

Yet another design for a brain? Review of Robert Port and Timothy van Gelder's *Mind as Motion: Explorations in the Dynamics of Cognition*. MIT/Bradford 1995.

To appear in *Philosophical Psychology*

Rick Grush

Philosophy-Neuroscience-Psychology Program, Washington University in St. Louis.

Introduction:

It is the aim of work in theoretical cognitive science to produce good theories of what exactly cognition amounts to, preferably theories which not only provide a framework for fruitful empirical investigation, but which also shed light on cognitive activity itself, which help us to understand our place, as cognitive agents, in a complex causally determined physical universe. The most recent such framework to gain significant fame is the so-called dynamical approach to cognition (henceforth DST, for *Dynamical Systems Theory*).

Explaining and exploring DST is the purpose of the collection *Mind as Motion:*

Explorations in the Dynamics of Cognition, edited by Robert Port and Timothy van Gelder.

This essay will investigate *Mind as Motion* (henceforth MM) as well as DST in some detail. In the first section I will briefly explain the DST paradigm and some of its historical precedents. The second section will focus on the treatment of time in DST and rival theories, which is singled out by Port and van Gelder as a, perhaps *the*, significant difference between DST and computational cognitive science (henceforth CCS). The third section will focus on the antirepresentationalism which is often embraced by DST researchers. Finally, section four will focus on the volume itself, and attempt to assess its import.

Section 1. The dynamical view of cognition

DST organizes itself around a set of mathematical modeling tools, and the conviction that the most revealing way to understand cognition is as embedded real-time adaptive activity. As such, DST is clearly a descendent of the cybernetics movement which reached its height the 50s and 60s. Indeed, Ashby's *Design for a Brain* (1952) explicitly adopts the mathematical apparatus now known as dynamical systems theory, as well as the view of cognitive behavior as adaptive real-time embedded activity. The similarities between Ashby's program and DST are astonishing, and it is no easy matter to isolate any significant conceptual progress. The most relevant difference is that the scientific

community now has technology, experimental equipment, and computer-facilitated modeling techniques which enable researchers to actually apply dynamical system theory to the analysis and modeling of behavior -- something that Ashby, Wiener, and their contemporaries could only dream of.

But this hardly marks a significant conceptual advance. Indeed, in *Design for a Brain*, Ashby outlined, 45 years ago, what would now be called the DST program (including a clear exposition of the mathematics of dynamical systems, an analysis of the Watt governor (one of van Gelder's favorite examples), and a statement and defense of the position that intelligence is manifested in real-time embedded activity). Though DST had not developed as a separate branch of mathematics when Ashby's book was published, most of the mathematical tools and techniques now known as DST had been around for quite a while. The second chapter of *Design for a Brain*, titled 'Dynamical Systems', explains the following notions with laudable clarity: a system as a set of state variables, phase spaces, trajectories through state space, state determined systems, and stability. In later chapters, he discusses the distinction between variables and parameters, and has a nice discussion of what amounts to parameter changes which can lead to drastic changes in the behavior of the system (these would now be called 'bifurcation points'). The only important elements missing in Ashby, but present in the current DST work, are chaos, and discrete systems, such as iterative maps. In my opinion, at least, when it comes to motivating the paradigm, and explaining the basics of the mathematical formalism, *Ashby simply did a better, clearer, and more interesting job.*

But enough history. Let's get to grips with today's DST paradigm. The view of cognition as embedded activity will be explored in the next section. For now we will look at the mathematical modeling apparatus of dynamical systems theory. Present-day dynamical systems theory is a very general mathematical framework which is used to describe how systems evolve over time. This involves three separate components: first, a description of a system as a set of state variables. The set of all possible state variable value combinations will form a *state space* (also known as a *phase space*). The state space can thus be represented as an n-dimensional space, where n is the number of state variables involved. The second component is an *ordering set*, which can be considered to be time, that is used to order successive stages in the evolution of the system. Two common choices are the positive integers (in which case the system will evolve in discrete steps, from t_0 to t_1 to t_2 , etc.), and the real numbers (in which case the system's evolution is continuous). The third component, often called the *dynamic*, is a rule for how the system evolves over time --

how, in other words, one determines the future positions in state space the system will occupy given the position it is currently in. The character of these rules will depend in large part on the choice of state variable types and the choice of the ordering set.

