need explicit supplementation in the form of timing assumptions to capture the temporal unfolding of cognitive processes, dynamical models incorporate time from the start. No *ad hoc* or otherwise suspect assumptions about the chronometry of component processes are needed. This argument from time is a novel challenge from dynamical systems theorists to their symbolic competitors.[1]

The structure of the paper is as follows. In section 2, I review the conceptual structure of the symbolic theory. In section 3, I present a qualitative summary of the mathematical tools necessary to understand what is at stake in the dynamics debate, focusing on the role of time and continuity of state in dynamical modeling. In section 4, I present two interpretations of dynamicists' argument from time, and argue that neither is successful. In section 5, I suggest that the argument from time is motivated by background disagreements over multiple realizability and the explanatory goals set by the two competing theories. I sketch these goals and argue that the symbolists' goal seems more plausibly achievable for cognitive science.

2 Elements of the Symbolic Theory

The symbolic theory consists of two interlocked claims, one about the structure of thoughts and the other about thought processes.[2] Thoughts are mental representations syntactically composed from discrete symbols. Basic symbols are assigned to syntactic categories that determine their combinatorial potential, and complex symbols are composed from appropriately combined basic symbols. Paralleling these syntactic combinatorics, each basic symbol has a semantic assignment, and the semantic assignments of complexes are composed from those of their syntactic components according to the rules of semantic combination; e.g., the meaning of SAD CLOWN is fixed by the meanings of SAD and CLOWN plus the phrase's syntactic profile, and it refers to the intersection of the two terms' extensions.[3] Having a thought is having such a complex representation tokened in some functionally defined subpart of one's cognitive system.

Thinking is computation, a series of causal transitions that are responsive exclusively to the formal properties of symbols.[4] All cognitive-level processes are composed from the basic operations and resources provided by the functional architecture of the organism under study. A functional architecture is just the set of such basic operations, resources, and constraints within which cognition takes place, and which are themselves immune to being altered by symbolic processes (Pylyshyn [1984]). The classic abstract model of the symbolic tradition is the Turing Machine, which engages in discrete state transitions determined by no more than the shapes of the symbols that are printed on its tape and the rules embodied in its program. Hypotheses about functional architecture usually suppose a richer body of primitive operations, a greater proliferation of memory stores, different methods of addressing stored data, and different ways of dictating the flow of control in executing stored programs than the Turing Machine model allows. Nevertheless, despite these differences, mental functioning, for symbolists, consists in the rule-governed transformation of structured mental representations in various places within an organism's functional architecture.[5]

Within these elements of the general symbolic approach, particular models of mental phenomena can be developed. When we apply a symbolic model to a set of experimentally derived data, we want to assess how good of a fit the model is. This assessment depends on the particular data set in question. For instance, a model of categorization might aim to predict, *inter alia*, the probability that certain presented items will be placed in one category versus another. By adjusting the model's parameters, the fit can be made better or worse.

Models do not predict every facet of human behavior, however. In particular, they omit one parameter that dynamicists hold is crucial: the time course of the cognitive process in question. So a categorization model may be silent about how long a subject takes to arrive at a category judgment, or a model of memory may not say precisely how long it takes a subject to retrieve an item that has been partially transferred to long-term memory. Not only is this data sometimes uncollected by experimenters, the models don't contain parameters for temporal information.

A notable exception to the general lack of concern with temporal phenomena is the use of reaction times (RTs) as a dependent measure.[6] RTs have been widely used and studied, but their precise role in relation to psychological theory remains controversial. Pylyshyn ([1984]) suggests three ways in which RTs may be understood. First, they might be taken to be just one more aspect of behavior that the model needs to entail or generate values for. A model, on Pylyshyn's understanding, generates the series of representations that the system being studied uses to carry out a task. Cognitive models would then need to have as output not only the ordered sequence of representations and processes employed in the task being studied, but also a representation of the temporal profile of those representation-process pairs. Second, RTs might be taken to be essentially direct measures of the duration of component mental processes. In conjunction with hypotheses about which processes are reused in various tasks, RTs might reveal the precise chronometry of the functional architecture of the organism under study.

Pylyshyn rejects both of these interpretations. He holds instead that RTs are tools with which we can decide between otherwise hard to distinguish hypotheses about the computations an organism is carrying out. If we have two putative algorithms, one more computationally complex than the other, each of which maps the same inputs onto the same outputs, RTs might be able to distinguish between them on the assumption that the relationship between computational complexity and speed of execution is roughly linear. But the RTs themselves are neither entailed by the cognitive models, nor taken as measures of the chronometric properties of mental events themselves. They are properly understood as tools for distinguishing competing cognitive-level explanations.

