Mind as a Dynamical System

by

Chris Eliasmith

A thesis
presented to the University of Waterloo
in fulfillment of the
thesis requirement for the degree of
Master of Arts
in
Philosophy

Waterloo, Ontario, Canada, 1995

Current address:
Philosophy-Neuroscience-Psychology Program
Department of Philosophy
Washington University in St. Louis
One Brookings Drive
St. Louis, MO 63130-4899
chris@twinearth.wustl.edu
http://artsci.wustl.edu/~celiasmi/

Abstract

Recently, a new approach to modeling cognitive phenomena has been gaining recognition: the dynamical systems approach. Proponents of this theory claim to have identified a new paradigm for the study of cognition which is superior to both symbolicism and connectionism.

Though dynamicism is a recent addition to cognitive science, its supporters have leveled severe criticism of the theoretical commitments of both symbolicism and connectionism. They wish to remove traditional concepts such as computation and representation from the vocabulary of cognitive scientists. Instead, they wish to discuss cognition in dynamical systems terms, using mathematical concepts such as state space, chaos, and attractor.

Through an examination of what is expected of good cognitive theories, and a discussion of the strengths and weaknesses of symbolicism, connectionism and dynamicism, this thesis attempts to situate and evaluate the dynamicist project in the context of current cognitive science. It is concluded that, though dynamicist ideas may play an important normative role in the future development of cognitive models, they are best understood as high-level descriptions of the behavior of connectionist systems.

Acknowledgments

Thanks must first of all go to my supervisor, Paul Thagard, whose guidance has been invaluable in the writing of this thesis. I cannot express the depth of my indebtedness to Paul for exposing me to an area of research which addresses many of my interests so fully.

In addition, I wish to express my gratitude to my office mates Cameron Shelley, and Paul Rusnock for their comments and discussions which have challenged (and entertained) me, and have thus helped to shape this thesis.

Many thanks also go to the readers of this thesis Jim Van Evra and Bill Abbott, whose comments, pointers, and questions have led me to think that this research has a future as well as a past.

I also wish to thank Andrew Meade for taking the time to answer questions which have clarified elements of this thesis.

Dedication

I would like to sincerely dedicate this work to my family: Jerry and Janet Elias, Steven Elias, Kim Pendrith, and, of course, Jennifer Smith.

Table of Contents

List of Figures

Introduction

In the mid 1950s, conceptualization of cognition in terms of information structures was introduced into the psychological literature. Both neural-like processing, and symbolic processing began to be examined as theories explaining the nature of cognition. These approaches to understanding human thinking gave rise to the discipline known as cognitive science. Very recently, however, another approach to modeling the mind has been gaining recognition; the dynamical systems approach. Proponents of this theory have expressed important concerns about the shortcomings of current theoretical commitments and they wish to displace both symbolic and neural-like theories with their own. This thesis is an attempt to determine if this desire is either practically or theoretically justified.

In the late 1950s it became increasingly common to discuss higher human cognitive powers, such as language, reasoning and planning, in terms of purely symbolic representations (see Gardner, 1985; Baars 1986 as in Thagard, 1992). Commitment to the symbol-based view of cognition will be referred to as *symbolicism*. Symbolicists have employed various forms of symbolic representation and computation ranging from collections of if-then rules (referred to as production systems) and predicate calculus to conceptual frames and behavioral schemas. Symbolicists believe that all cognition is performed by manipulation of such symbolic representations. Prime examples of employing this approach to the study of cognition are found in Anderson's set of Adaptive Control of Thought (Act) models which culminated with Act* (pronounced ‘Act-star’) (Anderson, 1983) , and Newell's Soar architecture (Newell, 1990) .

However, the late 1970s saw the resurgence of the examination of cognitive functioning in terms of neural-like, or parallel distributed, processing (PDP). This theory, in its purer form, postulates a distributed representation of concepts across a network of highly interconnected simple nodes. Rather than consisting of atomic, language-like representations, PDP representations do not reside at a particular node but are spread out over an area of the network. The interconnected structure of these processing networks has resulted in this approach commonly being referred to as *connectionism*. Connectionists rely on models of often simplified neurons operating in parallel to generate a desired behavior. Numerous examples of the application of such models can be found in Paul and Pat Churchland's work on cognition (Churchland, 1989; Churchland and Sejnowski, 1992) .

