1 It's About Time: An Overview of the Dynamical Approach to Cognition

Timothy van Gelder and Robert F. Port

How do we do what we do? How do we play tennis, have conversations, go shopping? At a finer grain, how do we recognize familiar objects such as bouncing balls, words, smiles, faces, jokes? Carry out actions such as returning a serve, pronouncing a word, selecting a book off the shelf? Cognitive scientists are interested in explaining how these kinds of extraordinarily sophisticated behaviors come about. They aim to describe *cognition*: the underlying mechanisms, states, and processes.

For decades, cognitive science has been dominated by one broad approach. That approach takes cognition to be the operation of a special mental *computer*, located in the brain. Sensory organs deliver up to the mental computer representations of the state of its environment. The system computes a specification of an appropriate action. The body carries this action out.

According to this approach, when I return a serve in tennis, what happens is roughly as follows. Light from the approaching ball strikes my retina and my brain's visual mechanisms quickly compute what is being seen (a ball) and its direction and rate of approach. This information is fed to a planning system which holds representations of my current goals (win the game, return the serve, etc.) and other background knowledge (court conditions, weaknesses of the other player, etc.). The planning system then infers what I must do: hit the ball deep into my opponent's backhand. This command is issued to the motor system. My arms and legs move as required.

In its most familiar and successful applications, the computational approach makes a series of further assumptions. Representations are static structures of discrete symbols. Cognitive operations are transformations from one static symbol structure to the next. These transformations are discrete, effectively instantaneous, and sequential. The mental computer is broken down into a number of modules responsible for different symbol-processing tasks. A module takes symbolic representations as inputs and computes symbolic representations as outputs. At the periphery of the system are input and output transducers: systems which transform sensory stimulation into input representations, and output representations into physical movements. The whole system, and each of its modules, operates cyclically: input, internal symbol manipulation, output.

The computational approach provides a very powerful framework for developing theories and models of cognitive processes. The classic work of pioneers such as Newell, Simon, and Minsky was carried out within it. Literally thousands of models conforming to the above picture have been produced. Any given model may diverge from it in one respect or another, but all retain most of its deepest assumptions. The computational approach is nothing less than a research paradigm in Kuhn's classic sense. It defines a range of questions and the form of answers to those questions (i.e., computational models). It provides an array of exemplars—classic pieces of research which define how cognition is to be thought about and what counts as a successful model. Philosophical tomes have been devoted to its articulation and defense. Unfortunately, it has a major problem: Natural cognitive systems, such as people, aren't computers.

This need not be very surprising. The history of science is full of episodes in which good theories were developed within bad frameworks. The Ptolemaic earth-centered conception of the solar system spawned a succession of increasingly sophisticated theories of planetary motion, theories with remarkably good descriptive and predictive capabilities. Yet we now know that the whole framework was structurally misconceived, and that any theory developed within it would always contain anomalies and reach explanatory impasses. Mainstream cognitive science is in a similar situation. Many impressive models of cognitive processes have been developed within the computational framework, yet none of these models are wholly successful even in their own terms, and they completely sidestep numerous critical issues. Just as in the long run astronomy could only make progress by displacing the earth from the center of the universe, so cognitive science has to displace the inner computer from the center of cognitive performance.

The heart of the problem is *time. Cognitive processes and their context unfold continuously and simultaneously in real time.* Computational models specify a discrete sequence of static internal states in arbitrary "step" time (t_1, t_2, etc.). Imposing the latter onto the former is like wearing shoes on your hands. You can do it, but gloves fit a whole lot better.

This deep problem manifests itself in a host of difficulties confronting particular computational models throughout cognitive science. To give just one example, consider how you might come to a difficult decision. You have a range of options, and consider first one, then another. There is hesitation, vacillation, anxiety. Eventually you come to prefer one choice, but the attraction of the others remains. Now, how are decision-making processes conceptualized in the computational worldview? The system begins with symbolic representations of a range of choices and their possible outcomes, with associated likelihoods and values. In a sequence of symbol manipulations, the system calculates the overall expected value for each choice, and determines the choice with the highest expected value. The system adopts that choice. End of decision. There are many variations on this basic "expected utility" structure. Different models propose different rules for calculating the choice the system adopts. But none of these models accounts perfectly for all the data on the choices that humans actually make. Like Ptolemaic theories of the planets, they become increasingly complex in attempting to account for residual anomalies, but for every anomaly dealt with another crops up elsewhere. Further, 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 symbolic states.

What is the alternative to the computational approach? In recent years, many people have touted *connectionism*—the modeling of cognitive processes using networks of neural units—as a candidate. But such proposals often underestimate the depth and pervasiveness of computationalist 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 leap *out* of the computational mindset and *into* time (see section 1.3, Relation to Connectionism, for elaboration).

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.* It just happens to be the single most widely used, most powerful, most successful, most thoroughly developed and understood descriptive framework in all of natural science. It is used to explain and predict phenomena as diverse as subatomic motions and solar systems, neurons and 747s, fluid flow and ecosystems. Why not use it to describe cognitive processes as well?

The alternative, then, is the *dynamical* approach. Its core is the application of the mathematical tools of dynamics to the study of cognition. Dynamics provides for the dynamical approach what computer science provides for the computational approach: a vast resource of powerful concepts and modeling tools. But the dynamical approach is more than just powerful tools; like the computational approach, it is a worldview. The cognitive system is not a computer, it is a dynamical system. It is not the brain, inner and encapsulated; rather, it is the whole system comprised of nervous system, body, and environment. The cognitive system is not a discrete sequential manipulator of static representational structures; rather, it is a structure of mutually and simultaneously influencing *change.* Its processes do not take place in the arbitrary, discrete time of computer steps; rather, they unfold in the *real* time of ongoing change in the environment, the body, and the nervous system. The cognitive system does not interact with other aspects of the world by passing messages or commands; rather, it continuously coevolves with them.

The dynamical approach is not a new idea: dynamical theories have been a continuous undercurrent in cognitive science since the field began (see section 1.4). It is not just a vision of the way things *might* be done; it's the way a great deal of groundbreaking research *has already* been carried out, and the amount of dynamical research undertaken grows every month. Much of the more recent work carried out under the connectionist banner is thoroughly dynamical; the same is true of such diverse areas as neural modeling, cognitive neuroscience, situated robotics, motor control, and ecological psychology. Dynamical models are increasingly prominent in cognitive psychology, developmental psychology, and even some areas of linguistics. In short, the dynamical approach is not just some new kid on the block; rather, to see that there is a dynamical approach is to see a new way of conceptually reorganizing cognitive science as it is currently practiced.

This introductory chapter provides a general overview of the dynamical approach: its essential commitments, its strengths, its relationship to other approaches, its history. It attempts to present the dynamical approach as a unified, coherent, plausible research paradigm. It should be noted, however, that dynamicists are a highly diverse group, and no single characterization would describe all dynamicists perfectly. Consequently, our strategy in this chapter is to characterize a kind of *standard* dynamicist position, one which can serve as a useful point of reference in understanding dynamical research.

The chapter is generally pitched in a quite abstract terms. Space limitations prevent us from going into particular examples in much detail. We urge readers who are hungry for concrete illustrations to turn to any of the 15 chapters of this book which present examples of actual dynamical research in cognitive science. It is essential for the reader to understand that *detailed demonstrations of all major points made in this overview are contained in chapters of the book.*

Before proceeding we wish to stress that our primary concern is only to understand *natural* cognitive systems—evolved biological systems such as humans and other animals. While the book is generally critical of the mainstream computational approach to the study of cognitive systems, it has no objections at all to investigations into the nature of computation itself, and into the potential abilities of computational systems such as take place in many branches of artificial intelligence (AI). While we think it *unlikely* that it will be possible to reproduce the kind of intelligent capacities that are exhibited by natural cognitive systems without also reproducing their basic noncomputational architecture, we take no stand on whether it is possible to program computers to exhibit these, or other, intelligent capacities.

1.1 WHAT IS THE DYNAMICAL APPROACH?

The heart of the dynamical approach can be succinctly expressed in the form of a very broad empirical hypothesis about the nature of cognition. For decades, the philosophy of cognitive science has been dominated by the *computational* hypothesis, that cognitive systems are a special kind of computer. This hypothesis has been articulated in a number of ways, but perhaps the most famous statement is Newell and Simon's *Physical Symbol System Hypothesis*, the claim that physical symbol systems (computers) are necessary and sufficient for intelligent behavior (Newell and Simon, 1976). According to this hypothesis, natural cognitive systems are intelligent by virtue of being physical symbol systems of the right kind. At this same level of generality, dynamicists can be seen as embracing the *Dynamical Hypothesis:* Natural cognitive systems are dynamical systems, and are best understood from the perspective of dynamics. Like its computational counterpart, the Dynamical Hypothesis forms a general framework within which detailed theories of particular aspects of cognition can be constructed. It can be empirically vindicated or refuted, but not by direct tests. We will only know if the Dynamical Hypothesis is true if, in the long run, the best theories of cognitive processes are expressed in dynamical terms.

The following sections explore the various components of the Dynamical Hypothesis in more detail.

Natural Cognitive Systems Are Dynamical Systems

What Are Dynamical Systems? The notion of dynamical systems occurs in a wide range of mathematical and scientific contexts, and as a result the term has come to be used in many different ways. In this section our aim is simply to characterize dynamical systems in the way that is most useful for understanding the dynamical approach to cognition.

Roughly speaking, we take dynamical systems to be systems with numerical states that evolve over time according to some rule. Clarity is critical at this stage, however, so this characterization needs elaboration and refinement.

To begin with, a *system* is a set of changing aspects of the world. The overall *state* of the system at a given time is just the way these aspects happen to be at that time. The *behavior* of the system is the change over time in its overall state. The totality of overall states the system might be in makes up its *state set*, commonly referred to as its *state space*. Thus the behavior of the system can be thought of as a sequence of points in its state space.

Not just any set of aspects of the world constitutes a system. A system is distinguished by the fact that its aspects somehow belong together. This really has two sides. First, the aspects must interact with each other; the way any one of them changes must depend on the way the others are. Second, if there is some *further* aspect of the world that interacts in this sense with anything in the set, then clearly it too is really part of the same system. In short, for a set of aspects to qualify as a system, they must be interactive and self contained: change in any aspect must depend on, and only on, other aspects in the set.

For example, the solar system differs from, say, the set containing just the color of my car and the position of my pencil, in that the position of any one planet makes a difference to where the other planets will be. Moreover, to a first approximation at least, the future positions of the planets are affected *only* by the positions, masses, etc., of the sun and other planets; there is nothing else we need take into account. By contrast, the position of my pencil is affected by a variety of other factors; in fact, it is unlikely that there is *any* identifiable system to which the position of my pencil (in all the vicissitudes of its everyday use) belongs.

Dynamical systems are special kinds of systems. To see *what* kind, we first need another notion, that of *state-determined* systems (Ashby, 1952). A system is state-determined only when its current state always determines a unique future behavior. Three features of such systems are worth noting. First, in such systems, the future behavior cannot depend in any way on whatever states the system might have been in *before* the current state. In other words, past history is irrelevant (or at least, past history only makes a difference insofar as it has left an effect on the current state). Second, the fact that the current state determines future behavior implies the existence of some *rule of evolution* describing the behavior of the system as a function of its current state. For systems we wish to understand, we always hope that this rule can be specified in some reasonably succinct and useful fashion. One source of constant inspiration, of course, has been Newton's formulation of the laws governing the solar system. Third, the fact that future behaviors are uniquely determined means that state space sequences can never fork. Thus, if we observe some system that proceeds in different ways at different times from the same state, we know we do not have a state-determined system.

The core notion of a state-determined system, then, is that of a self-contained, interactive set of aspects of the world such that the future states of the system are always uniquely determined, according to some rule, by the current state. Before proceeding, we should note an important extension of this idea, for cases in which changing factors external to the system do in fact affect how the system behaves. Suppose we have a set S of aspects $\{s_1, \ldots, s_m\}$ whose change depends on some further aspect s_n of the world, but change in s_n does not in turn depend on the state of S, but on other things entirely. Then, strictly speaking, neither S nor $S + s_n$ form systems, since neither set is self contained. Yet we can *treat* S as a state-determined system by thinking of the influence of s_n as built into its rule of evolution. Then the current state of the system *in conjunction with the rule* can be thought of as uniquely determining future behaviors, while the rule changes as a function of time. For example, suppose scientists discovered that the force of gravity has actually been fluctuating over time, though not in a way that depends on the positions and motions of the sun and planets. Then the solar system still forms a state-determined system, but one in which the rules of planetary motion must build in a gravitational constant that is changing over time. Technically, factors that affect, but are not in turn affected by, the evolution of a system are known as *parameters*. If a parameter changes over time, its changing effect can be taken into account in the rule of evolution, but then the rule itself is a function of time and the system is known as *nonhomogeneous*.