Like many other symbolic theorists, Pylyshyn thinks that the character of the physical medium that realizes symbolic states (processes, events, etc.) is irrelevant to their identity as symbolic states. This follows from the multiple realizability of cognitive states. If it is possible that the very same cognitive states that we or other organisms possess could be had by entities with very different physical constitutions, then it should be possible that some ways of realizing cognitive states will result in their having different chronometric properties, as well as different chemical and biological properties.[7] Computation happens faster in silicon than in neurons, but symbolic theorists hold that the very same cognitive processes could be realized by both electronic circuits and neural wetware. Believers in multiple realizability may appeal to intuitions to motivate their thesis, but the final verdict depends on whether the same theoretical explanations can be applied to organisms whose cognitive substrates differ. The more robust these explanations are over a range of substrates, the more likely multiple realizability becomes.[8]

Given multiple realizability, Pylyshyn claims that although mental event tokens, like all other token events, have particular locations, durations, etc., mental event types do not, any more than they have, say, neurobiological properties. Duration is like other physical properties in belonging to the substrate, or realization base, of mental event types in various organisms (and possibly at different times within the same organism), but properties of the realization base do not belong to what is realized. Even if, in all metaphysically or nomically possible realizations, mentally dividing 256 by 16 had the same chronometric profile, having that profile would not be a property of the mental event type dividing 256 by 16. So, although temporal information is sometimes collected and used by experimentalists, its role is principally to constrain cognitive models that are themselves insensitive to the timing of the underlying systems that implement them.

3 Elements of Dynamical Systems Theory

The dynamical theory also comprises claims about structures and processes.[9] Dynamical systems are just state-dependent systems consisting of numerical quantities that vary over time according to a rule, such as a nonlinear differential equation. A system's state is the set of values had at a time by the properties of the objects that form the system. The current total state of a system, in conjunction with a rule of evolution, determines its total state at any arbitrary future moment. Systems can be modeled geometrically by depicting their state space, which is an *n*-dimensional mathematical space whose dimensions correspond to the state variables of the system. The total state of the system at a time can be represented as a single point in this space, and the evolution of the state variables of the system can be represented as the movement of the system point around in the space.[10] Mental states are identified with state-space points, and mental processing is the motion of a point through the geometric structure of the space. The structure of state space determines the system's potential range of states much as a symbolic system's functional architecture does (with one important difference to be noted shortly).

Dynamicists model cognition using mainly low-dimensional state spaces. This is an important empirical constraint motivated by the complexity of neurobiological systems. For all practical purposes, it is impossible to model every variable involved in the generation of even the simplest monosynaptic spinal reflex, even when the general principles are understood. And even if we knew all of the relevant state variables to examine, writing and solving the resultant equations is certainly beyond us. When one scales up to the complexity involved in activities such as playing squash, writing poetry, planning a vacation in the Ozarks, or doing mathematical proofs, it becomes clear why some reduced-dimensional model is necessary.[11]

Attempting to model a system, then, requires searching for the right collective variables with which to characterize its states. Collective variables, sometimes also called 'order parameters', are numerically measured quantities that supervene on the behavior of a system's lower-level constituents, which are often homogeneous or near-homogeneous. Such variables need not correspond neatly to any particular part of the system being modeled. Kelso points out that a salient feature of a collective variable is that it 'is created by the coordination between the parts, but in turn influences the behavior of the parts' (Kelso [1995], p. 16). This relation of interlevel reciprocal influence[12] purportedly offers a more natural integration of cognition with the implementing neural structures than does the symbolic model. Finding the right few collective variables to characterize a massively interactive, many-component system is the challenge of much empirical work in dynamics.

A system's behavior may also be guided by control parameters, which 'lead the system through different patterns, but that (unlike order parameters) are not typically dependent on the patterns themselves' (Kelso [1995], p. 16). Control parameter settings 'fix the dynamics of the system' (van Gelder [1995], p. 356) because they determine the way in which collective variables change their values over time while themselves remaining unaffected by those changes in state. While control parameters are similar to the symbolic theorist's functional architecture in that they determine the possibilities for the system's temporal evolution, they differ in that they may change more frequently and more reliably, in ways that are generally adaptive for the system. Functional architecture can change as an effect of implementation-level changes such as neurological damage, but it does not normally do so in a reliable way, and such changes are never part of the system's regular operation.

One set of control parameter values might place an attractor at a certain point of state space, disposing all the trajectories that pass within its basin of attraction to converge on it. Changing the parameters might induce a repeller to form within that region, which would lead systems that were formerly converging to diverge suddenly. This sort of rapid, global change in the dynamics of systems can sometimes be induced by comparatively simple means. Control parameters may be as elementary as the temperature on one side of a room, the mass of an infant's leg, or the rate at which one's finger is making a wagging motion. Parameters like these are elementary compared to a syntactically structured representation whose content characterizes the to-be-produced behavior. The ability to produce global shifts in complex patterns of behavior by manipulating such primitive properties underlies much of the dynamicists' suspicion about positing complex internal structures such as symbols.

Kelso's ([1995]) model of synchronous finger motion illustrates these properties. When subjects are asked to move their index fingers back and forth in time with a metronome, they usually fall into one of two stable patterns involving the fingers being in-phase or anti-phase with each other. At low rates either pattern is available. Above a certain rate, however, a sudden shift occurs and one pattern becomes strongly preferred. There is a sudden 'switch' in the shape of the system's behavior, controlled by two numerical control parameters corresponding to the movement frequency.