Since the introduction of connectionism in the late 1970s, the two main paradigms of cognitive science have been symbolicism, and connectionism. However, in recent years, a new approach to the study of cognition has attempted to challenge their dominance - namely, *dynamicism*. There have been a series of papers and books (Globus, 1992; Robertson, Cohen et al., 1993; Thelen and Smith, 1994; van Gelder and Port, in press) that have advanced the claim that cognition is not best understood as symbolic manipulation or connectionist processing, but as complex, dynamical interactions of a cognizer with its environment. Dynamicists have critically examined both symbolicism and connectionism and have decided to dismiss these theories of cognition and instead wish to propose a “radical departure from current cognitive theory,” one in which “there are no structures” and “there are no rules” (Thelen and Smith, 1994, p. xix) .

The dynamicist view relies heavily on an area of mathematics referred to as *dynamical systems theory*. The concepts of dynamical systems theory are applied by dynamicists to a description of cognition. Mathematical ideas such as *state space*, *attractor*, *trajectory*, and *deterministic chaos* are used to explain the internal processing which underlies an agent's interactions with the environment. Hence, systems of differential equations are used to represent an agent's cognitive trajectory through a state space. In other words, cognition is explained as a multi-dimensional space of all possible thoughts and behaviors that is traversed by a path of thinking followed by an agent under certain environmental and internal pressures: all of which is captured by sets of differential equations. Van Gelder (in press) elucidates this view in his introduction to the book *Mind as motion: Explorations in the dynamics of cognition*.

Fundamentally, dynamicists believe that the brain is continually changing as it intersects with information from its environment. There are no representations in any normal sense of the word, rather there “is state-space evolution in certain kinds of non-computational dynamical systems” (van Gelder and Port, in press, p. 1) . The temporal nature of cognition does not rely on “clock ticks” or the completion of a particular task, rather it is captured by a continual evolution of interacting system parts which are always reacting to, and interacting with, the environment. Hence, a dynamicist would claim that symbolicist and connectionist models are being incorrectly superimposed on cognitive mechanisms which can only be faithfully described in terms of dynamical systems.

Dynamicists claim that through their critique of the current state of cognitive science, they are challenging a conceptual framework which has been applied to the problem of cognition since the time of Descartes. Rather than a Cartesian distinction between the cognizer and its environment, dynamicists hold that “the human agent as essentially embedded in, and skillfully coping with, a changing world” (van Gelder and Port, in press, p. 39) . Thus, we are not governed by rules, but “global settlements”; the brain does not compute, it “participates” in the environment; we do not think per se, but are self-organizing, adaptive, nonlinear dynamical systems, whose complex subtleties cannot be captured by connectionist or symbolicist models of cognitive processes.

Despite the many, and substantial, disagreements amongst symbolicism, connectionism and dynamicism, as cognitive theories the three approaches also maintain important similarities. Possibly the most fundamental assumption of all three theories is best expressed by Churchland and Sejnowski: “At this stage in the evolution of science, it appears highly probable that psychological processes are in fact processes of the physical brain, not, as Descartes concluded, processes of a nonphysical soul or mind” (Churchland and Sejnowski, 1992, p. 1) . In sharing this tenet, it becomes necessary for each approach to describe the mind as a particular sort of physical process.

The introduction of the digital computer was seen by many researchers as an opportunity to clearly identify what type of physical processes may be going on in the brain. General acceptance of the computer/brain analogy signaled the rise of symbolicism to its position of power: “Critics need to identify the reigning paradigm, and they are pretty clear that it is symbolic cognitive psychology” (Newell, 1990, p. 24; similarly, Harnard, 1992, p. 77; van Gelder and Port, in press, p. 1) . The symbolic view of cognition has inspired construction of “artificially intelligent” agents from the 1960s to the present day.

Both connectionism and dynamicism have shorter histories than symbolicism and are often considered to be less empirically supported or theoretically well-grounded than symbolicism, even by their proponents (van Gelder, 1993, p. 1) . Of course, it is often argued that this deficiency is due to the relative youth of the connectionist and dynamicist approaches. Thus it is necessary to closely examine theoretical distinctions as well as empirical or pragmatic ones.