Now, according to some (e.g., Giunti, chapter 18), *dynamical systems* are really just *state-determined systems*. This identification is certainly valuable for some purposes. In fact, it is really this very inclusive category of systems (or at least, its abstract mathematical counterpart) that is studied by that branch of mathematics known as *dynamical systems theory*. Nevertheless, if our aim is to characterize the dynamical approach to cognition—and in particular, to contrast it with the computational approach—it turns out that a narrower definition is more useful. This narrower definition focuses on specifically *numerical* systems.

The word "dynamical" is derived from the Greek *dynamikos*, meaning "forceful" or "powerful." A system that is dynamical in this sense is one in which changes are a function of the *forces* operating within it. Whenever forces apply, we have accelerations or decelerations; i.e., there is change in the *rate* at which the states are changing at any given moment. The standard mathematical tools for describing rates of change are *differential equations*. These can be thought of as specifying the way a system is changing at any moment as a function of its state at that moment.[1] For example, the differential equation

$$\ddot{x} = -\frac{k}{m}x$$

describes the way (in ideal circumstances) a heavy object on the end of a spring will bounce back and forth by telling us the instantaneous acceleration ($\ddot{x}$) of the object as a function of its position (x); k and m are constants (parameters) for the spring tension and mass, respectively.

State-determined systems governed by differential equations are paradigm examples of dynamical systems in the current sense, but the latter category also includes other systems which are similar in important ways.

Whenever a system can be described by differential equations, it has n aspects or features (position, mass, etc.) evolving simultaneously and continuously in real time. Each of these features at a given point in time can be *measured* as corresponding to some real number. Consequently we can think of the overall state of the system as corresponding to an ordered set of n real numbers, and the state space of the system as isomorphic to a space of real numbers whose n dimensions are magnitudes corresponding (via measurement) to the changing aspects of the system. Sometimes this numerical space is also known as the system's state space, but for clarity we will refer to it as the system's *phase space*[2] (figure 1.1). The evolution of the system over time corresponds to a sequence of points, or trajectory, in its phase space. These sequences can often be described mathematically as functions of an independent variable, time. These functions are *solutions* to the differential equations which describe the behavior of the system.

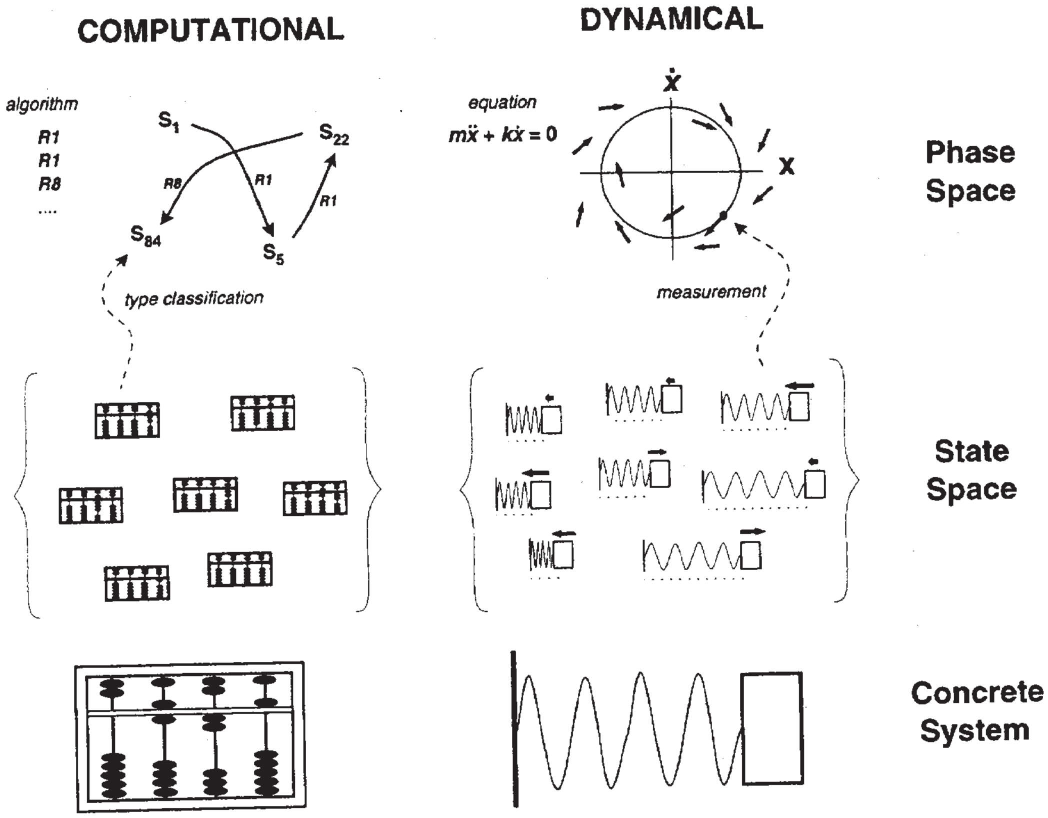

Figure 1.1 Mass-springs and computers are two different kinds of concrete state-determined system. (Our figure depicts an abacus; strictly speaking, the abacus would have to be automated to count as a computer.) Such systems are always in a particular state at a given point in time. This state is only one of many states that it *could* be in. The total set of possible states is commonly known as the system's *state space*. Corresponding to the state space is a set of abstract elements that is also commonly known as the system's state space, but which for clarity we refer to as its *phase space*. Possible states of the system are mapped onto elements of the phase space by some form of classification. In the computational case, tokens of symbols in the concrete system are classified into types, allowing the total state of the system to be classified as instantiating a particular configuration of symbol types. In the dynamical case, aspects of the system are measured (i.e., some yardstick is used to assign a number to each aspect), thereby allowing an ordered set of numbers to be assigned to the total state. Sequences of elements in the phase space can be specified by means of rules such as algorithms (in the computational case) and differential equations (in the dynamical case). A phase space and a rule are key elements of abstract state-determined systems. A concrete system realizes an abstract system when its states can be systematically classified such that the sequences of actual states it passes through mirror the phase space sequences determined by the rule. Typically, when cognitive scientists provide a model of some aspect of cognition, they provide an abstract state-determined system, such that the cognitive system is supposed to realize that abstract system or one relevantly like it.

Now, phase-space trajectories can be specified in a variety of ways. Differential equations constitute one particularly compact way of describing the shape of all possible trajectories in a given system. This kind of specification is useful for some purposes but not for others. A common alternative is to specify trajectories by means of a discrete mapping of any given point in the phase space onto another point. For example, perhaps the most-studied family of dynamical systems is the one whose rule is the "logistic equation" or "quadratic map" (Devaney, 1986)[3]:

$$F_\mu(x) = \mu x(1 - x)$$

For any particular value of the parameter μ, this equation determines a particular mapping of every point x in the phase space onto another point $F_\mu(x)$. A mapping like this can be regarded as giving us the state of a system at a subsequent point in time $(t + 1)$ if we know the state of the system at any given time (t). When the rule is written so as to bring this out, it is known as a *difference equation*, taking the general form

$$x(t + 1) = F(x(t))$$

If we take any given point in the phase space and apply ("iterate") the mapping many times, we obtain a phase-space trajectory.

Mathematicians and scientists often describe dynamical systems by means of discrete mappings rather than differential equations. In many cases these mappings are closely related to particular differential equations describing essentially the same behavior. This is not always the case, however. Consequently, a more liberal definition of *dynamical system* is: any state-determined system with a numerical phase space and a rule of evolution (including differential equations and discrete maps) specifying trajectories in this space.

These systems, while only a subset of state-determined systems in general, are the locus of dynamical research in cognitive science. They find their most relevant contrast with *computational* systems. These systems have states that are configurations of symbols,[4] and their rules of evolution specify transformations of one configuration of symbols into another. Whereas the phase space of a dynamical system is a *numerical* space, the phase space of a computational system is a space of configurations of *symbol types*, and trajectories are sequences of such configurations.

Why is it that dynamical systems (in our sense) are the ones chosen for study by dynamicists in cognitive science? Here we briefly return to the traditional idea that dynamics is a matter of forces, and therefore essentially involves rates of change. In order to talk about rates of change, we must be able to talk about *amounts* of change in *amounts* of time. Consequently, the phase space must be such as to allow us to say *how far* the state is changing, and the *time* in which states change must involve real *durations*, as opposed to a mere linear ordering of temporal points.

Now, these notions make real sense in the context of dynamical systems as defined here. Numerical phase spaces can have a metric that determines distances between points. Further, if the phase space is rich enough (e.g., dense) then between any two points in the phase space we can find other points, and so we can talk of the state of the system at any time between any two other times. Thus the notion of time in which the system operates is also one to which a substantial notion of "length" can be applied; in other words, it comes to possess some of the same key mathematical properties as *real* time. Note that neither of these properties is true of computational systems such as Turing machines; there, there is no natural notion of distance between any two total states of the system, and "time" (t_1, t_2, etc.) is nothing more than order. Consequently it is impossible to talk of how fast the state of the system is changing, and as a matter of fact, nobody ever tries; the issue is in a deep way irrelevant.

The importance of being able to talk about rates of change is that all actual processes in the real world (including cognitive processes) do in fact unfold at certain rates in real time. Further, for many such systems (including cognitive systems) *timing* is essential: they wouldn't be able to function properly unless they got the fine details of the timing right. Therefore, in order to provide adequate scientific descriptions of the behavior of such systems, we need to understand them as systems in which the notion of rates of change makes sense (see Cognition and Time, below). Dynamicists in cognitive science propose dynamical models in the current sense because they are such systems. It may well be that there are other, less well-known mathematical frameworks within which one could model change in real time without using specifically numerical systems. As things stand, however, dynamical systems in cognitive science are in fact state-determined numerical systems.

A wide variety of fascinating questions can be raised about the relations between dynamical and computational systems. For example, what is the relationship between an ordinary digital computer and the underlying electrical dynamical system that in some sense makes it up? Or, what is the relation between a dynamical system and a computational simulation or emulation of it? Even more abstractly, how "powerful" is the class of dynamical systems, in comparison with computational systems? However, we must be very careful not to allow the fact that *there are* many such relationships, some of them quite intimate, to blind us to an important philosophical, and ultimately practical, truth: dynamical and computational systems are fundamentally different *kinds* of systems, and hence the dynamical and computational approaches to cognition are fundamentally different in their deepest foundations.

Natural Cognitive Systems as Dynamical Systems Describing natural phenomena as the behavior of some dynamical system lies at the very heart of modern science. Ever since Newton, scientists have been discovering more and more aspects of the natural world that constitute dynamical systems of one kind or another. Dynamicists in cognitive science are claiming that yet another naturally occurring phenomenon, *cognition*, is the behavior of an appropriate kind of dynamical system. They are thus making exactly the same kind of claim for cognitive systems as scientists have been making for so many other aspects of the natural world. In the Dynamical Hypothesis, this is expressed as the idea that natural cognitive systems *are* dynamical systems.

Demonstrating that some aspect of the world constitutes a dynamical system requires picking out a relevant set of quantities, and ways of measuring them, such that the resulting phase-space trajectories conform to some specifiable rule. These trajectories must correspond to the behaviors of theoretical interest. So, if we are interested in *cognitive* systems, then the behaviors of interest are their *cognitive performances* (perceiving, remembering, conversing, etc.), and it is *these* behaviors, at their characteristic time scales, that must unfold in a way described by the rule of evolution. Consequently, the claim that cognitive systems are dynamical systems is certainly not trivial. Not everything is a dynamical system, and taking some novel phenomenon and showing that it *is* the behavior of a dynamical system is always a significant scientific achievement. If the Dynamical Hypothesis is in fact true, we will only know this as a result of much patient scientific work.[5]

Natural cognitive systems are enormously subtle and complex entities in constant interaction with their environments. It is the central conjecture of the Dynamical Hypothesis that these systems constitute single, unified dynamical systems. This conjecture provides a general theoretical orientation for dynamicists in cognitive science, but it has not been (and in fact may never be) demonstrated in detail, for nobody has specified the relevant magnitudes, phase space, and rules of evolution for the entire system. Like scientists confronting the physical universe as a whole, dynamicists in cognitive science strive to isolate particular aspects of the complex, interactive totality that are relatively self-contained and can be described mathematically. Thus, in practice, the Dynamical Hypothesis reduces to a series of more specific assertions, to the effect that particular aspects of cognition are the behavior of distinct, more localized systems. For example, Turvey and Carello (see chapter 13) focus on our ability to perceive the shape of an object such as a hammer simply by wielding it. They show how to think of the wielding itself as a dynamical system, and of perception of shape as attunement to key parameters of this system. The Dynamical Hypothesis, that entire cognitive systems constitute dynamical systems, is thus comparable to the Laplacean hypothesis that the entire physical world is a single dynamical system.