Consider the case of a simple mass-spring, in which a block of mass m , attached to a wall with a spring of stiffness k , and allowed to move without friction along one dimension of a surface. The state variables are position, x , and acceleration, $a = d^2x/dt^2$, and can take on continuous values. The ordering set is continuous real time, and the dynamic is given by the equation $d^2x/dt^2 = -(k/m)x$.

Dynamical systems theory is extremely general. In fact, to the degree that just about any system whatsoever is tractable it can be represented with this mathematical apparatus. But the reason we should adopt DST, its proponents argue, is not because of its generality, but rather because dynamical systems theory is better suited to describing and understanding cognition *as it really is* -- an environmentally embedded real time adaptive activity. That is, DST is not simply a movement in favor of a different mathematical apparatus, rather it is a paradigm for understanding cognition in a way very unlike the way it is modeled and understood by CCS.

Section 2. (*Ersatz*) time and (*ersatz*) time again.

The treatment of *time* in dynamical models of cognitive processes is paraded as *the* key difference, judging by the title of the editors' introductory article 'It's about time', setting them apart from computational models. Dynamical models, recall, include an ordering set, *time*, which is used to order successive stages in the evolution of the system. Presumably in most cases the real numbers are chosen for modeling cognitive processes, as well as other physical processes. In so doing, the argument continues, DST is able to do something which computational models cannot -- explain and model the real-time processes which result in adaptive behavior. Computational models, by contrast, will necessarily invoke time only in the weak sense of an ordering of discrete processing steps. In describing computational models of decision making, van Gelder and Port remark in their introduction:

...they say nothing at all about the temporal course of deliberation: how long it takes to reach a decision, how the decision one reaches depends on deliberation time, how a choice can appear more attractive at one time, less attractive at another, etc. They are intrinsically incapable of

such predictions, because *they leave time out of the picture*, replacing it only with *ersatz* "time": a bare, abstract sequence of symbol states...

The alternative must be an approach to the study of cognition which begins from the assumption that cognitive processes happen in time. *Real* time. Conveniently, there already is a mathematical framework for describing how processes in natural systems unfold in real time. It is *dynamics*.

This *ersatz* time, the editors correctly point out, is not essentially connected to real time, the time of real physical processes in the real world. Of course computational models may try to anchor themselves to real time by adding assumptions, such as that each processing step takes *n* msec. But, the editors argue, these assumptions are often unmotivated, and are typically chosen such as to make the model fit empirical results.

And indeed dynamical models are billed as something altogether different. Dynamical systems theory provides the possibility of using the reals as an ordering set, it has tools for understanding how processes unfold in time, the time of choice being the continuous time of physical processes. It is thus claimed that by weaving time into the theoretical fabric from stitch one, the *ad hoc* anchoring assumptions employed by computational models will be unnecessary.

Unfortunately, this characterization of the dynamical enterprise from the introduction has written a check which a number of the key contributions are unable to cash. A large portion of the models of 'higher' cognitive processes articulated in the book have exactly the same processing-step character as the vilified computational alternatives, even though the language, mathematics and illustrations used to present these models often obscure this fact. There are four contributions in MM which directly address aspects of so-called higher cognition, James T. Townsend and Jerome Busemeyer's 'Dynamic Representation of Decision-Making', Jeffrey L. Elman's 'Language as a Dynamical System', Jordan Pollack's 'The Induction of Dynamical Recognizers', and Jean Petitot's 'Morphodynamics and attractor syntax: Constituency in visual perception and cognitive grammar.' Of these, only the last is genuinely dynamic in the sense insisted upon by the editors.

To illustrate: Townsend and Busemeyer's model of decision making works by calculating a preference over its options, and comparing that preference to a threshold. If the preference exceeds the threshold for some action, then that action is taken. If the threshold is not exceeded, then a new preference is calculated. The interesting twist is that the model

typically takes many iterations before a decision threshold is exceeded, and the each preference calculation depends on the previous preference calculations. This dependence on previous outcomes allows one to plot the course of deliberation, that is, one can graph subsequent preference calculations as a function of the number of iterations computed before threshold is reached.