The model illustrates a dynamical versus a 'motor programming' approach to explaining behavior. '[O]ne of the main motivations behind these experiments was to counter the … notion of motor programs, which tries to explain switching (an abrupt shift in spatiotemporal order) by a device or mechanism that contains "switches"' (Kelso [1995], p. 57). Nothing in the finger-movement model corresponds to an internal state of the organism that represents the frequency of movement, or the state of the fingers' coordination with respect to one another. As Keijzer puts it, 'the control parameter does not "control" the system in any conventional sense … it does not prescribe what the system should do' (Keijzer [1998], p. 279).[13] A motor program, containing symbols that directly initiate the observed switching behaviors, therefore seems superfluous. Dynamicists hope to extend this explanatory model, on which complex interactions and patterns of behavior are controlled by simple, non-representational (and especially non-symbolic) parameters, to all of cognition.

4 The Argument from Time

Dynamicists argue that that symbolic models of mind are essentially incomplete, in that they have no way to capture interactions such as these, that occur in real time. The argument is simple. Computational state transitions are discrete. Dynamical systems theory, on the other hand, uses the apparatus of differential equations, which were designed explicitly for dealing with continuous quantitative changes in time. And since actual '[c]ognitive processes always unfold in real time' (van Gelder and Port [1995], p. 18), the dynamical approach must be the superior tool for understanding their behavior. There will always be an infinite number of systemic states that are uncaptured in a discrete computational model. So, for cognitive models, chronometry is destiny.

There are two major reasons that dynamicists give such emphasis to temporal considerations. The first is that models operating at the grain of real time will be able to predict the behavior of cognitive systems in far more detail than discrete-time models can. This follows straightforwardly from the fact that continuous models describe more states than any discrete models do. Given the general scientific goal of increasing descriptive and predictive adequacy, we should prefer more detailed models where we can get them. I will discuss some consequences of this claim in section 5. The second reason, which I will not discuss in detail, is that models cast in terms of nonlinear differential equations are more readily integrated with models from other sciences, such as biology and physics. They claim that we should prefer a psychological theory whose mathematical formalism is easily mapped onto the successful mathematical formalism of other sciences, since psychological systems are instantiated in and interact with systems described by these other sciences, and this easy mapping will facilitate the integration and unification of psychological theory with the rest of science.

The controversial premise in the dynamicists' argument is the claim that cognitive processes always unfold in real time. Van Gelder and Port offer two interpretations: first, 'real time is a continuous quantity best measured by real numbers, and for every point in time there is a state of the cognitive system' (p. 19); and second, that the timing of events in a cognitive system always matters to the system's operations. I will take these in turn.

4.1 First Interpretation of the Argument from Time

It's not up for debate that time is continuous. And there is a sense in which symbolic theorists might hold that for every point in time there is a (cognitive) state of a cognitive system, since cognitive states might persist across many instants, with no gaps between them. But dynamicists are committed to the stronger claim that at practically every point in time the system occupies a type-distinct cognitive state.[14] Consider two time-slices of me taken some fraction of a second apart. Numerous physical changes will have certainly occurred: neurotransmitters have diffused a minute amount across synaptic gaps, muscle fibers have contracted an imperceptible distance, and action potentials have jumped partway down axons. At lower levels of physical description there are many more changes. The important question for psychology is: do any of these count as cognitive changes? Van Gelder and Port are committed to thinking that all do, and that processes taking place within an arbitrarily small temporal interval will also count as cognitive changes. This claim, they say, 'is really just an obvious and elementary consequence of the fact that cognitive processes are ultimately physical processes taking place in real biological hardware' (p. 19).

I take it that the 'ultimate' identification of cognitive systems with physical/biological systems is equivalent to some form of materialism. Consider a weak materialism such as the token identity thesis: every mental state of an organism is identical with some physical state of that organism. It doesn't follow just from this weak materialism that cognitive states are equinumerous with physical states. One way that the equinumerosity of cognitive and physical states might be arranged is by the converse of the materialist claim: for every physical state of an organism there is some cognitive state with which it is identical. This converse claim, though, doesn't follow just from the materialist premise.

Without asserting the converse of weak materialism, there are other ways to establish the equinumerosity claim. A cognitive dynamical system is a system whose collective variables supervene on some underlying physical variables. As these underlying physical variables change, the cognitive collective variables also change, driving the system's progress through state space. But it need not be the case that every physical change results in a change in the cognitive variables, even when this physical change is in a part of the system on which those variables supervene. Many values of physical state variables may correspond to single values of cognitive state variables. What matters is that there still be a continuum of cognitive states. This way of implementing the equinumerosity claim has the advantage of not being implausible on its face, but it is not a claim that van Gelder and Port are in a position to advance, since it presupposes what is to be argued, namely that the cognitive state variables occupy a continuum of positions.

Finally, there is another sense in which it is true that, given token identity, every physical change is also a cognitive change, since when there is a new token physical state there will *ipso facto* be a new token cognitive state. But this is not clearly a form of cognitive change that is interesting to psychology, since a physical change may result only in a token of a cognitive state type being replaced by another token of the same type.[15] So there would be no change in cognitive state type across these physical changes. Arguably, changes in state type are all that psychology (or any other science) is directly concerned with, since the dynamical laws of sciences make reference to transitions among state types.