Many cognitive processes are thought to be distinguished from other kinds of processes in the natural world by the fact that they appear to depend crucially on *knowledge* which must somehow be stored and utilized. At the heart of the computational approach is the idea that this knowledge must be *represented*, and that cognitive processes must therefore be operations on these representations. Further, the most powerful known medium of representation is symbolic, and hence cognitive processes must manipulate symbols, i.e., must be computational in nature.

In view of this rather compelling line of thought, it is natural to ask: How can dynamicists, whose models do *not* centrally invoke the notion of representation, hope to provide theories of paradigmatically *cognitive* processes? If cognition depends on *knowledge*, how can there be a dynamical approach to *cognition*? The answer is that, while dynamical models are not *based* on transformations of representational structures, they allow plenty of room for representation. A wide variety of aspects of dynamical models can be regarded as having a representational status: these include states, attractors, trajectories, bifurcations, and parameter settings. So dynamical systems can store knowledge and have this stored knowledge influence their behavior. The crucial difference between computational models and dynamical models is that in the former, the rules that govern how the system behaves are defined over the entities that have representational status, whereas in dynamical models, the rules are defined over numerical states.[6] That is, dynamical systems can be representational without having their rules of evolution defined over representations. For example, in simple connectionist associative memories such as that described in Hopfield (1982), representations of stored items are point attractors in the phase space of the system. Recalling or recognizing an item is a matter of settling into its attractor, a process that is governed by purely numerical dynamical rules.

The Nature of Cognitive Systems The claim that cognitive systems are computers, and the competing claim that natural cognitive systems are dynamical systems, each forms the technical core of a highly distinctive vision of the nature of cognitive systems.

For the computationalist, the cognitive system is the brain, which is a kind of control unit located inside a body which in turn is located in an external environment. The cognitive system interacts with the outside world *via* its more direct interaction with the body. Interaction with the environment is handled by sensory and motor *transducers*, whose function is to translate between the *physical* events in the body and the environment, on the one hand, and the *symbolic* states, which are the medium of cognitive processing. Thus the sense organs convert physical stimulation into elementary symbolic representations of events in the body and in the environment, and the motor system converts symbolic representations of actions into movements of the muscles. Cognitive episodes take place in a cyclic and sequential fashion; *first* there is sensory input to the cognitive system, *then* the cognitive system algorithmically manipulates symbols, coming up with an output which *then* causes movement of the body; the whole cycle then begins again. Internally, the cognitive system has a modular, hierarchical construction; at the highest level, there are modules corresponding to vision, language, planning, etc., and each of these modules breaks down into simpler modules for more elementary tasks. Each module replicates in basic structure the cognitive system as a whole; thus, the modules take symbolic representations as inputs, algorithmically manipulate those representations, and deliver a symbolic specification as output. Note that because the cognitive system traffics only in symbolic representations, the body and the physical environment can be dropped from consideration; it is possible to study the cognitive system as an autonomous, bodiless, and worldless system whose function is to transform input representations into output representations.

Now, the dynamical vision differs from this picture at almost every point. As we have seen, dynamical systems are complexes of parts or aspects which are all evolving in a continuous, simultaneous, and mutually determining fashion. If cognitive systems are dynamical systems, then they must likewise be complexes of interacting change. 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. The cognitive system does not interact with the body and the external world by means of periodic symbolic inputs and outputs; rather, inner and outer processes are *coupled*, so that both sets of processes are continually influencing each other. Cognitive processing is not cyclic and sequential, for all aspects of the cognitive system are undergoing change all the time. There is a *sense* in which the system is modular, since for theoretical purposes the total system can be broken down into smaller dynamical subsystems responsible for distinct cognitive phenomena. Standardly these smaller systems are coupled, and hence coevolving, with others, but significant insight can be obtained by "freezing" this interaction and studying their independent dynamics. Of course, cognitive performances do exhibit many kinds of sequential character. Speaking a sentence, for example, is behavior that has a highly distinctive sequential structure. However, in the dynamical conception, any such sequential character is something that emerges over time as the overall trajectory of change in an entire system (or relevant subsystem) whose rules of evolution specify not sequential change but rather simultaneous, mutual coevolution.

Natural Cognitive Systems Are Best Understood Using Dynamics

In science, as in home repair, the most rapid progress is made when you have the right tools for the job. Science is in the business of describing and explaining the natural world, and has a very wide range of conceptual and methodological tools at its disposal. Computer science provides one very powerful collection of tools, and these are optimally suited for understanding complex systems of a particular kind, namely *computational* systems. If *cognitive* systems are computational systems, then they will be best understood by bringing these tools to bear. If the Dynamical Hypothesis is right, however, then the most suitable conceptual tools will be those of *dynamics*. So, whereas in the previous sections we described *what it is* for natural cognitive systems to *be* dynamical systems, in the following discussion we describe what is involved in applying dynamics in *understanding* such systems.

What Is Dynamics? Dynamics is a very broad field overlapping both pure and applied mathematics. For current purposes, it can be broken down into two broad subdivisions. *Dynamical modeling* is describing natural phenomena *as* the behavior of a dynamical system in the sense outlined in the previous discussion. It involves finding a way of isolating the relevant system, a way of measuring the states of the system, and a mathematical rule, such that the phenomena of interest unfold in exactly the way described by the rule. Obviously, effective dynamical modeling involves considerable exploration of both the real system being studied, and the mathematical properties of the governing equations.

Dynamical systems theory is the general study of dynamical systems. As a branch of pure mathematics, it is not directly concerned with the empirical description of natural phenomena, but rather with abstract mathematical structures. Dynamical systems theory is particularly concerned with complex systems for which the *solutions* of the defining equations (i.e., functions that specify trajectories as a function of time) are difficult or impossible to write down. It offers a wide variety of powerful concepts and tools for describing the general properties of such systems. Perhaps the most distinctive feature of dynamical systems theory is that it provides a *geometric* form of understanding: behaviors are thought of in terms of locations, paths, and landscapes in the phase space of the system.[7]

Some natural phenomena can be described as the evolution of a dynamical system governed by particularly straightforward equations. For such systems, the traditional techniques of dynamical modeling are sufficient for most explanatory purposes. Other phenomena, however, can only be described as the behavior of systems governed by nonlinear equations for which solutions may be unavailable. Dynamical systems theory is essential for the study of such systems. With the rapid development in the twentieth century of the mathematics of dynamical systems theory, an enormous range of natural systems have been opened up to scientific description. There is no sharp division between dynamical modeling and dynamical systems theory, and gaining a full understanding of most natural systems requires relying on both bodies of knowledge.

Understanding Cognitive Phenomena Dynamically Dynamics is a large and diverse set of concepts and methods, and consequently there are many different ways that cognitive phenomena can be understood dynamically. Yet they all occupy a broadly dynamical perspective, with certain key elements.

At the heart of the dynamical perspective is *time*. Dynamicists always focus on the details of how behavior unfolds in real time; their aim is to describe and explain the temporal course of this behavior. The beginning point and the endpoint of cognitive processing are usually of only secondary interest, if indeed they matter at all. This is in stark contrast with the computationalist orientation, in which the primary focus is on input-output relations, i.e., on what output the system delivers for any given input.

A second key element of the dynamical perspective is an emphasis on *total state*. Dynamicists assume that all aspects of a system are changing simultaneously, and so think about the behavior of the system as a matter of how the total state of a system is changing from one time to the next. Computationalists, by contrast, tend to suppose that most aspects of a system (e.g., the symbols stored in memory) do *not* change from one moment to the next. Change is assumed to be a local affair, a matter of replacement of one symbol by another.

Because dynamicists focus on how a system changes from one total state to another, it is natural for them to think of that change as a matter of movements in the *space* of all possible total states of the system; and since the phase spaces of their systems are numerical, natural notions of *distance* apply. Thus, dynamicists conceptualize cognitive processes in *geometric* terms. The distinctive character of some cognitive process as it unfolds over time is a matter of how the total states the system passes through are spatially located with respect to one another and the dynamical landscape of the system.

Quantitative Modeling Precise, quantitative modeling of some aspect of cognitive performance is always the ultimate goal of dynamical theorizing in cognitive science. Such research always requires two basic components: data and model. The data take the form of a time series: a series of measurements of the phenomenon to be understood, taken as that phenomenon unfolds over time. The model is a set of equations and associated phase space. The modeling process is a matter of distilling out the phenomenon to be understood, obtaining the time-series data, developing a model, and *interpreting* that model as capturing the data (i.e., setting up correspondences between the numerical sequences contained in the model and those in the data). When carried out successfully, the modeling process yields not only precise *descriptions* of the existing data but also *predictions* which can be used in evaluating the model.

For an excellent example of quantitative dynamical modeling, recall the process of reaching a decision described briefly in the introductory paragraphs. We saw that traditional computational (expected-utility theory) approaches to decision-making have had some measure of success in accounting for what decisions are actually reached, but say nothing at all about any of the temporal aspects of the deliberation process. For Busemeyer and Townsend (Busemeyer and Townsend, 1993; see also chapter 4), by contrast, describing these temporal aspects is a central goal. Their model of decision-making is a dynamical system with variables corresponding to quantities such as values of consequences and choice preferences. The model describes the multiple simultaneous *changes* that go on in an individual decision-maker in the process of coming to a decision. It turns out that this model not only recapitulates the known data on *outcomes* as well as or better than traditional computational models; it also explains a range of temporal phenomena such as the dependence of preference on deliberation time, and makes precise predictions which can be experimentally tested.

Qualitative Modeling Human cognitive performance is extraordinarily diverse, subtle, complex, and interactive. Every human behaves in a somewhat different way, and is embedded in a rich, constantly changing environment. For these kinds of reasons (among others), science has been slow in coming to be able to apply to cognition the kinds of explanatory techniques that have worked so successfully elsewhere. Even now, only a relatively small number of cognitive phenomena have been demonstrated to be amenable to precise, quantitative dynamical modeling. Fortunately, however, there are other ways in which dynamics can be used to shed light on cognitive phenomena. Both the data time series and the mathematical model that dynamical modeling requires can be very difficult to obtain. Even without an elaborate data time series, one can study a mathematical model which exhibits behavior that is at least *qualitatively* similar to the phenomena being studied. Alternatively, in the absence of a precise mathematical model, the language of dynamics can be used to develop qualitative dynamical descriptions of phenomena that may have been recorded in a precise data time series (see Dynamical Description, below).

Cognitive scientists can often develop a sophisticated understanding of an area of cognitive functioning independently of having any elaborate data time series in hand. The problem is then to understand what kind of system might be capable of exhibiting that kind of cognitive performance. It can be addressed by specifying a mathematical dynamical model and comparing its behavior with the known empirical facts. If the dynamical model and the observed phenomena agree sufficiently in broad qualitative outline, then insight into the nature of the system has been gained.

Elman's investigations into language processing are a good example of qualitative dynamical modeling (Elman, 1991; see also chapter 8). In broad outline, at least, the distinctive complexity of sentences of natural language is well understood, and psycholinguistics has uncovered a wide range of information on human abilities to process sentences. For example, it is a widely known fact that most people have trouble processing sentences that have three or more subsidiary clauses embedded centrally within them. In an attempt to understand the internal mechanisms responsible for language use, Elman investigates the properties of a particular class of connectionist dynamical systems. When analyzed using dynamical concepts, these models turn out to be in broad agreement with a variety of general constraints in the data, such as the center-embedding limitation. This kind of agreement demonstrates that it is possible to think of aspects of our linguistic subsystems in dynamical terms, and to find there a basis for some of the regularities. This model does not make precise temporal predictions about the changing values of observable variables, but it does make testable *qualitative* predictions about human performance.