It should be immediately obvious from this description that this model has nothing at all to do with real time. Indeed, it operates as a sequence of calculations, each iteration of which can take as long or as short a time as the computer running the simulation requires. Of course, if one illicitly equates *processing iterations* with *time*, then the model looks like it is in fact doing something genuinely dynamic. And this is in fact what is done. All of the descriptions of the model's behavior are cleverly couched in terms of *time*, and only when one digs beneath the surface of the mathematics is it apparent that this is being used as a synonym for *processing step*. Consider the editors' introduction to this model (p. 101):

The DFT framework [Decision Field Theory, the model explicated by Townsend and Busemeyer -- RG], by contrast, sets out with the explicit aim of modeling the psychological processes involved in decision making. In this framework the system begins in a certain preference state with regard to certain choices, and this state evolves over time according to the dynamical equations which govern the relationship among factors such as the motivational value of an outcome and the momentary anticipated value of making a particular choice.

This description makes exactly the same error that the editors, in their introduction to the volume, hammer computational theories of cognition for making -- the substitution of *ersatz* time for *real* time.

The bait-and-switch is not limited to this particular contribution, however. Elman's model of syntactic processing has exactly the same character. It is a connectionist network which is fed, one at a time, representations of words. The output at each iteration is a prediction of the category of the next word. The output at each step depends not only on the input on that trial, but the inputs on the previous steps. But although the processing step, *ersatz* time character of this model is plainly evident, the text and graphs consistently, and frustratingly, refer to processing steps as time. This is not too surprising, given that in their introduction the editors themselves characterize connectionist work as "often underestimate[ing] the depth and pervasiveness of computational assumptions. Much standard connectionist work (e.g. modeling with layered backprop networks) is just a variation on computationalism, substituting activation patterns for symbols. This kind of

connectionism took some steps in the right direction, but mostly **failed to take the needed step *out of the computational mindset into time*...**" (p. 3, italic emphasis original, bold emphasis added).

Indeed, one graph (page 212) which plots the changing value of a hidden-unit principle component through 3 processing steps, labels the horizontal axis, which represents number of processing steps, as 'time', and more amazingly labels them in increments of one-half! That is, the horizontal 'time' axis labels are 1.0, 1.5, 2.0, 2.5, 3.0, even though this axis labels processing steps, and even though it makes no sense whatsoever to label an activation at 1.5 processing steps. Not surprisingly, the data points skip these *ersatz* half-intervals, because, of course, there is no data for these non-existent points. But the fact that the axis is so labeled, and so *misleadingly* labeled, is interesting. One wonders how this distinction, which is so crucial to the editors' own rhetoric, was simply missed.

Indeed, in their introduction, the editors' suggest the following diagnostic questions for determining when a theorist has illicitly substituted *ersatz* time for real time: "To see that integer 'times' in the Turing machine are not real times consider the following questions: What state was the machine in at time 1.5? How long was the machine in state 1? How long did it take for the machine to change from state 1 to state 2? *None of these questions are appropriate*, though they would be if we were talking about real amounts of time." (Emphasis original.) I leave to the readers as a take-home assignment the application of these questions to the models I have been discussing.

Care is needed to understand my point: these criticisms do not, I think, detract at all from the importance of Townsend and Busemeyer's model, or Elman's, or the insights behind them. Indeed, these models are quite interesting and perhaps even represent significant steps forward in understanding deliberative and linguistic processes. My point is the more specific one that the strong claims made by the editors concerning the character of these models, and their ability to address issues unaddressable by computational models, are simply inaccurate.

It will be pointed out that not all the contributions in this book have this *ersatz*-time character. To take but two examples, Randall Beer's models of insect leg motion, and Esther Thelen's analysis of infant motor control development, are both clearly and genuinely dynamical and real-time. But interestingly it is, by-and-large, the models which address higher cognitive processes which cling to *ersatz* time, while those true to the DST

ideals typically address so-called lower functions, such as motor control. For instance Thelen's discussion of the development of infant motor control, Beer's models of insect leg motion, Turvey and Carello's treatment of haptic perception and motor control, all fit nicely within the paradigm of real time continuous dynamical analyses. Townsend and Busemeyer's model of decision making and Elman's model of sentence parsing, as well as Pollack's model of legal-sequence recognizers, process information in discrete atemporal steps which fail the real-time challenges posed by the editors as miserably as do computational models.