Of these two approaches, connectionism has had a longer life span. Thus, significantly more connectionist models than dynamicist models have been developed. In many cases, however, the same model is claimed to support both dynamicist and connectionist interests (*e.g.* Skarda and Freeman, 1987) . Thus, a distinction between connectionism and dynamicism is often quite subtle. Commonly, discriminating these approaches relies on the dynamicist belief that connectionist networks are not adequate for reproducing the behavior of living neural nets which rely on “ongoing chemical tuning of the input/output transfer function at the nodes, connection weights, network parameters, and connectivity” (Globus, 1992, p. 299) . So, it seems if a model is complicated enough, the dynamicist would claim it is no longer connectionist, but dynamicist. It is uncertain how connectionists who posit such models (see Churchland and Sejnowski, 1992) would react to an insistence that they are not connectionists. Less controversially, dynamicists have attempted to generalize concepts found in connectionism in order to move cognitive science away from any sort of reliance on connectionist networks as the basis for nonlinear dynamical modeling of cognition (van Gelder and Port, in press) . So, the dynamicists wish to distinguish themselves, not on the basis of the physical processes being modeled, but on the basis of a specific *type* of dynamical system which best describes those processes.

This commitment to describing cognition as a physical process allies all three approaches in their acceptance of the enterprise of cognitive science. As Newell (1990) has noted, cognitive science is not simply interested in designing *any* intelligent agent, but more specifically an agent with *human-like* intelligence. Thus, cognitive science attempts to posit theories of problem solving that explain the observed *human* behavior. It is unanimous among the three approaches to cognitive modeling that their research interests include at least: problem solving, decision making, routine action, memory, learning, skill, perception, motor behavior, language, motivation, emotion, imagining and dreaming (Newell, 1990, p. 15; Churchland and Sejnowski, 1992; van Gelder and Port, in press, p. 1) .

It is also agreed that the more of these areas that are addressed by a single theory, the more powerful that theory is: “ultimately what is wanted is a story, unified from top to bottom - from behavior, through systems to networks to neurons and molecules: a unified science of the mind-brain" (Churchland and Sejnowski, 1992, p. 13) ; or similarly from a symbolicist: "the hallmark of a unified theory is the range of central cognition and its surround that it addresses" (Newell, 1990, p. 31) ; or lastly from the dynamicist: "a picture of cognitive processing in its entirety, from peripheral input systems to peripheral output systems and everything in between" (van Gelder and Port, in press) . In other words, each approach to cognitive science is attempting to unify our understanding of cognition.

The purpose of this thesis is to compare these three theories of mind in an attempt to clearly understand their relations to one another and their respective strengths and weaknesses. Moreover, consideration of dynamicism will be emphasized and I will attempt to determine exactly what dynamicism has to offer the study of cognition. If dynamicism is truly a new paradigm, ready to replace the out-of-date theories of symbolicism and connectionism, as its proponents claim (Thelen and Smith, 1994; van Gelder and Port, in press) , it not only deserves consideration, but all of the resources of cognitive science should be dedicated to its advancement. Such a contention is well worth consideration.

However, before this issue can be addressed, we need to consider what constitutes a good model of mind in order to establish criteria by which to evaluate these three competing approaches. In doing so, it is necessary to elucidate the many uses of the term model and understand their relation to metaphor, analogy and theory especially in the context of cognitive science. The first chapter is dedicated to this examination.

With this discussion in mind, a comparison of dynamicism, connectionism, and symbolicism is detailed in chapter two, allowing an overall framework for understanding the strengths and weaknesses of each of these approaches. The theories will be compared in their ability to provide good accounts of the temporal, architectural, computational, and representational nature of cognition.

Ultimately, the third chapter is dedicated to a discussion of the dynamicist view of mind. Because dynamicism is the most recently proposed theory and is, in a sense, attempting to carve out a place for itself in the field of cognitive science, it warrants a fuller discussion. Specific examples, and general theoretical discussions elucidate the limitations and strengths of the dynamicist approach to cognition, allowing us to determine exactly how and where dynamicism fits into the discipline of cognitive science.