It may be that van Gelder and Port think that allowing cognitive or mental states to be characterized as discrete at any level whatsoever threatens the superiority of dynamical systems modeling. If mental states are discrete, then it follows that they can be modeled discretely, and therefore they could be modeled symbolically.[16] But while this would explain the assertion of the extreme thesis, it does not justify it. Van Gelder and Port can't simply assume that there are as many cognitive states as there are physical ones, since that's what they're trying to establish with the argument from time. If there were a continuum of cognitive states, then, trivially, symbolic theories would be inadequate insofar as they failed to make distinctions among them that are predictively and explanatorily significant. This, though, is just what is at issue.

Van Gelder and Port's other considerations do not fare much better. They are willing to concede that some cognitive processes, such as mentally carrying out long division, can be thought of as discrete state transitions. But they argue that there are 'innumerable kinds of tasks that cognitive systems face which appear to demand a continuum of states in any system that can carry them out' (p. 22). The example they give is sensorimotor coordination such as what takes place while playing tennis, when the organism must react to events in a continuous surrounding medium by occupying 'states that are equally rich [as those in the environment] and subtly distinct' (p. 22). I can see no reason for this strong assumption, however. Sensorimotor coordination requires that an organism be good enough at distinguishing and manipulating objects in its environment to achieve its ends, survive, and so on. It need not distinguish them in all their resplendent and variegated detail, or at least we cannot assume that it must. It doesn't seem implausible that to make so many distinctions would swamp the system's processing capacity. And even if organisms make an extremely high number of perceptual discriminations, it's not obvious that all of those discriminations function in the cognitive processing that leads to action. A great deal of cognition seems to be directed towards selecting salient commonalities among perceptual stimuli and grouping them into equivalence classes that are more coarse-grained than those of perception. It is at least as important in cognition to throw away information as it is to collect it (Clark and Toribio [1994]).

It's worth noting, in any case, that that our own nervous systems do not seem to respond to such a continuum of details. While the impinging energy on sensory receptors might be continuous-valued, and the receptor potentials are also analog, the afferent signals are coded by digital, all-or-nothing action potentials. Probably, part of the explanation for the use of such digital signals is that neural channels tend to be noisy and lossy, hence a message that is too fine-grained will tend to get lost in transit. Analog signals that are sensitive to fine temporal information cannot travel the length of the typical axon. Large, discrete signals are better vehicles of information than subtly distinct signals in such conditions (Anderson [1995], pp. 32-3).[17]

4.2 Second Interpretation of the Argument from Time

So much for the first interpretation. The second interpretation of the premise that '[c]ognitive processes always unfold in real time' is considerably less controversial. It says only that how long a system spends in a particular state, how rapidly it makes a state transition, and how quickly it can interact with other states all matter to the system's overall performance. Parsing a heard sentence, for example, takes a certain amount of time for a system. There is an effective upper limit on how long it can take, given that the system will constantly be barraged with new strings to parse, as well as other tasks that must be simultaneously carried out. Some strings are parsed faster than others, and this fact is important for how the system's overall behavior is guided by its verbal inputs; and so on. Timing considerations are clearly important to organisms that need to react in order to take advantage of the opportunities their environment offers and avoid its hazards.

Dynamicists rightly point out that insofar as their models are cast in terms of differential equations they will automatically have a definite answer to the question of how long a certain cognitive process takes, how long a system spends in the region of a certain state, and so on. Furthermore, on the standard Turing machine model of computation, 'time' means no more than the order in which a system passes through its states. It makes no sense to ask of a Turing Machine how long it spent in machine state 209, or what state it was at between times 54 and 55. Turing Machine time is just order. A Turing Machine's program does not entail anything about the timing of its operation when it is implemented in biological or silicon hardware, nor do more sophisticated symbolic models constructed by psychologists.

It is obviously true that the timing of cognitive processes matters to the practical success of the organism. This is one sense in which time matters. But it's unclear whether it should matter to the theorist interested in modeling and understanding the cognition of the organism. Not everything that matters to the organism and its success needs to matter to the theorist's principled individuation of the organism's states for purposes of psychological explanation.

For the moment, consider whether a supplemented symbolic model might answer dynamicists' criticism. A supplemented symbolic model is a symbolic model plus a set of auxiliary assumptions about the time it takes the brain (or whatever realizes the cognitive system) to carry out the basic operations posited by the model.[18] With these auxiliary assumptions we might extract predictions about the (more or less) precise timing of processes from a supplemented computational model. The introduction of parallel processing into symbolic models complicates things considerably, but the basic assumption is still that there is a fixed speed of execution for the basic operations in whatever processors may exist.

The objection to supplemented models is that the assignments of times to basic operations is radically unconstrained. In effect, it is a free parameter that can be set in any way that fits the data. In assessing the charge we have to distinguish between whether an assignment of durations to basic mental processes is *ad hoc* in practice versus whether it is *ad hoc* in principle. In many cases, decisions about the time basic operations take in a model are in a sense *ad hoc*. Even here, though, we should temper the charge. Such precise assignments are usually derived by fitting the model, augmented with a set of time parameters, to the observed temporal data. The sense in which this is *ad hoc* is only the sense in which any parameter-fitting in a complex model is *ad hoc*. That is, it depends on the plausibility of the assignment of parameter values given everything else that is known about the process under study. If another task suggests a different assignment of times to the same primitive operations, then the original assignment may be disconfirmed. This is a rather weak sense of '*ad hoc*'.