Of course the irony is that Port and van Gelder take themselves to be overturning a historically rooted and inadequate *split* between higher cognition on the one hand, and perception and motor control on the other, a split aided and abetted by a generation of computational models of reasoning which treat cognition as a bare sequence of processing steps. But no sooner does one peek beneath the rhetoric and mathematical formalisms of the models presented in the volume, than one sees that they have simply, unwittingly, reaffirmed this rift.

Section 3. *Being-in vs. thinking about*

Even though the treatment of time in dynamical models is given the central role in defining DST's conceptual break with computational models, it is the paradigm's affinity for anti-representationalism which is, in philosophical circles at least, most infamous. While not taken as a defining or even necessary feature of DST, many of its most well-known proponents, including van Gelder, Randall Beer and Esther Thelen, have pushed anti-representationalism, to one degree or another, as a consequence of the dynamical view (See, e.g. van Gelder (1995), Beer (1995), Thelen and Smith (1994).) The path from dynamics to antirepresentationalism is not a clear one, and in fact I think that the representationalism/antirepresentationalism issue is orthogonal to the dynamicism/computationalism issue. But for now let's look at the antirepresentational arguments.

As I will reconstruct it, there are two parts to this argument. The first is a certain vision of cognition which seems to harmonize with seeing cognitive agents as dynamical systems. I will call this the *strong coupling thesis*. The second part consists of arguments to the effect that the sort of strong coupling envisioned is so tight as to leave no room for an

explanatorily significant notion of representation -- the tools of dynamical systems theory, together with some notion of adaptive embedded activity, being preferred.

The strong coupling thesis is that cognitive agents are embedded in an environment in such a way that both are in constant mutual interaction, and furthermore that cognition is most fruitfully seen as adaptive activity in an appropriate environment. The states of an agent which are responsible for its adaptive activity will be dynamically coupled to environmental states and features, with the consequence that any change in a state of the cognizer induces, in real continuous time, changes in some states of the environment (or more accurately, induces changes in those states' time derivatives). The affected environmental states, in turn, influence the temporal evolution of states in the cognizer in a similar way. Such an arrangement is sometimes called circular causation, or closed-loop control. The term of choice in DST circles being that the agent and its environment are dynamically *coupled*:

Since the nervous system, body, and environment are all continuously evolving and simultaneously influencing one another, the cognitive system cannot be simply the encapsulated brain; rather, it is a single unified system embracing all three. ...inner and outer process are *coupled*, so that both sets of processes are continually influencing each other. (editors' introduction, page 13)

As Randall Beer nicely puts it, in explaining a diagram similar to Figure 1 (where E stands for 'environment', A stands for 'agent', and S and M stand for 'sensory input' and 'motor output'): p. 131:

Note that feedback plays a fundamental role in the relationship between an agent and its environment. Any action that an agent takes affects its environment in some way through M, which in turn affects the agent itself through the feedback it receives from its environment via S. Likewise, the environment's effects on an agent through S are fed back through M in turn to affect the environment. Thus each of the two dynamical systems is continuously deforming the flow of the other (perhaps drastically if any coupling parameters cross bifurcation points in the receiving systems' parameter space), and therefore influencing its subsequent trajectory. ..

...an equally legitimate view is that the two coupled nonautonomous systems A and E are merely components of a single autonomous dynamical system U whose state variables are the union of the state variables of A and E and whose dynamical laws are given by all of the internal relations (including S and M) among the larger set of state variables and their derivatives.




Figure 1: An agent embedded in an environment.