Chapter 1

Theories, Models and Cognition

In this chapter we will survey general distinctions between models, metaphors, analogies and theories, and discuss these differences in the context of cognitive science. Subsequently, we will examine the *application* of such models and theories in cognitive science in an attempt to generate criteria for evaluating dynamicist theory and its resulting models.

1.1 Distinguishing Metaphor, Analogy, Model and Theory

As Churchland has noted, problems in the philosophy of mind have recently found themselves reconstructed as problems in the philosophy of science (Churchland, 1989) . The result, has been an infusion of both philosophy of science and philosophy of mind into the multidisciplinary field of cognitive science. With this shift in attitudes, it becomes necessary to establish a clear understanding of foundational philosophy of science problems including the role of metaphor, analogy, model, and theory in order to develop a convincing framework in which to examine the proposed dynamical systems theory of cognition.

1.1.1 Metaphor and Analogy

Metaphor and analogy both involve a transfer of meaning from one situation to another; this type of attribution is commonly referred to as a *mapping* from a *source* to a *target* (Holyoak and Thagard, 1995; *c.f.* Beardsley, 1972) . For the purposes of this thesis, there is no need to distinguish metaphor from analogy; hence the terms may be used interchangeably. However, because the term *metaphor* is often used in a less rigorous sense (*i.e.* in a literary, aesthetic, or rhetoric context), *analogy* will most often be our term of choice, as we are examining these sorts of relations in a scientific context (as per Hesse, 1972) .

In forming an analogy, one attempts to express similarities between situations in the world, say A and B (Holyoak and Thagard, 1995, p. 5) . If the analogy is good, these similarities will allow for knowledge of the better understood situation, say A, to be transferred effectively to the lesser understood situation, that is B. Thus, an analogy is formed upon the utterance of a simple statement such as “B is like A”. However, such a statement is best regarded as a sort of shorthand which often requires considerable interpretation. In claiming, for example, “The atom is like the solar system,” we must explicitly come to understand the structural and relational mappings between the elements of the solar system and those of the atom as in figure 1.1. Mappings are often complex and it is sometimes difficult to distinguish between valid and invalid mappings. For this reason, Beardsley cautions: “just because of the metaphors complexity of meaning, it is especially susceptible to misunderstanding...and cannot safely be used in inductive or deductive argument “ (Beardsley, 1972, p. 286; *c.f.* Holyoak and Thagard, 1995, p. 208) . In our example, we can see that the concept of encircling for the solar system (*i.e.* orbit) is superficially like, but very unlike, the quantum mechanical probability distribution of electrons around a nucleus in an atom. Equating the orbiting of planets with the motion of electrons is simply wrong.

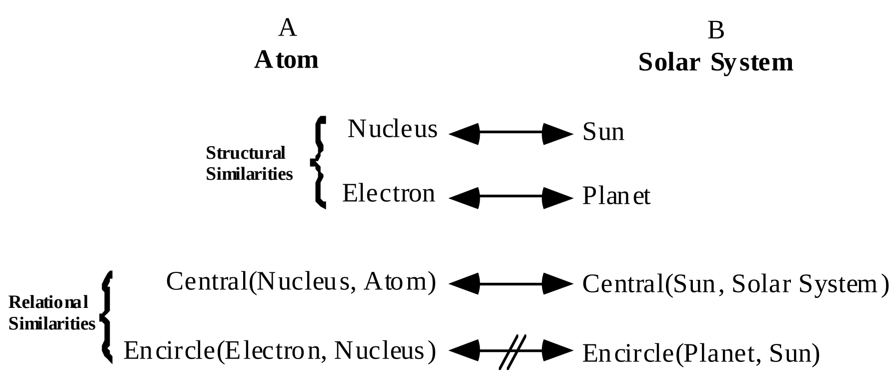

Figure 1.1: Mappings in the atom/solar system analogy.

In other words, analogies are not rigorous. There is ample room for various interpretation or extrapolation from what is provided by the original analogy. The analogy is essentially a limited description of B *in terms of* A. Therefore, the description is not a precise one. There are certain details of B which will not be captured and others which may even be falsely implied.