These considerations suggest that timing assignments do not need to be *ad hoc* in principle, either. Convergence with other psychological tasks that tap the same hypothesized processes would be one source of constraint and convergent evidence. Neurophysiological considerations are another obvious possible source of constraint. We might begin with the assumption that basic cognitive operations supervene on neurophysiological processes that have a stable temporal course. Comparatively little is yet known about the computational properties of actual neurons, as distinguished from their oversimplified counterparts (e.g., connectionist nodes). Still, the possibility of reducing the *ad hoc* parameters to zero is real. The mere fact that, at present, timing assumptions are often unmotivated does not by itself argue for the abandonment of the supplemented symbolic framework.

In fact, this point about timing has not gone unnoticed by symbolic theorists. Allan Newell, in his discussion of unified theories of cognition, notes that cognitive models operate under what he calls the ‘real-time constraint’ ([Newell, 1990], pp. 129-31). The real-time constraint states, in effect, that given how long it takes humans to carry out simple behaviors and given how long it takes neurons to carry out elementary processing, there can only be a certain number of operations and levels that implement those behaviors. Newell’s constraint can be understood as a proposal about how to use information about the implementation of cognition to determine the structure of cognition itself. It is thus of a piece with Pylyshyn’s proposal about the proper use of RTs. Both point out how many sources of data (behavioral and neurophysiological) can converge to winnow down the space of possible cognitive models. Neither is committed to an essential role for temporal considerations in cognitive models themselves, however. Timing remains strictly supplementary.

Furthermore, there are paradigm examples of successful dynamical systems modeling that contain open parameters for the timing of basic operations that are exactly the same as those left open in symbolic models. Busemeyer and Townsend’s ([1993]; Townsend and Busemeyer [1995]) Decision Field Theory is one such model, touted by van Gelder ([1995]) as the flagship example of a dynamical model that deals directly with processes of central cognition. The model employs seven parameters to account for a variety of important empirical phenomena in research on human decision making, including violations of the transitivity of preference, speed-accuracy tradeoffs, and the inverse relationship between choice probability and decision time. Most important in the present context is the parameter h, which they introduce in the last stage of developing the model. The previous models they consider ‘are all discrete-time dynamic processes, and, consequently, they cannot be used to make quantitative predictions for decision time' ([1993], p. 444). In this they are open to the criticisms that face symbolic models. The new parameter *h* 'represents the amount of time used to process each sample valence difference. In other words, h represents that amount of time that it takes to retrieve and process one pair of anticipated consequences before shifting attention to another pair of consequences' ([1993], p. 444). The time unit parameter thus stands for how long the basic operation of comparing potential choices takes the cognitive system. But just as in symbolic models, the parameter is set simply by fitting the empirical data. No attention is paid to underlying neurophysiological factors or other 'objective' sources of constraint.

The fact that at least one highly touted dynamical model employs the same method for making precise temporal predictions as do symbolic models leads to two possible conclusions. It might be that both symbolic models and Decision Field Theory are badly off, or it might be that neither are as badly off as dynamicists have sometimes claimed. The latter seems to be a more appealing alternative. Precise timing predictions for all models are presently derived by data-fitting. Data-fitting temporal parameters is no more *ad hoc* than any other kind of parameter fitting, however. So neither dynamical models nor symbolic models should be reviled for employing such techniques.

I conclude that on neither interpretation of their most controversial, and crucial, premise do van Gelder and Port make the case that the symbolic theory is weakened by the argument about time.

5 Limits of Dynamical Systems Theory

The argument from time is unsuccessful, but the argument itself is best understood when set against the broad background of the differences between dynamical and symbolic theorists over what the proper explanatory structure of cognitive science is.

Recall that according to Pylyshyn and the orthodox symbolic theorists, it isn't the task of a cognitive model to entail anything about mental chronometry. This is because cognitive models aim to tell us only what can be systematically known about mental states, and (type) mental states don't have temporal properties—their various realizers do, but they inherit none of them. Given this, the dynamicist's charge that symbolic models leave out time seems strikingly misdirected. Symbolists just think that timing is a matter of implementation.

There are ways to rationalize the charge, though. One way would be for dynamicists to deny multiple realizability, and therefore to squash together the cognitive level and the level of implementation. Multiple realizability entails that psychological states are widely sharable, and thus guarantee that a robust psychological theory that generalizes over individuals that differ in many non-psychological respects is possible. If the very identity of a mental process depends on the extremely fine details of its temporal unfolding, then it is highly unlikely that these states will be shared (even if it is distantly possible that they are sharable).

Dynamicists sometimes espouse such cross-level modeling. For example, Thelen and Smith ([1995]) suggest that:

> If we want to explain the dynamics of cognitive structures—how they emerge and change and break apart—we cannot write theories at only the macrostructural level. Nor will we succeed only by looking at the microlevel. […] Explanation requires that we keep both the view from above and the view from below (p. 41).