Often, the system one wants to understand can be observed to exhibit any of a variety of highly distinctive *dynamical* properties: asymptotic approach to a fixed point, the presence or disappearance of maxima or minima, catastrophic jumps caused by small changes in control variables, oscillations, chaotic behavior, hysteresis, resistance to perturbation, and so on. Such properties can be observed even without knowing the specific equations which in fact govern the evolution of the system. They are, however, a particularly rich source of constraints for the process of qualitative dynamical modeling, for they narrow down considerably the classes of equations that can exhibit qualitatively similar behavior.

Dynamical Description In another kind of situation, we may or may not have good time-series data available for modeling, but the complexity of the phenomena is such that laying down the equations of a formal model adequate to the data is currently not feasible. However, even here dynamics may hold the key to advances in understanding, because it provides a general conceptual apparatus for understanding the way systems—including, in particular, nonlinear systems—change over time. In this kind of scenario it is *dynamical systems theory* which turns out to be particularly useful.

For example, Thelen (see chapter 3) is concerned with understanding the development, over periods of months and even years, of basic motor skills such as reaching out for an object. At this stage, no satisfactory mathematical model of this developmental process is available. Indeed, it is still a major problem to write down equations describing just the basic movements themselves! Nevertheless, adopting a dynamical perspective can make possible descriptions which cumulatively amount to a whole new way of understanding how motor skills can emerge and change, and how the long-term developmental process is interdependent with the actual exercise of the developing skills themselves. From this perspective, particular actions are conceptualized as attractors in a space of possible bodily movements, and development of bodily skills is the emergence, and change in nature, of these attractors over time under the influence of factors such as bodily growth and the practice of the action itself. Adopting this general perspective entails significant changes in research methods. For example, Thelen pays close attention to the exact shape of individual gestures at particular intervals in the developmental process, and focuses on the specific changes that occur in each individual subject rather than the gross changes that are inferred by averaging over many subjects. 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.

1.2 WHY DYNAMICS?

Why should we believe the Dynamical Hypothesis? Ultimately, as mentioned above, the proof of the pudding will be in the eating. The Dynamical Hypothesis is correct only if sustained empirical investigation shows that the most powerful models of cognitive processes take dynamical form. Although there are already dynamical models—including many described in this book

—which are currently the best available in their particular area, the jury is still out on the general issue. Even if the day of final reckoning is a long way off, however, we can still ask whether the dynamical approach is *likely* to be the more correct, and if so, why.

The dynamical approach certainly begins with a huge head start. Dynamics provides a vast resource of extremely powerful concepts and tools. Their usefulness in offering the best scientific explanations of phenomena throughout the natural world has been proved again and again. It would hardly be a surprise if dynamics turned out to be the framework within which the most powerful descriptions of cognitive processes were also forthcoming. The conceptual resources of the computational approach, on the other hand, are known to describe only one category of things in the physical universe: manmade digital computers. Even this success is hardly remarkable: digital computers were designed and constructed by us in accordance with the computational blueprint. It is a bold and highly controversial speculation that these same resources might also be applicable to natural cognitive systems, which are evolved biological systems in constant causal interaction with a messy environment.

This argument for the dynamical approach is certainly attractive, but it is not grounded in any way in the specific nature of cognitive systems. What we really want to know is: What general things do we *already know* about the nature of *cognitive* systems that suggest that dynamics will be the framework within which the most powerful models are developed?

We know, at least, these very basic facts: that cognitive processes always unfold in real time; that their behaviors are pervaded by *both* continuities and discretenesses; that they are composed of multiple subsystems which are simultaneously active and interacting; that their distinctive kinds of structure and complexity are not present from the very first moment, but emerge over time; that cognitive processes operate over many time scales, and events at different time scales interact; and that they are embedded in a real body and environment. The dynamical approach provides a natural framework for the description and explanation of phenomena with these broad properties. The computational approach, by contrast, either ignores them entirely or handles them only in clumsy, ad hoc ways.[8]

Cognition and Time

The argument presented here is simple. Cognitive processes always unfold in real time. Now, computational models specify only a postulated *sequence* of states that a system passes through. Dynamical models, by contrast, specify in detail not only *what* states the system passes through, but also how those states unfold in real time. This enables dynamical models to explain a wider range of data for any cognitive functions, and to explain cognitive functions whose dependence on real time is essential (e.g., temporal pattern processing).

When we say that cognitive processes unfold in real time, we are really saying two distinct things. 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. For an example of a process unfolding in real time, consider the movement of your arm as it swings beside you. At every one of an infinite number of instants in time from the beginning to the end of the motion, there is a position which your arm occupies. No matter how finely time is sampled, it makes sense to ask what position your arm occupies at every sampled point. The same is true of cognitive processes. As you recognize a face, or reason through a problem, or throw a ball, various aspects of your total cognitive system are undergoing change in real time, and no matter how finely time is sampled, there is a state of the cognitive system at each point. This is really just an obvious and elementary consequence of the fact that cognitive processes are ultimately physical processes taking place in real biological hardware.

The second thing we mean by saying that cognitive processes unfold in real time is that—as a consequence of the first point—*timing* always matters. A host of questions about the way the processes happen *in* time make perfect sense: questions about rates, durations, periods, synchrony, and so forth. Because cognitive processes happen in time, they cannot take too little time or too much time. The system must spend an appropriate amount of time in the vicinity of any given state. The timing of any particular operation must respect the rate at which other cognitive, bodily, and environmental processes are taking place. There are numerous subtleties involved in correct timing, and they are all real issues when we consider real cognitive processing.

Since cognitive processes unfold in real time, any framework for the description of cognitive processes that hopes to be fully adequate to the nature of the phenomena must be able to describe not merely *what* processes occur but *how* those processes unfold in time. Now, dynamical models based on differential equations are the preeminent mathematical framework science uses to describe how things happen in time. Such models specify how change in state variables at any instant depends on the current values of those variables themselves and on other parameters. Solutions to the governing equations tell you the state that the system will be in at any point in time, as long as the starting state and the amount of elapsed time are known. The use of differential equations presupposes that the variables change smoothly and continuously, and that time itself is a real-valued quantity. It is, in short, of the *essence* of dynamical models of this kind to describe how processes unfold, moment by moment, in real time.

Computational models, by contrast, specify only a bare sequence of states that the cognitive system goes through, and tell us nothing about the timing of those states over and above their mere order. Consider, for example, that paradigm of computational systems, the Turing machine.[9] Every Turing machine passes through a series of discrete symbolic states, one after another.

We talk about the state of the machine at time 1, time 2, and so forth. However, these "times" are not points in real time; they are merely indices which help us keep track of the order that states fall into as the machine carries out its sequence of computational steps. We use the integers to index states because they have a very familiar order and there are always as many of them as we need. However, we mustn't be misled into supposing that we are talking about *amounts* of time or *durations* here. Any other ordered set (e.g., people who ran the Boston Marathon, in the order they finished) would, in theory, do just as well for indexing the states of a Turing machine, though in practice they would be very difficult to use. To see that the 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.

Now, let us suppose we have a particular Turing machine which adds numbers, and we propose this machine as a model of the cognitive processes going on in real people when they add numbers in their heads. The model specifies a sequence of symbol manipulations, passing from one discrete state to another; we suppose that a person passes through essentially the same sequence of discrete states. Note, however, that the Turing machine model is inherently incapable of telling us anything at all about the *timing* of these states and the transitions from one state to another. The model just tells us "first this state, then that state ..."; it makes no stand on how long the person will be in the first state, how fast the transition to the second state is, and so forth; it cannot even tell us what state the person will be in halfway between the time it enters the first state and the time it enters the second state, for questions such as these make no sense in the model.

Of course, even as far as computational models go, Turing machines do not make good models of cognitive processes. But the same basic points hold true for all standard computational models. LISP programs, production systems, generative grammars, and so forth, are all intrinsically incapable of describing the fine temporal structure of the way cognitive processes unfold, because all they specify—indeed, all they *can* specify—is *which* states the system will go through, and in what order. To see this, just try picking up any mainstream computational model of a cognitive process—of parsing, or planning, for example—and try to find any place where the model makes any commitment at all about such elementary temporal issues as how much time each symbolic manipulation takes. One quickly discovers that computational models simply aren't in that business; they're not dealing with time. "Time" in a computational model is not real time, it is mere order.

Computationalists do sometimes attempt to extract from their models implications for the timing of the target cognitive processes. The standard and most appropriate way to do this is to assume that each computational step takes a certain chunk of real time (say, 10 ms).[10] By adding assumptions of this kind we can begin to make some temporal predictions, such as that a particular computational process will take a certain amount of time, and that a particular step will take place some number of milliseconds after some other event. Yet the additional temporal assumptions are completely ad hoc; the theorist is free to choose the step time, for example, in any way that renders the model more consistent with the psychological data.[11] In the long run, it is futile to attempt to weld temporal considerations onto an essentially atemporal kind of model. If one professes to be concerned with temporal issues, one may as well adopt a modeling framework which builds temporal issues in from the very beginning—i.e., take up the dynamical approach.

One refuge for the computationalist from these arguments is to insist that certain physical systems are such that they can be described at an abstract level where temporal issues can be safely ignored, and that the most tractable descriptions of these systems must in fact take place at that level. This claim is clearly true of ordinary desktop digital computers; we standardly describe their behavior in algorithmic terms in which the precise details of timing are completely irrelevant, and these algorithmic descriptions are the most tractable given our high-level theoretical purposes. The computationalist *conjecture* is that cognitive systems will be like computers in this regard; high-level cognitive processes can, and indeed can *only* be tractably described in computational terms which ignore fine-grained temporal issues. Note, however, that this response concedes that computational models are inherently incapable of being fully adequate to the nature of the cognitive processes themselves, since these processes always do unfold in real time. Further, this response concedes that *if there were* a tractable dynamical model of some cognitive process, it would be inherently superior, since it describes aspects of the processes which are out of reach of the computational model. Finally, computationalists have not as yet done enough to convince us that the only tractable models of these high-level processes will be computational ones. Dynamicists, at least, are still working on the assumption that it *will* someday be possible to produce fully adequate models of cognitive processes.

Computationalists sometimes point out that dynamical models of cognitive processes are themselves typically "run" or simulated on digital computers. Does this not establish that computational models are not inherently limited in the way these arguments seem to suggest? Our answer, of course, is no, and the reason is simple: a computational simulation of a dynamical model of some cognitive process is not itself a model of that cognitive process in anything like the manner of standard computational models in cognitive science. Thus, the cognitive system is not being hypothesized to pass through a sequence of symbol structures of the kind that evolve in the computational simulation, any more than a weather pattern is thought to pass through a sequence of discrete symbolic states just because we can simulate a dynamical model of the weather. Rather, all the computational simulation delivers is a sequence of symbolic *descriptions* of points in the dynamical model (and thereby, indirectly, of states of the cognitive system). What we have in such situations is a dynamical model plus an atemporal computational approximation to it.[12]

Continuity in State

Natural cognitive systems sometimes change state in continuous ways; sometimes, on the other hand, they change state in ways that can appear discrete. Dynamics provides a framework within which continuity *and* discreteness can be accounted for, even within the same model. The computational approach, by contrast, can only model a system as changing state from one discrete state to another. Consequently, the dynamical approach is inherently more flexible—and hence more powerful—than the computational approach.

This argument must be carefully distinguished from the previous one. There, the focus was continuity in *time*; the claim was that models must be able to specify the state of the system at every point in time. Here, the focus is continuity in *state*; the claim is that models must be capable of describing change from one state to another arbitrarily close to it, *as well as* sudden change from one state to another discretely distinct from it.

Standard computational systems only change from one discrete state to another.[13] Think again of a Turing machine. Its possible (total) states are configurations of symbols on the tape, the condition of the head, and the position of the head. Every state transition is a matter of adding or deleting a symbol, changing the head condition, and changing its position. The possibilities, however, are all discrete; the system always jumps directly from one state to another without passing through any in-between. There simply *are no* states in between; they are just not defined for the system. The situation is like scoring points in basketball: the ball either goes through the hoop or it doesn't. In basketball, you can't have fractions of points.

When a computational system is used as a model for a natural cognitive process, the natural cognitive system is hypothesized to go through the same state transitions as the model. So a computational model can only attribute discrete states, and discrete state transitions, to the cognitive system.