The argument from this strong coupling thesis to antirepresentationalism is a bit murky, but the core thought is that the coupling which characterizes the relationship between cognitive agent and environment is such as to make dynamical systems theory the only genuinely enlightening set of tools for explaining adaptive success, representational and computational tools being inadequate. The best version of the argument I have seen (from van Gelder 1995) starts by analyzing a successful control mechanism, the Watt governor. In brief (see van Gelder 1995 for more details), the Watt Governor is a device which keeps the speed of a steam driven turbine constant by linking the turbine to a spindle from which hang two arms, and these two arms are linked to a throttle. As the turbine speeds up, the arms swing outward causing the throttle to close, and hence the speed of the turbine to decrease. But as the speed decreases, the arms fall, opening the throttle. There is a stable equilibrium point at which the speed of the turbine is close to constant, even under changing loads. After suggesting that its operation is a prime example of real-time problem solving activity and explaining its operation, van Gelder goes on to show how trying to explain the governor's success in terms of representations and computations is not only unnecessary, but crude and misleading.

The...deepest reason for supposing that the centrifugal governor is not representational is that, when we fully understand the relationship between engine speed and arm angle, we see that the notion of representation is just the wrong sort of conceptual tool to apply. There is no doubt that at all time the arm angle is in some interesting way related to the speed of the engine. This is the insight which leads people to suppose that the arm angle is a representation. Yet appropriately close examination shows exactly why the relationship cannot be one of representation... The arm angle and engine speed are at all times both determined

by, and determining, each other's behavior... There is nothing mysterious about this relationship: it is quite amenable to mathematical description. Yet it is much more subtle and complex than the standard concept of representation can handle... The real problem with describing the governor as a representational device, then, is that the relation of representing -- something standing in for some other state of affairs -- is too simple to capture the actual interaction between the governor and the engine. (p. 353)

Going into more of the details of van Gelder's argument would take us too far afield, and this would be unnecessary in any case. I believe that van Gelder's argument from tight coupling to antirepresentationalism is a good one. The problem is that the tight coupling thesis is simply and demonstrably (and, I would have thought, *obviously*) false -- which is to say among other things that the Watt governor is not a good model for cognitive activity. It is possible to demonstrate exactly how certain classes of dynamical systems can fail to be strongly coupled, and to do so in a way which makes it clear that parts of such systems are representations. Furthermore, there is abundant evidence that this manner of decoupling is actually employed by brains to represent and cognize. In brief, one can accept that cognition is a dynamical process, and also accept that it is representational, provided one denies the strong coupling thesis. I will now briefly outline how this can be done (more detailed accounts of the thesis which follows can be found in Grush (1995, 1997, in preparation); this thesis was outlined in qualitative, but still compelling, terms in Kenneth Craik's (1943) *The Nature of Explanation*).

Pre-theoretically (the supporting theory will be here in a few moments), many paradigmatically cognitive capacities seem to have nothing at all to do with being in a tightly coupled relationship with the environment. I can think about the St. Louis Arch while I'm sitting in a hot tub in southern California or while flying over the Atlantic ocean. In any non-distracting surroundings I can reason about which retirement fund would be best. Imagery, also, seems like it has nothing to do with being in a certain sort of environment, but rather has everything to do with purely internal processes.

The lesson is that in at least some cases, representing and cognizing is a matter of breaking the coupling, of getting out of the causal loop. But one cannot just split a coupled dynamic system in half and expect each half, especially the 'agent' half, to do anything at all (this is the thought behind the strong coupling thesis). So what is needed, in slightly more refined terms, is an executive part, C (for Controller), of an agent, A, which is in an environment E, decoupling from E, and coupling instead to some other system E' that stands in for E, in order for the agent to 'think about' E (see Figure 2). Cognitive agents are exactly those

which can selectively couple to either the 'real' environment, or to an environment model, or *emulator*, perhaps internally supported, in order to reason about what would happen if certain actions were undertaken with the real environment.

There is significant evidence that mental imagery works in exactly this way. During normal sensorimotor activity in an environment, subsystems of the central nervous system monitor the input-output relations of various aspects of the environment (including the musculoskeletal system). In terms of Beer's diagram, this would amount to anticipating what S will be at time t on the basis of M and what S was at $t-1$. It has been argued that there are CNS circuits, musculoskeletal emulators, which represent the musculoskeletal system (see Gerdes and Happee (1994); Kawato (1990); Wolpert et al. (1995)), and that these circuits are run off-line to support motor imagery. It has been shown how similar mechanisms can explain visual imagery, and counterfactual reasoning in general (see Grush (1995), (1997), (in preparation)).