If cognitive theories require that we attend to both the level of realization and the higher levels that supervene on it, then there will *ipso facto* be different cognitive theories for organisms whose cognitively pertinent material structures differ. Unless there are ways of type-identifying neurobiological structures that make, say, humans and cats come out as neurobiologically type-identical, then there will be no cognitive theory that is common to both. Plausibly, the same theory will not even cover human adults and human infants, given the radical changes that occur in neural organization during normal development (for instance, myelinization of frontal cortex continues up to adolescence, resulting in increases in processing speed for some tasks).[19]

An alternative approach, although one closely related to outright denying multiple realizablity, would be to accept that cognitive states are multiply realized, but maintain that the proper study of cognition requires attending to both the higher and lower levels in particular kinds of organism. So, for instance, both humans and cats might be able to entertain the same cognitive state, realized differently in each case, yet the proper study of that cognitive state requires pairing study of the higher level with various kind-specific studies of its realization bases. The quote from Thelen and Smith is ambiguous between these two possibilities. This position, though, still results in a proliferation of theories tailored to however many different realizations of the same cognitive states there are. In that, it is vulnerable to the criticisms leveled against the outright denial of multiple realizability. While a complete understanding of mind and behavior requires explanations pitched at multiple levels, the working hypothesis of symbolic cognitive science is that there is a level of study that is robust and implementation-independent.

The commitment to real time modeling, and thus to a continuum of states for cognitive systems, poses related difficulties. Dynamical systems models individuate cognitive states in an extremely fine-grained way. Dynamical models of motor systems make highly specific predictions about the bodily movements the organisms will make at any given time. Kelso's ([1995]) finger motion experiments might be taken as a paradigm here, as might Thelen and Smith's ([1995]) research on infant reaching. Is such highly fine-grained description always a virtue? Van Gelder and Port think so: 'It is only in the fine details of an individual subject's movements and their change over time that the real shape of the dynamics of development is revealed' (p. 17). And Thelen and Smith ([1995], p. 342) emphasize that detailed studies of individual subjects are the ideal tool for revealing the forces that drive development.

Pursuing this line, however, runs the risk of committing us to thinking of psychology as the science of finding laws tailored to fit the total behavior of individual organisms over their whole lifespan. Such laws might take the form, for example, of a set of dynamical equations precisely tailored to cover one of Thelen and Smith's infant subjects' motor behavior. These equations might provide some explanatory purchase insofar as they are counterfactual-supporting. Nevertheless, they seem little better than mathematical biography. Here dynamicists emphasize the theoretical virtue of predictive fecundity to the exclusion of the virtue of predictive robustness. We want our psychological generalizations not just to predict individual behavior in detail, where possible, but also to predict what many individuals will do. And we will trade fine-grained prediction off against broad generalization, other things being equal. This only makes sense, given the variety of individuals we will encounter over our lives. Close studies of individuals may be valuable, but to achieve generality we may have to sacrifice fineness of detail.

An analogous point has been made by Fodor and McLaughlin ([1998]) in their critique of Smolensky's connectionist program: '… the moral is that whether you have a level of causal explanation is a question, not just of how much precision you are able to achieve, but also of what generalizations you are able to express. The price you pay for doing psychology at the level of units is that you lose causal generalizations that symbol-level theories are able to state' (p. 108). For 'units', read 'particular bodily movements', and the same critique would seem to apply to dynamicists. Dynamicists need to make it plausible that as we broaden the scope of our laws to include more individuals in their generalizations we do not find ourselves forced back into the comparatively abstract inner structures that symbolists posit. This requires presenting more well-articulated models of the symbolic theory's central phenomena than are currently on offer.

It is tempting to suggest that an ecumenical approach to cognitive science might be able to adopt tools of both dynamical systems theory and symbolic computation. Strongly favoring one over the other for all tasks might unnecessarily limit our potentinal explanatory resources. While this cautionary note is laudable, our wholehearted endorsement should wait on seeing the shape of potential models that can integrate symbolic and dynamical components. Clark ([1998]), for example, suggests that we can honor the dynamicist's urging that time is important to cognitive modeling by allowing temporal properties of brain states to partially individuate representational vehicles. He points out that oscillatory neural processes might themselves be 'high-level syntactic items whose functional role is to carry, communicate, and transform specific bodies of information' (p. 371). The suggestion is one that symbolic theorists should seriously consider. But if this is the shape of the ecumenical model, dynamicists will be unlikely to sign on, since this is tantamount to confining the importance of time to the level of implementation, rather than allowing it into the cognitive models themselves.

The future of mixed models thus remains unclear. What is clear, though, is that the negative arguments of dynamical systems theory underestimate the considerable conceptual resources of the symbolic theory and rest on dubious assumptions about cognition. The continuity assumption, in particular, is troubling. It might turn out, if the best psychological models were dynamical, that we have reason to believe that mental states occupy a continuum, but we cannot simply assume it to be so.

Acknowledgements

I would like to thank Andy Clark, Philip Robbins, Chase Wrenn, an anonymous reviewer, and audience members at the 2001 meeting of the Society for Philosophy and Psychology for their helpful comments on earlier versions of this paper.

Daniel A. Weiskopf

Department of Philosophy

Washington University in St. Louis

Campus Box 1073

One Brookings Drive

St. Louis, MO 63130–4899

Email: daweisko@artsci.wustl.edu

References

Anderson, J. A. [1995]: *An Introduction to Neural Networks*, Cambridge, MA: MIT Press.