Now, quite often, state transitions in natural cognitive systems can be thought of as discrete. For example, in trying to understand how people carry out long division in their heads, the internal processes can be thought of as passing through a number of discrete states corresponding to stages in carrying out the division. However, 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. For example, most real problems of sensorimotor coordination deal with a world in which objects and events can come in virtually any shape, size, position, orientation, and motion. A system which can flexibly deal with such a world must be able to occupy states that are equally rich and subtly distinct. Similarly, everyday words as simple as *truck* seem to know no limit in the fineness of contextual shading they can take on. Any system that can understand *Billy drove the truck* must be able to accommodate this spectrum of senses. Only a system that can occupy a continuum of states with respect to word meanings stands a real chance of success.

Many dynamical systems, in the core sense that we have adopted in this chapter, change in continuous phase spaces, and so the dynamical approach is inherently well-suited to describing how cognitive systems might change in continuous ways (see, e.g., Port, Cummins, and McAuley, this volume, chapter 12). However,—and this is the key point—it can also describe discrete transitions in a number of ways. The dynamical approach is therefore more flexible—and hence, again, more powerful—than the computational approach, which can only attribute discrete states to a system.

The dynamical approach can accommodate discrete state transitions in two ways. First, the concepts and tools of dynamics can be used to describe the behavior of systems with only discrete states. A dynamical model of an ecosystem, for example, assumes that its populations always come in discrete amounts; you can have 10 or 11 rabbits, but not 10.5 rabbits. However, perhaps the most interesting respect in which dynamics can handle discreteness is in being able to describe how a continuous system can undergo changes that look discrete from a distance. This is more interesting because cognitive systems appear to be thoroughly pervaded by *both* continuity and discreteness; the ideal model would be one which could account for both together. One kind of discrete change in a continuous system is a *catastrophe*: a sudden, dramatic change in the state of a system when a small change in the parameters of the equations defining the system lead to a qualitative change—a bifurcation—in the "dynamics" or structure of forces operating in that system (Zeeman, 1977; see also Petitot, chapter 9).[14] Thus, high-level, apparently discrete changes of state can be accounted for within a dynamical framework in which continuity and discreteness coexist; indeed, the former is the precondition and explanation for the emergence of the latter.

Multiple Simultaneous Interactions

Consider again the process of returning a serve in tennis. The ball is approaching; you are perceiving its approach, are aware of the other player's movements, are considering the best strategy for the return, and are shifting into position to play the stroke. *All this is happening at the same time.* As you move into place, your perspective on the approaching ball is changing, and hence so is activity on your retina and in your visual system. It is your evolving sense of how to play the point that is affecting your movement. The path of the approaching ball affects which strategy would be best and hence how you move. *Everything is simultaneously affecting everything else.*

Consider natural cognitive systems from another direction entirely. Neurons are complex systems with hundreds, perhaps thousands of synaptic connections. There is some kind of activity in every one of these, all the time. From all this activity, the cell body manages to put together a firing rate. Each cell forms part of a network of neurons, all of which are active (to a greater or lesser degree) all the time, and the activity in each is directly affecting hundreds, perhaps thousands of others, and indirectly affecting countless more. The networks form into maps, the maps into systems, and systems into the central nervous system (CNS), but at every level we have the same principle, that there is constant activity in all components at once, and components are simultaneously affecting one another. No part of the nervous system is ever completely inactive. As neurophysiologist Karl Lashley (1960) put it, "Every bit of evidence available indicates a dynamic, constantly active system, or, rather, a composite of many interacting systems ..." (p. 526).

Clearly, any fully adequate approach to the study of cognitive systems must be one that can handle multiple, simultaneous interactive activity. Yet doing this is the essence of dynamics. Dynamical systems *are* just the simultaneous, mutually influencing activity of multiple parts or aspects. The dynamical approach is therefore inherently well-suited to describe cognitive systems.

A classic example of a dynamical model in this sense is McClelland and Rumelhart's "interactive activation network" (McClelland and Rumelhart, 1981). This model was designed to account for how a letter embedded in the context of a five-letter word of English could be recognized faster than the same letter embedded within a nonword string of letters and even better than the single letter presented by itself. This "word superiority effect" suggested that somehow the whole word was being recognized at the same time as the individual letters that make up the word. Thus, it implied a mechanism where recognition of the word and the letters takes place simultaneously and in such a way that each process influences the other. McClelland and Rumelhart proposed separate cliques of nodes in their network that mutually influence one another by means of coupled difference equations. The output activation of some nodes served as an excitatory or inhibitory input to certain other nodes. This model turned out to capture the word superiority effect and a number of other related effects as well.

Almost all computational approaches attempt to superimpose on this multiple, simultaneous, interactive behavior a sequential, step-by-step structure. They thereby appear to assume that nothing of interest is going on in any component other than the one responsible for carrying out the next stage in the algorithm. It is true, as computationalists will point out, that a computational model can—in principle—run in parallel, though it is devilishly difficult to write such a code. The "blackboard model" of the Hearsay-II speech recognition system (Erman, Hayes-Roth, Lesser, et al. 1980) represents one attempt at approaching parallelism by working within the constraints of serial computationalism. The "blackboard," however, was just a huge, static data structure on which various independent analysis modules might asynchronously post messages, thereby making partial analyses of each module available for other modules to interpret. This is a step in the right direction, but it is a far cry from simultaneous interactive activation. Each module in Hearsay-II can do no more than say "Here is what I have found so far, as stated in terms of my own vocabulary," rather than "Here is exactly how your activity should change on the basis of what has happened in my part of the system,"—the kind of interaction that components governed by coupled equations have with one another. Other methods of parallelism more sophisticated than this may certainly be postulated in principle, but apparently await further technological developments.

Multiple Time Scales

Cognitive processes always take place at many time scales. Changes in the state of neurons can take just a few milliseconds, visual or auditory recognition half a second or less, coordinated movement a few seconds, conversation and story understanding minutes or even hours, and the emergence of sophisticated capacities can take months and years. Further, these time scales are interrelated; processes at one time scale affect processes at another. For example, Esther Thelen (see chapter 3) has shown how actually engaging in coordinated movement promotes the development of coordination, and yet development itself shapes the movements that are possible; it is in this interactive process, moreover, that we find the emergence of concepts such as *space* and *force*. At finer scales, what we see (at the hundreds-of-milliseconds time scale) affects how we move (at the seconds scale) and vice versa.

The dynamical approach provides ways of handling this variety and interdependence of time scales. For example, the equations governing a dynamical system typically include two kinds of variables: state variables and parameters. The way the system changes state depends on both, but only the state variables take on new values; the parameters are standardly fixed. However, it is possible to think of the parameters as not fixed but rather changing as well, though over a considerably longer time scale than the state variables. Thus we can have a single system with both a "fast" dynamics of state variables on a short time scale and a "slow" dynamics of parameters on a long time scale, such that the slow dynamics helps shape the fast dynamics. It is even possible to link the equations such that the fast dynamics shapes the slow dynamics; in such a case, we have true interdependence of time scales.

Note that it is other features of the dynamical approach, such as continuity in space and time, and multiple simultaneous interactive aspects, which make possible its account of the interdependence of time scales. The computational approach, by contrast, has no natural methods of handling this pervasive structural feature of natural cognitive systems.

Self-Organization and the Emergence of Structure

Cognitive systems are highly structured, in both their behavior and their internal spatial and temporal organization. One kind of challenge for cognitive science is to *describe* that structure. Another kind of challenge is to explain *how it got to be there*. Since the computational framework takes inspiration from the organization of formal systems like logic and mathematics, the traditional framework characteristically tackles only the problem of describing the structure that exists. Models in this framework typically postulate some initial set of a priori structures from which more complex structures may be derived by application of rules. The question of *emergence*—of where the initial elements or structures come from—always remains a problem, usually ignored.

A major advantage of the dynamical approach is that dynamical systems are known to be able to create structure both in space and in time. By structure, we mean something nonrandom in form that endures or recurs in time. Thus an archetypal physical object, such as a chair, is invariant in form over time, while a transient event, like a wave breaking on a beach, may recur with temporal regularity. The words in human languages tend to be constructed out of units of speech sound that are reused in different sequences (e.g., *gnat, tan, ant*, etc.), much like the printed letters with which we write words down. But where do *any* such structures come from if they are not either assumed or somehow fashioned from preexisting primitive parts? This is the question of "morphogenesis," the creation of forms. It has counterparts in many branches of science, including cosmology. Why are matter and energy not uniformly distributed in the universe? Study of the physics of relatively homogeneous physical systems, like the ocean, the atmosphere, or a tank of fluid, can begin to provide answers. Some form of energy input is required plus some appropriate dynamical laws. Under these circumstances most systems will tend to generate regular structure of some sort under a broad range of conditions.

The atmosphere exhibits not only its all-too-familiar chaotic properties, but it can also display many kinds of highly regular spatiotemporal structures that can be modeled by the use of differential equations. For example, over the Great Plains in the summer, one sometimes observes long "streets" of parallel clouds with smooth edges like the waves of sand found in shallow water along a beach or in the corduroy ridges on a well-traveled dirt road. How are these parallel ridges created? Not with any form of rake or plow. These patterns all depend on some degree of homogeneity of medium and a consistently applied influx of energy. In other conditions (involving higher energy levels), a fluid medium may, in small regions, structure itself into a highly regular tornado or whirlpool. Although these "objects" are very simple structures, it is still astonishing that any medium so unstructured and so linear in its behavior could somehow constrain itself over vast distances in such a way that regular structures in space and time are produced. The ability of one part of a system to "enslave" other parts, i.e., restrict the degrees of freedom of other, distant parts, is now understood, at least for fairly simple systems (Haken, 1988, 1991; Kelso, Ding, and Schöner, 1992; Thom, 1975).

The demonstration that structure can come into existence without either a specific plan or an independent builder raises the possibility that many structures in physical bodies as well as in cognition might occur without any externally imposed shaping forces. Perhaps cognitive structures, like embryological structures, the weather and many other examples, simply *organize themselves* (Kugler and Turvey, 1987; Thelen and Smith, 1994). Dynamical models are now known to account for many spatial and temporal structures in a very direct way (Madore and Freeman, 1987; Murray, 1989). They enable us to understand how such apparently unlikely structures could come to exist and retain their morphology for some extended period of time. We assume that cognition is a particular structure in space and time—one that supports intelligent interaction with the world. So our job is to discover how such a structure could turn out to be a stable state of the brain in the context of the body and environment. The answer to this question depends both on structure that comes from the genes and on structure that is imposed by the world. No theoretical distinction need be drawn between learning and evolution—they are both, by hypothesis, examples of adaptation toward stable, cognitively effective states of a brain (or an artificial system). The primary difference is that they operate on different time scales.

In both computer science and in cognitive science, the role of adaptation as a source of appropriate structure is under serious development (Forrest, 1991; Holland, 1975; Kauffman, 1993). Most of these methods depend on differential or difference equations for optimization. Thus, a final reason to adopt the dynamical perspective is the possibility of eventually accounting for how the structures that support intelligent behavior could have come about. Detailed models for specific instances of structure creation present many questions and will continue to be developed. But the possibility of such accounts developing from dynamical models can no longer be denied.

Embeddedness

If we follow common usage and use the term *cognitive system* to refer primarily to the internal mechanisms that underwrite sophisticated performance, then cognitive systems are essentially embedded, both in a nervous system and, in a different sense, in a body and environment. Any adequate account of cognitive functioning must be able to describe and explain this embeddedness. Now, the behavior of the nervous system, of bodies (limbs, muscles, bone, blood), and of the immediate physical environment, are all best described in dynamical terms. An advantage of the dynamical conception of cognition is that, by describing cognitive processing in fundamentally similar terms, it minimizes difficulties in accounting for embeddedness.

The embeddedness of cognitive systems has two rather different aspects. The first is the relation of the cognitive system to its neural substrate. The cognitive system somehow *is* the CNS, but what are the architectural and processing principles, and level relationships, that allow us to understand how the one can be the other? The other aspect is the relation of the cognitive system to its essential surrounds—the rest of the body, and the physical environment. How do internal cognitive mechanisms "interact" with the body and the environment?

A computational perspective gives a very different kind of understanding of the behavior of a complex system than a dynamical perspective. Given that the behavior of the nervous system, the body, and the environment are best described in dynamical terms, adopting the computational perspective for internal cognitive mechanisms transforms the *issue* of embedding into a *problem*: how can two kinds of systems, which are described in fundamentally different terms, be related? That is, describing cognition in computational terms automatically creates a theoretical gap between cognitive systems and their surrounds, a gap which must then somehow be bridged.