Given this, with the differences between Figures 1 and 2 it becomes possible to illustrate exactly where DST goes wrong in both its antirepresentationalism and the strong coupling thesis. In diagram N, as in M, we have an agent in an environment. The difference is that the agent is not compelled to act only in coordination with that environment, but rather the agent can couple instead to an internal representation of the environment (E'). This coupling takes place via mock motor commands, M' (efferent copies), and mock feedback, S' (imagery). While it might be thought that this is an idea from computational cognitive science, specifically mental models, which is quite out of place here, this proposal can be cashed out readily in terms of dynamical systems theory. The agent A is a system consisting of state variables $(a_1, a_2, \dots a_i)$, etc. The environment consists of state variables $(e_1, e_2, \dots e_j)$, etc. But in the case I am supposing, there is an internal system E' , composed of systems variables $(e'_1, e'_2, \dots e'_j)$, etc., which are such that the executive part of A (C in the diagram, for *controller*) can couple to E' in exactly the same way it normally couples to E .




Figure 2: A cognitive agent in an environment which the agent can also internally represent.

Thinking about cognition and representation in this way has a number of advantages. First, it makes sense of intuitions regarding cognition and representation to the effect that they do not depend on being coupled to an environment. Second, it makes sense of the very reasonable theory of mental models, but without the computational baggage. Third, it makes no reference to ontologically magical barriers, such as skulls, in deciding what can and what cannot be a representation. Wherever an agent, or a subcomponent of an agent, can decouple from the normal environment and couple instead to a plug-compatible model, whether this model is external (as in the case of a flight simulator), or internal (as in the case of V-II), makes no difference. Representations are, on this view, quite sensibly entities which are used to stand in for something else. Fourth, the representations are semantically evaluable, they have a content, which is simply those entities of the real target which they stand in for. This is attractive because it counters a trend in cognitive science to see any theoretical machinery, such as attractors or principal components, as a representation.

Finally, regarding cognition and representation as manifestations of the use of internal emulators allows for a genuine distinction of kind between those systems which are really cognitive (like many animals), and those which are merely complex (like jet engines) and/or adaptive (like plants) -- those systems which can selectively de- and re-couple control loops from the environment to an internal emulator and back will be capable of internal representation and cognition.

In a similar way it allows us to distinguish those abilities of cognitive systems which are cognitive from those which are not. The bare apparatus of DST does not allow this, and is forced to treat all such abilities in the same way. As van Gelder and Port put it in their introduction: "Inner reasoning processes are no more essentially cognitive than the skillful execution of coordinated movement or the nature of the environment in which cognition takes place." (p. viii-ix). DST rhetoric notwithstanding, there can be a genuine distinction between cognitive and noncognitive adaptive behavior, even from within a generally dynamical perspective, once the requirement of strong coupling has been dropped so as to allow for internal emulation. E.g., on the present account, mental imagery is representational, while most motor control is not.

4. Assessing the book

The previous sections were devoted to some ax-grinding aimed at DST and some of its more radical claims. I would like now to turn to some more general issues, such as the likely impact of DST on the future of cognitive science, and (relatedly) the import of MM.

Regarding the importance of the DST paradigm in cognitive science, if one believes some of its adherents, it is the connectionism of the 90s, the radical new framework which will set the cognitive science community on fire and provide the long-sought-for tools for seeing clearly into the heart of cognition. This vision is over-optimistic. While DST is likely to remain on the cognitive science scene as an important source of modeling tools, along with connectionist and computational approaches, it will not sweep them aside. The fact that brain and its abilities are so incredibly complex that we have almost no idea how any of it works, together with the fact that these other approaches continue to turn out productive research, guarantee that, even if they turn out in the long run to be misguided or inadequate, there will be no compelling reason to abandon them for some time to come.

Furthermore, it seems unlikely that real continuous time will play as universal a role in understanding cognition as the editors assume -- many of the tasks that brains have evolved to accomplish are perhaps best described in algorithmic terms rather than dynamical ones. While examples of complex motor control, such as walking or catching a fly ball, surely rely on exact timing information, many other cases of adaptive activity do not. When building a nest, or a web, or an animal skin shelter, the exact timing of what one does is less important than getting

the various steps completed in the right order. Indeed, as I have pointed out in the second section, many of the models in MM rely only on algorithms invoking discrete step sequences.