Bechtel, W. [1998]: 'Representations and Cognitive Explanations: Assessing the Dynamicist's Challenge in Cognitive Science', *Cognitive Science*, **22**, pp. 295-318.

Block, N. [1980]: 'Troubles with Functionalism', in N. Block (*ed.*), *Readings in Philosophy of Psychology* (Vol. 1), Cambridge, MA: Harvard University Press, pp. 268-305.

Busemeyer, J. R., and Townsend, J. T. [1993]: ‘Decision Field Theory: A Dynamic-Cognitive Approach to Decision Making in an Uncertain Environment’, *Psychological Review*, **100**, pp. 432-59.

Churchland, P. M. [1989]: ‘On the Nature of Theories: A Neurocomputational Perspective’, in P. M. Churchland, *A Neurocomputational Perspective: The Nature of Mind and the Structure of Science*, Cambridge, MA: MIT Press, pp. 153-96.

Clark, A. [1997]: ‘The Dynamical Challenge’, *Cognitive Science*, **21**, pp. 461-81.

Clark, A. [1998]: ‘Time and Mind’, *Journal of Philosophy*, **95**, pp. 354-76.

Clark, A., and Toribio, J. [1994]: ‘Doing Without Representing?’, *Synthese*, **101**, pp. 401-31.

Eliasmith, C. [1996]: ‘The Third Contender: A Critical Examination of the Dynamicist Theory of Cognition’, *Philosophical Psychology*, **9**, pp. 441-63.

Fodor, J. [1975]: *The Language of Thought*, Cambridge, MA: Harvard University Press.

Fodor, J. [2000]: *The Mind Doesn’t Work That Way*, Cambridge, MA: MIT Press.

Fodor, J., and McLaughlin, B. [1998]: ‘Connectionism and the Problem of Systematicity’, in J. Fodor, *In Critical Condition*, Cambridge, MA: MIT Press, pp. 91-111.

Garson, J. W. [1996]: ‘Cognition Poised at the Edge of Chaos: A Complex Alternative to a Symbolic Mind’, *Philosophical Psychology*, **9**, pp. 301-22.

Garson, J. W. [1997]: ‘Syntax in a Dynamic Brain’, *Synthese*, **110**, pp. 343-55.

Keijzer, F. [1998]: ‘Doing Without Representations Which Specify What to Do’, *Philosophical Psychology*, **11**, pp. 269-302.

Kelso, J. A. S. [1995]: *Dynamic Patterns*, Cambridge, MA: MIT Press.

Luce, R. D. [1984]: *Reaction Times*, Oxford, UK: Oxford University Press.

Newell, A. [1990]: *Unified Theories of Cognition*, Cambridge, MA: Harvard University Press.

Norton, A. [1995]: ‘Dynamics: An Introduction’, in R. Port and T. van Gelder (*eds.*), *Mind As Motion: Explorations In The Dynamics Of Cognition*, Cambridge, MA: MIT Press, pp. 45-68.

Port, R., and van Gelder, T. (*eds.*) [1995]: *Mind as Motion: Explorations in the Dynamics of Cognition*, Cambridge, MA: MIT Press.

Pylyshyn, Z. [1984]: *Computation and Cognition: Towards a Foundation for Cognitive Science*, Cambridge, MA: MIT Press.

Shapiro, L. A. [2000]: ‘Multiple Realizations’, *Journal of Philosophy*, **97**, pp. 635-54.

Smolensky, P. [1988]: ‘On the Proper Treatment of Connectionism’, *Behavioral and Brain Sciences*, **11**, pp. 1-74.

Thelen, E., and Smith, L. B. [1995]: *A Dynamic Systems Approach to the Development of Cognition and Action*. Cambridge, MA: MIT Press.

Townsend, J. T., and Busemeyer, J. R. [1995]: ‘Dynamic Representation of Decision Making’, in R. Port and T. van Gelder (*eds.*), *Mind as Motion: Explorations in the Dynamics of Cognition*, Cambridge, MA: MIT Press, pp. 101-20.

van Gelder, T. [1995]: ‘What Might Cognition Be, If Not Computation?’, *Journal of Philosophy*, **111**, pp. 345-81.

van Gelder, T., and Port, R. [1995]: ‘It’s About Time’, in R. Port and T. van Gelder (*eds.*), *Mind as Motion: Explorations in the Dynamics of Cognition*, Cambridge, MA: MIT Press, pp. 1-43.

---

[1] It is also a challenge to at least some forms of connectionist modeling. Connectionist networks have a dynamical systems interpretation, as Smolensky ([1988]) and van Gelder ([1995]) note, but they are also interpreted in computational terms. Furthermore, ‘time’ in a network simulation often means no more than the arbitrary counting off of steps during which activation is passed throughout

the network from one layer to the next. Insofar as time means only this, and there are few constraints placed on how step time is mapped onto the real physical time in which organisms behave, it seems that the dynamicists' arguments should be just as applicable to connectionism as to the symbolic theory. Since I contend that the argument from time is unsuccessful, this paper can also be read as a defense of connectionism against a possible dynamicist line of attack.

[2] Because this theory is orthodox in cognitive science, I will be brief in presenting it. For more detailed presentations, see Fodor ([1975]) and Pylyshyn ([1984]).