In the case of the embeddedness of the cognitive system in a nervous system, the problem is to account for how a system that is fundamentally dynamical at one level can simultaneously be a computational system considered at another level. The challenge for the computationalist is to show how such a dynamical system configures itself into a classical computational system. It is a challenge because the two kinds of system are so deeply different. Of course, it is not *impossible* to meet a challenge of this kind; standard digital computers are systems that are continuous dynamical systems at one level and discrete computational systems at another, and we can explain how one realizes the other. However, this provides little reason to believe that a similar cross-level, cross-kind explanation will be feasible in the case of natural cognitive systems, since computers were constructed precisely so that the low-level dynamics would be severely, artificially constrained in exactly the right way. Finding the components of a computational cognitive architecture in the actual dynamical neural hardware of real brains is a challenge of an altogether different order. It is a challenge that computationalists have not even begun to meet.

The embeddedness of the cognitive system within a body and an environment is equally a problem for the computational approach. Again, the problem arises because we are trying to describe the relationship between systems described in fundamentally different terms. The crux of the problem here is time. Most of what organisms deal with happens essentially in time. Most of the critical features of the environment which must be perceived—including events of "high-level" cognitive significance, such as linguistic communication—unfold over time, and so produce changes in the body over time. In action, the movement of the body, and its effects on the environment, happen in time. This poses a real problem for models of cognitive processes which are, in a deep way, atemporal. For the most part, computational approaches have dealt with this problem by simply avoiding it. They have assumed that cognition constitutes an autonomous domain that can be studied entirely independently of embeddedness. The problem of how an atemporal cognitive system interacts with a temporal world is shunted off to supposedly noncognitive transduction systems (i.e., somebody else's problem). When computationalists do face up to problems of embeddedness, the interaction of the cognitive system with the body and world is usually handled in ad hoc, biologically implausible ways. Thus inputs are immediately "detemporalized" by transformation into static structures, as when speech signals are transcribed into a spatial buffer. Outputs are handled by periodic intervention in the environment, with the hope that these interventions will keep nudging things in the right direction. Both methods require the addition to the model of some independent timing device or clock, yet natural cognitive systems don't have clocks in anything like the required sense (Glass and Mackey, 1988; Winfree, 1980). The diurnal clocks observed in many animals, including humans, do not help address the problem of rapid regular sampling that would appear to be required to recognize speech (or a bird song or any other distinctive pattern that is complex in time) using a buffered representation in which time is translated into a labeled spatial axis.

The dynamical approach to cognition handles the embeddedness problem by refusing to create it. The same basic mathematical and conceptual tools are used to describe cognitive processes on the one hand and the nervous system and the body and environment on the other. Though accounting for the embeddedness of cognitive systems is still by no means trivial, at least the dynamical approach to cognition does not face the problem of attempting to overcome the differences between two very different general frameworks. Thus the dynamics of central cognitive processes are nothing more than aggregate dynamics of low-level neural processes, redescribed in higher-level, lower-dimensional terms (see Relation to Neural Processes, below). Dynamical systems theory provides a framework for understanding these level relationships and the emergence of macroscopic order and complexity from microscopic behavior. Similarly, a dynamical account of cognitive processes is directly compatible with dynamical descriptions of the body and the environment, since the dynamical account never steps outside time in the first place. It describes cognitive processes as essentially unfolding over time, and can therefore describe them as occurring in the very same time frame as the movement of the body itself and physical events that occur in the environment.

That cognitive processes must, for this general reason, ultimately be understood dynamically can be appreciated by observing what happens when researchers attempt to build serious models at the interface between internal cognitive mechanisms and the body and environment. Thus Port et al. (see chapter 12) aim to describe how it is possible to handle auditory patterns, with all their complexities of sequence, rhythm, and rate, without biologically implausible artificialities such as static input buffers or a rapid time-sampling system. They find that the inner, cognitive processes themselves must unfold over time with the auditory sequence, and that their qualitative properties (like invariance of perception despite change in rate of presentation) are best described in dynamical terms. In other words, attempting to describe how a cognitive system might perceive its essentially temporal environment drives dynamical conceptualizations inward, into the cognitive system itself. Similarly, researchers interested in the production of speech (see Saltzman, chapter 6; Browman and Goldstein, chapter 7) find that to understand the control of muscle, jaw, etc., we need models of cognitive mechanisms underlying motor control that unfold dynamically in time. That is, attempts to describe how a cognitive system might control essentially temporal bodily movements also drives dynamics inward into the cognitive system. In short, whenever confronted with the problem of explaining how a natural cognitive system might interact with another system that is essentially temporal, one finds that the relevant aspect of the cognitive system itself must be given a dynamical account. It then becomes a problem how this dynamical component of the cognitive system interacts with even more "central" processes. The situation repeats itself, and dynamics is driven further inward. The natural outcome of this progression is a picture of cognitive processing in its entirety, from peripheral input systems to peripheral output systems and everything in between, as all unfolding dynamically in real time: *mind as motion.*

1.3 RELATION TO OTHER APPROACHES

A careful study of the relation of the dynamical conception of cognition to the various other research enterprises in cognitive science would require a book of its own. Here we just make some brief comments on the relation of the dynamical approach to what are currently the two most prominent alternative approaches, mainstream computationalism and connectionism. In addition, we discuss how the dynamical approach relates to the modeling of neural processes and to chaos theory.

Relation to the Computational Approach

Much has already been said about the relation between the computational and dynamical approaches. In this section we add some clarifying remarks on the nature of the empirical competition between the two approaches.

Earlier we characterized cognition in the broadest possible terms as all the processes that are causally implicated in our sophisticated behaviors. Now, it has always been the computationalist position that *some* of these processes are computational in nature and many others are not. For these other processes, traditional dynamical modes of explanation would presumably be quite appropriate. For example, our engaging in an ordinary conversation depends not only on thought processes which enable us to decide what to say next but also on correct movements of lips, tongue, and jaw. Only the former processes would be a matter of internal symbol manipulation; the muscular movements would be dynamical processes best described by differential equations of some sort. In other words, computationalists have always been ready to accept a form of *peaceful coexistence* with alternative forms of explanation targeted at a different selection of the processes underlying sophisticated performance. As we mentioned in section 1.2, the computationalist position is that the processes that must be computational in nature are distinguished by their dependence on "knowledge"; this knowledge must be represented somehow, and the best candidate is symbolically; hence the processes must be computational (symbol manipulation). In fact, from this perspective these knowledge-dependent, symbolic processes are the only *genuinely* cognitive ones; all other processes are peripheral, or implementational, or otherwise ancillary to real cognition.

Now, it has never been entirely clear exactly where the boundary between the two domains actually lies. The conflict between the computational and dynamical approaches can thus be seen as a kind of boundary dispute. The most extreme form of the computationalist hypothesis places the boundary in such a way as to include *all* processes underlying our sophisticated behaviors in the computational domain. Probably nobody has ever maintained such a position, but during the heyday of AI and computational cognitive science in the 1960s and 1970s many more processes were thought to have computational explanations than anyone now supposes. Similarly, the dynamical hypothesis draws the boundary to include *all* processes within the dynamical domain. According to this ambitious doctrine the domain of the computational approach is empty, and dynamical accounts will *eliminate* their computational competitors across all aspects of cognition. It remains to be seen to what extent this is true, but dynamicists in cognitive science are busily attempting to extend the boundary as far as possible, tackling phenomena that were previously assumed to lie squarely within the computational purview.

There is another sense in which computationalists have always been prepared to concede that cognitive systems are dynamical systems. They have accepted that all cognitive processes, including those centrally located in the computational domain, are *implemented* as dynamical processes at a lower level. The situation is exactly analogous to that of a digital desktop computer. The best high-level descriptions of these physical systems are cast in terms of the algorithmic manipulation of symbols. Now, each such manipulation is simply a dynamical process at the level of the electrical circuitry, and there is a sense in which the whole computer is a massively complex dynamical system that is amenable (in principle at least) to a dynamical description. However, any such description would be hopelessly intractable, and would fail to shed any light on the operation of the system *as computing*. Likewise, human thought processes are based ultimately on the firing of neurons and myriad other low-level processes that are best modeled in dynamical terms; nevertheless, the computationalist claims that only high-level computational models will provide tractable, revealing descriptions at the level at which these processes can be seen as *cognitive* performances.

It may even turn out to be the case that there is a high-level computational account of some cognitive phenomenon, *and* a lower-level dynamical account that is *also* theoretically tractable and illuminating. If they are both targeted on essentially the same phenomenon, and there is some precise, systematic mapping between their states and processes, then the computational account would not be eliminated but simply implemented. A relationship of this kind has been recently been advocated for certain psycholinguistic phenomena by Smolensky, Legendre, and Miyata (1992). An alternative possibility is that a high-level, computational description of some phenomenon turns out to be an *approximation*, framed in discrete, sequential, symbol-manipulating terms, of a process whose most powerful and accurate description is in dynamical terms. In such a case only certain of the states and processes in the computational model would stand in a kind of rough correspondence with features of the dynamical model.

Relation to Connectionism

For the purposes of this discussion, we take connectionism to be that rather broad and diverse research program which investigates cognitive processes using artificial neural network models. Defined this way, connectionism is perfectly compatible with the dynamical approach. Indeed, neural networks, which are themselves typically continuous nonlinear dynamical systems, constitute an excellent medium for dynamical modeling.

Thus the two approaches overlap, but only partially. On the one hand, despite the fact that all connectionist networks are dynamical systems, many connectionists have not been utilizing dynamical concepts and tools to any significant degree. At one extreme, connectionists have used their networks to directly implement computational architectures (e.g., Touretzky, 1990). More commonly, they have molded their networks to conform to a broadly computational outlook. In standard feedforward backpropagation networks, for example, processing is seen as the sequential transformation, from one layer to the next, of static representations. Such networks are little more than sophisticated devices for mapping static inputs into static outputs. No dynamics or temporal considerations are deployed in understanding the behavior of the network or the nature of the cognitive task itself. For example, in the famous NETtalk network (Rosenberg and Sejnowski, 1987) the text to be "pronounced" is sequentially fed in via a spatial input buffer and the output is a phonemic specification; all the network does is sequentially transform static input representations into static output representations. To the extent that the difficult temporal problems of speech production are solved at all, these solutions are entirely external to the network. Research of this kind is really more computational than dynamical in basic orientation.

On the other hand, many dynamicists are not connectionists. This is obvious enough on the surface; their intellectual background, focus, and methods are very different (see, e.g., Turvey and Carello, chapter 13; Reidbord and Redington, chapter 17). But what, more precisely, is it that distinguishes the two kinds of dynamicist? If we compare the various contributions to this book, some features of a distinctively connectionist approach emerge. Most obviously, connectionists deploy network dynamical models; they can thus immediately be contrasted with dynamicists whose main contribution is dynamical description (see discussion in section 1.1). Even among those that offer formal dynamical models, there are contrasts between connectionists and others, though the distinction is more one of degree and emphasis.

One kind of contrast is in the nature of the formal model deployed. Connectionists standardly operate with relatively high-dimensional systems that can be broken down into component systems, each of which is just a parametric variation on a common theme (i.e., the artificial neural units). Thus, for example, the connectionist systems used by Randy Beer in his studies of simple autonomous agents are defined by the following general differential equation:

$$\tau_i \dot{y}_i = -y_i + \sum_{j=1}^{N} w_{ji}\sigma(y_j - \theta_j) + I_i(t) \qquad i = 1, 2, \ldots, N$$

In this equation each y_i designates the activation level of i-th of the N individual neural units, and w_{ji} the weight which connects the i-th unit to the j-th unit.[15] This equation is thus really a schema, and if we were to write all the equations out fully, we would have one each for $\dot{y}_1$, $\dot{y}_2$, etc. All these equations take the same form, which is to say that each of the component subsystems (the neural units) are just variations on a common type.

Now, the models deployed by nonconnectionist dynamicists typically cannot be broken down in this way; they are not made up of individual subsystems that have essentially the same dynamical form. For example, the model system deployed by Turvey and Carello (chapter 13) to describe coordination patterns among human oscillating limbs

$$\dot{\phi} = \Delta\omega - a\sin(\phi) - 2y\sin(2\phi) + \sqrt{Q}\xi t$$

has only one state variable (ϕ, the phase difference between the limbs). (See Norton [chapter 2] for plenty of other examples of dynamical systems—including multivariable systems—that cannot be broken down in his way.)