Though the bulk of this review has been highly critical, I want to emphasize that my criticisms are aimed almost exclusively at what I have taken to be some rhetorical excesses committed by the book's editors and some members of the DST community at large. Perhaps I have made too big an issue of these excesses, and as a result failed to give due credit to the exciting and interesting features of DST in general, and MM in particular. If there is a justification for this, it is that, my criticisms notwithstanding, I recommend the book to anyone interested in theoretical cognitive science, and one can get, as it were, the 'pro' side from the book, its editors, and the researchers themselves. Even so, I would be culpably remiss if I did not at least briefly run through some of these points myself.

First off, the sheer generality, power, and richness of metaphors offered by DST probably insure that it will be an increasingly important tool, especially for understanding peripheral processes such as motor control and perception, and to a lesser degree core cognitive processes as well. Also, DST is likely to provide increasingly sophisticated models of neural dynamics, and as such is sure to play an important role in our eventual understanding of cognition in terms of neurophysiology. Indeed, the brief account of how cognizers represent through internal emulation, which I sketched in Section 3, is entirely compatible with much of the DST paradigm, and could even be seen as a DST account of internal representation. The bottom line is that DST is something which anyone with a serious interest in theoretical cognitive science should familiarize themselves with. Not because it is a completely radical break with other approaches, or will provide all the answers, but because, like connectionism, its approach is different, and a significant portion of work in the field will be couched in dynamical terms in the future.

And aside from the likely impact that DST will have on the future of theoretical cognitive science, there is an independent and more immediate reason to be enthusiastic about MM -- the quality, breadth and inherent interest of its contributions; features for which much credit must be given to the editors. Whether or not the details of each of the contributions consistently confirm the more radical DST claims is ultimately secondary to the fact that without exception the contributions are top-notch. Not only do they cover a broad range of research domains, including human motor control development, perception, language, and reasoning, but they often explore new approaches to problems in these domains in a lively and engaging way. And in the final analysis, it is hard to overestimate the value of new ideas such as these.

Acknowledgments:

I would like to thank the McDonnell foundation for financial support.

References:

Ashby, W.R. (1952) *Design for a brain*. London: Chapman and Hall.

Beer, R.D. (1995) A dynamical systems perspective on agent-environment interactions. *Artificial Intelligence* 72:173-215.

Craik, Kenneth (1943) *The Nature of Explanation*. Cambridge University Press.

Gerdes, V.G.J., and Happee, R. (1994) The use of an internal representation in fast goal-directed movements: a modeling approach. *Biological cybernetics* 70:513 - 524.

Grush, Rick (1995) *Emulation and Cognition*. Doctoral Dissertation, University of California, San Diego. At URL <http://www.artsci.wustl.edu/~rgrush/>

Grush, Rick (1997) The architecture of representation. *Philosophical Psychology* 10(1):5-23.

Grush, Rick (in preparation) *The neural construction of mind, language and reality. Volume I: Representation, objectivity, and content*.

Kawato, Mitsuo (1990). Computational schemes and neural network models for formation and control of multijoint arm trajectories. in Miller, W.T., Sutton, R.S. and Werbos, P.J., eds. *Neural Networks for Control*. MIT/Bradford.

Port, Robert, and van Gelder, Timothy (1995) *Mind as Motion: Explorations in the dynamics of cognition*. MIT/Bradford.

Thelen, E, and Smith, L.B. (1994) *A dynamic systems approach to the development of cognition and action*. Cambridge: MIT/Bradford.

van Gelder, Timothy (1995) What might cognition be, if not computation? *The Journal of Philosophy* 91(7):345-381.

van Gelder, Timothy (to appear) The Dynamical Hypothesis in Cognitive Science. *Behavioral and Brain Sciences*.

Wiener, Norbert (1948) *Cybernetics: or control and communication in the animal and machine*. New York: Wiley.

Wolpert, Daniel, Ghahramani, Zoubin, and Jordan, Michael (1995). An internal model for sensorimotor integration. *Science* 269:1880-1882.