[3] I use all caps to mark the names of mental representations.

[4] I don't intend this to be a general definition of 'computation', since one of the things some dynamicists argue for is an analog conception of computation that does not require the manipulation of such symbolic tokens (Clark [1998]). Rather, this is computation as symbolic theorists understand it. No notion of computation more general than either the analog or the symbolic conception should be needed here.

[5] At least, this is what as much of mental functioning as can be captured by cognitive psychology consists in. There is certainly a need for some causal generalizations that do not govern representation to representation transitions; e.g., the transduction of ambient energy and the subsequent generation of representations of external stimuli. Further, there may be causes and effects of thoughts that are not representational but neurobiological. The more abundant these are, the smaller the explanatory domain of cognitive psychology proper. See Fodor ([1975]), pp. 200-1.

[6] Luce ([1984]) presents a commanding study of reaction and response time measures in psychology. He distinguishes reaction times from response times as species from genus, on the basis of whether the subject is informed that her time to respond is the focus of the study. Herein I will treat ‘reaction time’ as a generic term covering the use of chronometric dependent measures in psychological experiments; the response/reaction distinction will not be important for my purposes.

[7] The relevant grade of possibility here might be metaphysical or nomological. Nothing turns on whether it receives the weak or the strong interpretation, although historically symbolic theorists, like functionalists generally, have supposed the relevant possibility to be metaphysical. See Block ([1980]). On the other hand, Shapiro ([2000]) takes multiple realizability to be an empirical hypothesis, hence to be mainly about physical or nomological possibility.

[8] This passes over the difficult question of when two mechanisms should count as realizing the same cognitive kind. For discussion, see Shapiro ([2000]).

[9] My introduction to dynamics here will necessarily be compact, both for reasons of space and because the concepts are increasingly becoming familiar to philosophers of psychology. For further mathematical background, see Norton ([1995]). For other philosophical discussion, see Bechtel ([1998]), Clark ([1997], [1998]), Eliasmith ([1996]), Garson ([1996], [1997]), Keijzer ([1998]), and van Gelder ([1995]).

[10] This style of explanation leads Bechtel ([1998]) to call dynamical explanations 'total state' explanations.

[11] The emphasis on reduced-dimensional spaces also distinguishes pure dynamical systems theorists from connectionists. Connectionist networks can also be understood in state-space terms, but the dimensions of the spaces are often very high—in the case of a network's activation space, for example, as many dimensions are required as there are units. Smolensky ([1988]) was an early advocate of a dynamical systems understanding of connectionist networks. See Eliasmith ([1996]) for a comparison of the virtues of connectionist and pure dynamical models.

[12] Kelso actually calls this 'circular causality'. However, that unattractive term is also used to describe the intralevel behavior of two or more coupled dynamical systems. So, to forestall confusion, I will refer to the relations between the numerous interacting lower-level constituents of a system and their collective variables as 'interlevel influence'.

[13] Keijzer's formulation here is slightly tendentious, for two reasons. First, symbolic representations don't *prescribe* what the system ought to do, either. They cause it to behave as it does in virtue of having certain formal and semantic properties, but this is not obviously prescription, even though the content of a motor representation will partially describe the behavior that the command causes. Second, there is a sense in which control parameters do control behavior in a fairly straightforward way: how the organism behaves depends on their value, and the behavior would be different if the parameters had different values. Given this counterfactual test for

influence, control parameters are in a straightforward sense controllers, even if they lack rich structure and content.

[14] There are two caveats. First, this is only true as long as the system has not become stuck in a point attractor. I take it that this represents an undesirable state for a cognitive system, though. Second, a system may, in time, loop back and reoccupy the very same point in state space, thus reoccupying the same cognitive state. But the individuation conditions on cognitive states make this very unlikely.

[15] There is even a question of whether we should count the re-tokening of a cognitive state type at the next moment as a novel tokening or just the persisting of the previous token. But suppose for the sake of argument that there are principled grounds for making such a distinction here.

[16] The 'could' here is to remind us that more than discreteness is needed for symbolic modeling. The symbolic model also requires that the discrete states be combinable into syntactically structured wholes.

[17] Garson ([1996]) points out that analog neural networks that are somewhat noisy need not suffer reduced overall computational power. Hence, the noisiness of neurons might not impugn dynamical models whose computations rely on analog states. But the issue here isn't just whether noise degrades the performance of the system as a whole, but whether noise in a perceptual channel interferes with the delicate sensitivity that van Gelder and Port are imputing to that channel. The

system might suffer no overall loss of computational power even if its perceptual capacities were degraded.

[18] However, the operations that are basic to some model of a fragment of cognition are rarely thought to be the truly basic operations of the organism's overall cognitive architecture. Typically they are themselves tacitly assumed to be composed from still more basic operations of an unspecified sort.

[19] Churchland ([1989]) argues that the Language of Thought hypothesis (roughly the symbolic conception as I have sketched it here) is flawed because it purportedly entails a radical phylogenetic discontinuity between humans, whose cognition plausibly is linguiform in many respects, and the lower animals, whose cognition plausibly is not. It would be a terrific irony if dynamical systems theorists were also open to the charge of requiring at least two different theories in order to explain behavior across phylogenetic discontinuity.