Another kind of contrast is the connectionist tendency to focus on learning and adaptation rather than on mathematical proofs to demonstrate critical properties. Much effort in connectionist modeling is devoted to finding ways to modify parameter settings (e.g., the connection weights) for networks of various architectures so as to exhibit a certain desired behavior, using techniques like backpropagation and genetic algorithms. Nonconnectionists, by contrast, rely on equations using many fewer parameters, with their parameter settings often determined by hand, and typically concentrate proportionately more attention on the fine detail of the dynamics of the resulting system.

In section 1.1 we claimed that connectionism should not be thought of as constituting an alternative to the computational research paradigm in cognitive science. The reason is that there is a much deeper fault line running between the computational approach and the dynamical approach. In our opinion, connectionists have often been attempting, unwittingly and unsuccessfully, to straddle this line: to use dynamical machinery to implement ideas about the nature of cognitive processes which owe more to computationalism. From the perspective of a genuinely dynamical conception of cognition, classic PDP-style connectionism (as contained in, for example, the well-known volumes Rumelhart and McClelland, 1986, and McClelland and Rumelhart, 1986) is little more than an ill-fated attempt to find a halfway house between the two worldviews. This diagnosis is borne out by recent developments. Since its heyday in the mid- to late-1980s, this style of connectionist work has been gradually disappearing, either collapsing back in the computational direction (hybrid networks, and straightforward implementations of computational mechanisms), or becoming increasingly dynamic (e.g., the shift to recurrent networks analyzed with dynamical systems techniques). Connectionist researchers who take the latter path are, of course, welcome participants in the dynamical approach.

Relation to Neural Processes

All cognitive scientists agree that cognition depends critically on neural processes; indeed, it is customary to simply *identify* internal cognitive processing with the activity of the CNS. Neuroscientists are making rapid progress investigating these neural processes. Moreover, the predominant mathematical framework among neuroscientists for the detailed description of neural processes is dynamics, at levels ranging from subcellular chemical transactions to the activity of single neurons and the behavior of whole neural assemblies. The CNS can therefore be considered a single dynamical system with a vast number of state variables. This makes it tempting to suggest that dynamical theories of cognition must be high-level accounts of the very same phenomena that neuroscientists study in fine detail.

This would only be partially true, however. Not all dynamicists in cognitive science are aiming to describe internal neural processes, even at a high level. A central element of the dynamical perspective (see The Nature of Cognitive Systems, above) is that cognitive processes span the nervous system, the body, and the environment; hence cognition cannot be thought of as wholly contained *within* the nervous system. Thus, in modeling cognition, dynamicists select aspects from a spectrum ranging from purely environmental processes (e.g., Bingham, chapter 14) at one extreme to purely intracranial processes (e.g., Petitot, chapter 9) at the other; in between are bodily movements (e.g., Saltzman, chapter 6) and processes which straddle the division between the intracranial and the body or environment (e.g., Turvey and Carello, chapter 13). To select some local aspect of the total cognitive system on which to focus is not to deny the importance or interest of other aspects; choices about which aspect to study are made on the basis of factors such as background, available tools, and hunches about where the most real progress is likely to be made.

Clearly, the idea that the dynamical approach to cognition is just the high-level study of the same processes studied by the neuroscientists is applicable only to those dynamicists whose focus is on processes that are completely or largely within the CNS. Other dynamicists are equally studying *cognition*, but by focusing on other aspects of the large system in which cognitive performance is realized.

What is involved in studying processes *at a higher level*? This simple phrase covers a number of different shifts in focus. Most obviously, dynamical cognitive scientists are attempting to describe systems and behaviors that are *aggregates* of vast numbers of systems and behaviors as described at the neural level. Whereas the neuroscientist may be attempting to describe the dynamics of a single neuron, the dynamicist is interested in the dynamics of whole subsystems of the nervous system, comprised of millions, perhaps billions of neurons. Second, the dynamicist obviously does not study this aggregate system by means of a mathematical model with billions of dimensions. Rather, the aim is to provide a *low-dimensional* model that provides a scientifically tractable description of the same qualitative dynamics as is exhibited by the high-dimensional system. Thus, studying systems at a higher level corresponds to studying them in terms of lower-dimensional mathematical models. Third, dynamical cognitive scientists often attempt to describe the neural processes at a larger *time scale* (see Multiple Time Scales, above). The cognitive time scale is typically assumed to lie between roughly a fifth of a second (the duration of an eyeblink) on up to hours and years. It happens to be approximately the range of time scales over which people have awareness of some of their own states and about which they can talk in natural languages. Neuroscientists, by contrast, typically study processes that occur on a scale of fractions of a second.

Relation to Chaos Theory

Chaos theory is a branch of dynamical systems theory concerned with systems that exhibit chaotic behavior, which for current purposes can be loosely identified with sensitivity to initial conditions (see Norton, chapter 2, for further discussion). Sometimes, especially in popular discussions, the term *chaos theory* is even used to refer to dynamical systems theory in general, though this blurs important distinctions. Chaos theory has been one of the most rapidly developing branches of nonlinear dynamical systems theory, and developments in both pure mathematics and computer simulation have revealed the chaotic nature of a wide variety of physical systems. Chaos theory has even come to provide inspiration and metaphors for many outside the mathematical sciences. It is therefore natural to ask what connection there might be between chaos theory and the dynamical approach to cognition.

The answer is simply that there is *no* essential connection between the two. Rather, chaos theory is just one more conceptual resource offered by dynamical systems theory, a resource that might be usefully applied in the study of cognition, but only if warranted by the data. None of the contributors to this volume have deployed chaos theory in any substantial sense. In this early stage in the development of the dynamical approach, researchers are still exploring how to apply simpler, more manageable models and concepts. The very features that make a system chaotic constitute obvious difficulties for anyone wanting to use that system as a model of cognitive processes.

On the other hand, there are reasons to believe that chaos theory will play some role in a fully developed dynamical account of cognition. Generically, the kinds of systems that dynamicists tend to deploy in modeling cognitive processes (typically continuous and nonlinear) are the home of chaotic processes. Not surprisingly, certain classes of neural networks have been mathematically demonstrated to exhibit chaotic behavior. Of more interest, perhaps, chaos has been empirically observed in brain processes (Basar, 1990; Basar and Bullock, 1989). In one well-known research program, chaotic behavior has been an integral part of a model of the neural processes underlying olfaction (Skarda and Freeman, 1987) (though here the role of chaos was to provide a kind of optimal background or "ready" state rather than the processes of scent recognition themselves). There have been fascinating initial explorations of the idea that highly distinctive kinds of complexity in cognitive performance, such as the productivity of linguistic capacities, might be grounded in chaotic or near-chaos behavior (see, e.g., Pollack, chapter 10). Accounting for such indications of chaos as already exist, and the further uncovering of any role that chaotic notions might play in the heart of cognitive processes, are clearly significant open challenges for the dynamical approach.

1.4 A HISTORICAL SKETCH

The origins of the contemporary dynamical approach to cognition can be traced at least as far back as the 1940s and 1950s, and in particular to that extraordinary flux of ideas loosely gathered around what came to be known as *cybernetics* (Wiener, 1948). At that time the new disciplines of computation theory and information theory were being combined with elements of electrical engineering, control theory, logic, neural network theory, and neurophysiology to open up whole new ways of thinking about systems that can behave in adaptive, purposeful, or other mindlike ways (McCulloch, 1965; Shannon and Weaver, 1949; von Neumann, 1958). There was a pervasive sense at the time that somewhere in this tangled maze of ideas was the path to a rigorous new scientific understanding of both biological and mental phenomena. The problem for those wanting to understand cognition was to identify this path, and follow it beyond toy examples to a deep understanding of natural cognitive systems. What were the really crucial theoretical resources, and how might they be forged into a paradigm for the study of cognition?

Dynamics was an important resource in this period. It was the basis of control theory and the study of feedback mechanisms, and was critical to the theory of analog computation. It figured centrally in neuroscience and in the study of neural networks. The idea that dynamics might form the general framework for a unified science of cognition was the basis of one of the most important books of the period, Ashby's *Design for a Brain* (Ashby, 1952). Interestingly, it was so obvious to Ashby that cognitive systems should be studied from a dynamical perspective that he hardly even bothered to explicitly assert it. Unfortunately, the book was mainly foundational and programmatic; it was short on explicit demonstrations of the utility of this framework in psychological modeling or AI.

During this period two other important strands from the web of cybernetic ideas were under intensive development. One was the theory of neural networks and its use in constructing "brain models" (abstract models of how neural mechanisms might exhibit cognitive functions) (Rosenblatt, 1962). The other was the theory of symbolic computation as manifested in the creation and dominance of LISP as a programming language for AI and for models of psychological processes. Potted histories of the subsequent relationship between these two approaches have become part of the folklore of cognitive science, and the details have been traced in other places (e.g., Dreyfus, 1992). For current purposes, it suffices to say that, although they were initially seen as natural partners, and although research of both types was sometimes even conducted by the same researchers, beginning in the late 1950s, neural network research and computationalism separated into distinct and competing research paradigms. The computational approach scored some early successes and managed to grab the spotlight, appropriating to itself the vivid phrase "artificial intelligence" and the lion's share of research funding. In this way computer science came to provide the theoretical core of mainstream cognitive science for a generation. Neural network research, nevertheless, did continue throughout this period, and much of it was strongly dynamical in flavor. Of particular note here is the work of Stephen Grossberg and colleagues, in which dynamical ideas were being applied in a neural network context to a wide range of aspects of cognitive functioning (see Grossberg, chapter 15).

By the early 1980s mainstream computational AI and cognitive science had begun to lose steam, and a new generation of cognitive scientists began casting around for other frameworks within which to tackle some of the issues that caused problems for the computational approach. As is well known, this is when neural network research burgeoned in popularity and came to be known as *connectionism* (Hinton and Anderson, 1981; Rumelhart and McClelland, 1986; McClelland and Rumelhart, 1986; Quinlan, 1991). Since connectionist networks are dynamical systems, it was inevitable that dynamical tools would become important for understanding their behavior and thereby the nature of cognitive functions. The recent rapid emergence of the dynamical approach is thus due, in large measure, to this reemergence of connectionism and its development in a dynamical direction.

Apart from cybernetics and neural network research, at least three other research programs deserve mention as antecedents to the contemporary dynamical approach. One is derived from the physical sciences via biology, another from pure mathematics, and the third from experimental psychology. The first began with the question: Can the basic principles of description and explanation applied with such success in the physical sciences to simple closed systems be somehow extended or developed to yield an understanding of complex, open systems? In particular, can general mathematical laws be deployed in understanding the kinds of behaviors exhibited by *biological* systems? One natural target was the biological phenomenon of coordinated movement, since it involves regular, mathematically describable motion. Yet the study of coordinated movement cannot avoid eventually invoking notions such as intention, information, and perception, and so must overlap with psychology. At this nexus arose a distinctive program of research into human motor and perceptual skills which relied on resources proposed by physicists and mathematicians such as Pattee, Prigogine, Rosen, and Haken, and was inspired by Bernstein's insights into motor control (Bernstein, 1967). This program is exemplified in the work of Turvey, Kugler, and Kelso (Kelso and Kay, 1987; Kugler and Turvey, 1987).

Dynamics is, in the first instance, a branch of mathematics. Applications of dynamics in various areas of science have often flowed directly from developments in pure mathematics. A particularly dramatic example of this phenomenon has been applications derived from the development, principally by René Thom, of catastrophe theory. This theory is an extension of dynamics, in combination with topology, to describe situations in which there arise discontinuities, i.e., sudden, dramatic changes in the state of a system.[16] Discontinuities are common in the physical, biological, cognitive, and social domains, and are the basis for the formation of temporal *structures*, and so the development of catastrophe theory led directly to new attempts to describe and explain phenomena that had been beyond the scope of existing mathematical methods. Of particular relevance here is application of catastrophe theory to the investigation of language and cognition. Initial proposals by Thom and Zeeman (Thom, 1975; Thom, 1983; Zeeman, 1977) have been taken up and developed by Wildgen (1982) and Petitot (1985a, b) among others. This work has involved some radical and ambitious rethinking of problems of perception and the nature of language.

A third major source of inspiration for dynamical modeling came from Gibson's work in the psychology of perception (Gibson, 1979). Gibson asserted that it was a mistake to devote too much attention to models of internal mechanisms when the structure of stimulus information remained so poorly understood. Since both the world and our bodies move about, it seemed likely to Gibson that the structuring of stimulus energy (such as light) by dynamical environmental events would play a crucial role in the achievement of successful real-time interaction with the environment. The resulting focus on discovery of the sources of high-level information in the stimulus turns out to dovetail nicely with continuous-time theories of dynamic perception and dynamic action. The inheritors of Gibson's baton have had many successes at specifying the dynamic information that underlies perceptual and motor achievement (e.g., see Turvey and Carello, chapter, 13; Bingham, chapter 14). The work of the ecological psychologists has been a key influence in encouraging researchers to adopt a dynamical perspective in various other areas of cognitive science.

These five lines of research have recently been joined by other dynamics-based investigations into a wide variety of aspects of cognition. What explains this recent surge in activity? Partly, of course, it is the spreading influence of the research programs just described. But another important factor has been the rapid development in mathematics of nonlinear dynamical systems theory in the 1970s and 1980s, providing contemporary scientists with a much more extensive and powerful repertoire of conceptual and analytical tools than were available to Ashby or McCulloch, for example. At a more prosaic level, the continuing exponential growth in computing resources available to scientists has provided cognitive scientists with the computational muscle required to explore complex dynamical systems. In recent years a number of new software packages have made dynamical modeling feasible even for researchers whose primary training is not in mathematics or computer science.

Finally, of course, there is the nature of cognition itself. If the dynamical conception of cognition is largely correct, then a partial explanation of why researchers are, increasingly, applying dynamical tools may lie simply in the fact that cognitive systems are the kind of systems that call out for a dynamical treatment.

ACKNOWLEDGMENTS

This research was supported by a Queen Elizabeth II Research Fellowship from the Australian Research Council to the first author, and the Office of Naval Research grants N0001491-J-1261, N0001493, and N0001492-J-1029 to the second author. Critical feedback from John Haugeland, Esther Thelen, and James Townsend was especially useful in its preparation.

NOTES

1. Technically, a differential equation is any equation involving a function and one or more of its derivatives. For more details on differential equations, and the mass-spring equation in particular, see Norton, chapter 2.

2. The notion of phase, like that of dynamical system itself, differs from one context to another. In some contexts, a phase space is taken to be one in which one of the dimensions is a time derivative such as velocity. In other contexts, phase is taken to refer to position in a periodic pattern, as when we talk of the phase of an oscillating signal. Our notion of phase here is a generalization of this latter sense. Since the rule governing a state-determined system determines a unique sequence of points for any given point, every point in the space can be understood as occupying a position (or "phase") in the total pattern (or "dynamic") fixed by the rule. Our use thus accords with the common description of diagrams that sketch the overall behavior of a dynamical system as *phase portraits* (see, e.g., Abraham and Shaw, 1982).

3. For an example of the use of forms of the logistic equation, as a difference equation, in cognitive modeling, see van Geert, chapter 11.

4. In fact, the total state of a computational system is more than just a configuration of symbols. A Turing machine, for example, has at any time a configuration of symbols on its tape, but it is also in a certain head state, and the head occupies a certain position; these must also be counted as components of the total state of the system.

5. In particular, one could not demonstrate that cognitive systems are dynamical systems merely by showing that any given natural cognitive system is governed by some dynamical rule or other. Certainly, all people and animals obey the laws of classical mechanics; drop any one from a high place, and it will accelerate at a rate determined by the force of gravitational attraction. However, this does not show that *cognitive* systems are dynamical systems; it merely illustrates the fact that heavy objects belong to dynamical systems.

6. A more radical possibility is that dynamical systems can behave in a way that depends on knowledge without actually *representing* that knowledge by means of any particular, identifiable aspect of the system.

7. For a more detailed introduction to dynamics, see Norton, chapter 2.

8. Of course, a range of general and quite powerful arguments have been put forward as demonstrating that cognitive systems must be computational in nature (see, e.g., Fodor, 1975; Newell and Simon, 1976; Pylyshyn, 1984). Dynamicists remain unconvinced by these arguments, but we do not have space here to cover the arguments and the dynamicists' responses to them.

9. Turing machines are a particularly simple kind of computer, consisting of one long tape marked with squares that can contain symbols, and a "head" (a central processing unit) which moves from square to square making changes to the symbols. They are very often used in discussions of foundational issues in cognitive science because they are widely known and, despite their simplicity, can (in principle) perform computations just as complex as any other computational system. For a very accessible introduction to Turing machines, see Haugeland (1985).

10. One *inappropriate* way to extract temporal considerations from a computational model is to rely on the timing of operations that follow from the model's being *implemented* in real physical hardware. This is inappropriate because the particular details of a model's hardware implementation are irrelevant to the nature of the model, and the choice of a particular implementation is theoretically completely arbitrary.

11. Ironically, these kinds of assumptions have often been the basis for attacks on the plausibility of computational models. If you assume that each computational step must take some certain minimum amount of time, it is not difficult to convince yourself that the typical computational model has no hope of completing its operations within a psychologically realistic amount of time.

12. Precisely because discrete models are only an approximation of an underlying continuous one, there are hard limits on how well the continuous function can be modeled. Thus, it is well known to communications engineers that one must have at least two discrete samples for each event of interest in the signal (often called Nyquist's theorem). The cognitive corollary of this is that to model dynamical cognitive events that last on the order of a half-second and longer, one must discretely compute the trajectory at least four times a second. Anything less may result in artifactual characterization of the events. Since the time scale of cognitive events is relatively slow compared to modern computers, this limit on discrete modeling of cognition would not itself serve as a limiting constraint on real-time modeling of human cognitive processes.

13. This is true for computational systems when they are considered at the level at which we understand them as computational. The same object (e.g., a desktop computer) can be seen as undergoing continuous state changes when understood at some different level, e.g., the level of electric circuits.

14. Note that when continuous systems bifurcate there can be *genuinely* discrete changes in the attractor landscape of the system.

15. For a more detailed explanation of this equation, see Beer, chapter 5.

16. Note that this informal notion of *discontinuity* should not be confused with the precise mathematical notion. It is a central feature of catastrophe theory that systems that are continuous in the strict mathematical sense can exhibit discontinuities—dramatic, sharp changes—in the more informal sense.

REFERENCES

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

Basar, E. (Ed.). (1990). *Chaos in brain function.* Berlin: Springer-Verlag.

Basar, E., and Bullock, T. H. (Ed.). (1989). *Brain dynamics: progress and perspectives.* Berlin: Springer-Verlag.

Bernstein, N. A. (1967). *The control and regulation of movement.* London: Pergamon.

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,* 432–459.

Devaney, R. L. (1986). *An introduction to chaotic dynamical systems.* Menlo Park, CA: Benjamin/Cummings.

Dreyfus, H. L. (1992). *What computers still can't do: a critique of artificial reason.* Cambridge, MA: MIT Press.

Elman, J. L. (1991). Distributed representations, simple recurrent networks, and grammatical structure. *Machine Learning, 7,* 195–225.

Erman, L. D., Hayes-Roth, F., Lesser, V. R., et al. (1980). The HEARSAY-II speech understanding system: integrating knowledge to resolve uncertainty. *Computing Surveys, 12,* 213–253.

Fodor, J. A. (1975). *The language of thought.* Cambridge, MA: Harvard University Press.

Forrest, S. (Ed.). (1991). *Emergent computation: self-organizing, collective, and cooperative phenomena in natural and artificial computing networks.* Cambridge, MA: MIT Press.

Gibson, J. J. (1979). *The ecological approach to visual perception.* Boston: Houghton-Mifflin.

Glass, L., and Mackey, M. (1988). *From clocks to chaos: the rhythms of life.* Princeton NJ: Princeton University Press.

Haken, H. (1988). *Information and self-organization: a macroscopic approach to complex systems.* Berlin: Springer-Verlag.

Haken, H. (1991). *Synergetics, computers and cognition.* Berlin: Springer-Verlag.

Haken, H., Kelso, J. A. S., and Bunz, H. (1985). A theoretical model of phase transitions in human hand movements. *Biological Cybernetics, 51,* 347–356.

Haugeland, J. (1985). *Artificial intelligence: the very idea.* Cambridge, MA: MIT Press.

Hinton, G. E., and Anderson, J. A. (Ed.). (1981). *Parallel models of associative memory.* Hillsdale, NJ: Erlbaum.

Holland, J. H. (1975). *Adaptation in natural and artificial systems.* Ann Arbor: University of Michigan Press.

Hopfield, J. (1982). Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the National Academy of Sciences of the United States of America, 79,* 2554–2558.

Kauffman, S. A. (1993). *The origins of order: self-organization and selection in evolution.* New York: Oxford University Press.

Kelso, J. A., and Kay, B. A. (1987). Information and control: a macroscopic analysis of perception-action coupling. In H. Heuer and A. F. Sanders (Eds.), *Perspectives on perception and action.* Hillsdale, NJ: Erlbaum.

Kelso, J. A. S., Ding, M., and Schöner, G. (1992). Dynamic pattern formation: a primer. In J. E. Mittenthal and A. B. Baskin (Eds.), *Principles of organization in organisms.* Reading, MA: Addison-Wesley.

Kugler, P. N., and Turvey, M. T. (1987). *Information, natural law, and the self-assembly of rhythmic movement.* Hillsdale, NJ: Erlbaum.

Lashley, K. S. (1960). The problem of serial order in behavior. In F. A. Beach, D. O. Hebb, C. T. Morgan, et al. (Eds.), *The neuropsychology of Lashley.* New York: McGraw-Hill.

Madore, B. F., and Freeman, W. L. (1987). Self-organizing structures. *American Scientist, 75,* 253–259.

McClelland, J. L., and Rumelhart, D. E. (1981). An interactive-activation model of context effects in letter perception: Part 1, an account of basic findings. *Psychological Review, 88,* 375–407.

McClelland, J. L., and Rumelhart, D. E. (Eds.). (1986). *Parallel distributed processing: explorations in the microstructure of cognition,* Vol. 2: *Psychological and biological models.* Cambridge, MA: MIT Press.

McCulloch, W. S. (1965). *Embodiments of mind.* Cambridge, MA: MIT Press.

Murray, J. D. (1989). *Mathematical biology.* Berlin: Springer-Verlag.

Newell, A., and Simon, H. (1976). Computer science as empirical enquiry: symbols and search. *Communications of the Association for Computing Machinery, 19,* 113–126.

Petitot, J. (1985a). *Les catastrophes de la parole.* Paris: Maloine.

Petitot, J. (1985b). *Morphogenèse du sens.* Paris: Presses Universitaires de France.

Pylyshyn, Z. W. (1984). *Computation and cognition: toward a foundation for cognitive science.* Cambridge, MA: Bradford/MIT Press.

Quinlan, P. (1991). *Connectionism and psychology.* Chicago: University of Chicago Press.

Rosenberg, C. R., and Sejnowski, T. J. (1987). Parallel networks that learn to pronounce English text. *Complex Systems, 1,* 145–168.

Rosenblatt, F. (1962). *Principles of neurodynamics: perceptrons and the theory of brain mechanisms.* New York: Spartan Books.

Rumelhart, D. E., and McClelland, J. L. (Eds.). (1986). *Parallel Distributed Processing: explorations in the microstructure of cognition,* Vol 1: *Foundations.* Cambridge, MA: MIT Press.

Shannon, C. E., and Weaver, W. (1949). *The mathematical theory of communication.* Urbana: University of Illinois Press.

Skarda, C. A., and Freeman, W. J. (1987). How brains make chaos to make sense of the world. *Behavior and Brain Sciences, 10,* 161–195.

Smolensky, P., Legendre, G., and Miyata, Y. (1992). *Principles for an integrated connectionist/symbolic theory of higher cognition.* (No. CU-CS-600-92.) Computer Science Department, University of Colorado, Boulder.

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

Thom, R. (1975). *Structural stability and morphogenesis* (Fowler, D. H., Trans.). Reading, MA: Benjamin.

Thom, R. (1983). *Mathematical models of morphogenesis.* Chichester, England: Ellis Horwood.

Touretzky, D. S. (1990). BoltzCONS: dynamic symbol structures in a connectionist network. *Artificial Intelligence, 46,* 5–46.

Von Neumann, J. (1958). *The computer and the brain.* New Haven: Yale University Press.

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

Wildgen, W. (1982). *Catastrophe theoretic semantics: an elaboration and extension of René Thom's theory.* Amsterdam: Benjamins.

Winfree, A. T. (1980). *The geometry of biological time.* New York: Springer-Verlag.

Zeeman, C. (1977). *Catastrophe theory: selected papers 1972–1977.* Redwood City, CA: Addison-Wesley.