Nevertheless, analogy plays an extremely important role in the sciences. Scientific analogical thinking has often led researchers to a novel perspective resulting in valuable insights (Holyoak and Thagard, 1995, pp. 186-188) . Similarly, analogy in cognitive science is often central to a particular view of cognition. For the symbolicist, the mind is a specific type of computer; a symbol manipulator. For the connectionists, cognition is still computational, but rather as a property of highly connected networks of nodes which are analogous to neurons in the human brain. For the dynamicist, mentality is like a complex temporal system, such as weather patterns or a Watt governor (van Gelder, 1993) , traveling from attractor to attractor through the terrain of a high-dimensional state space. All three of these analogies have had theories and rigorous models based on them: “Perhaps every science must start with metaphor and end with algebra; and perhaps without the metaphor there would never have been any algebra.” (Black (1962) as quoted in Beardsley, 1972) .

1.1.2 Model and Analogy

Unfortunately, the term *model* has a wide variety of colloquial and scientific uses, resulting in an obfuscation of the term. For example, we have in our conceptual repertoire for the term *model* the likes of fashion models, model airplanes, model behavior, mathematical models and computer models to mention a few. The common relation between all such uses lies in their reference to a resemblance of some kind (*i.e.* an *ideal* clothes wearer, a *real* airplane, *ideal* behavior, a description of *behavior of a real system* and an *implementation* of a mathematically described system, respectively): “The relation between model and thing modeled can be said generally to be a relation of analogy” (Hesse, 1972, p. 355) .

In science, however, a model is no longer simply a resemblance, but rather a precise description of the properties of the system being modeled. The more important properties of the source that are exactly demonstrated by the model, the better the model (see section 1.2.1). So, to differentiate between model and analogy in science, one can determine if the mapping of these important properties is explicit, leaving no room for interpretation; if so, one is dealing with a model. In this sense, a model can be thought of as “a kind of controlled metaphor” (Beardsley, 1972, p. 287) [1]. Thus, where an analogy may consist of

---

[1] In contrast to an explicitly controlled metaphor, Beardsley describes to the use of a *normal* metaphor: "But of course the (normal) metaphorical description, as its implications are pursued, can be checked at each step, and we need not feel committed to all of its implications merely because it has a general appropriateness. So the metaphorical description may least misleadingly, perhaps, be considered as an aid to thought rather than a special mode of thinking" (Beardsley, 1972, p. 287).

the statement: “The atom is like the solar system” - leaving room for the listener to fill in the details, and possibly to infer wrongly that the orbits of electrons and planets are similar - a model would consist of a picture, physical prototype, or mathematical description in which each element of the source would be explicitly represented by some particular aspect of the model. In other words, a model presents a precisely constrained analogy.

The above example brings to light a further distinction in dealing with scientific models; that is, the distinction between physical and computer (or mathematical[2]) models. A physical model is most common in the engineering sciences and normally consists in the construction of a prototype. In some instances, as with a wind tunnel model, the dimensions are in exact proportion to what is being modeled. The materials used may, or may not be the same, but it is definitely known that model's aerodynamic characteristics are meant to demonstrate those of the target object. In cases like the atom model, however, the representation is not one of exact proportions, but rather an explicit mapping from one physical domain to another (see Hesse, 1972) .

On the other hand, computer and mathematical models have no similar sort of physical existence[3]. So, our distinction between physical and computer models comes down to one of the activities of the modeler. It is exactly the physical nature of the prototype, or model, that the modeler is concerned with in the first instance, whereas the computer modeler is not interested in manipulating the physical instantiation of his or her model, but rather the functions which give rise to a particular physical instantiation. In other words, you are quantitatively changing the properties of the model with respect to the source in a computational model whereas the prototype simply changes quantitative properties of the source.

Thus, computer and mathematical modeling use equations and symbol manipulation to describe the system being modeled, rather than relying on a physical reproduction of the system to provide the description. More precisely a

---

[2] Technically, any mathematical model can be *approximated* by a computer model, and any computer model can be exactly represented as a set of equations.

[3] That is, the model itself has a qualitatively different physical instantiation. Obviously the computer has a physical existence but it is not the model. Technically the computer's program exists in its storage medium and in its particular activation of electrical impulses within the computer. However, a physical system/program mapping has a non-physical target, namely the behaviors of particular equations or, more generally, functions.

functioning computer model can be viewed as an *implementation* of the mathematical model (Ubbink, 1961) . This type of computer modeling, commonly in the form of a *program*, is often referred to as *simulation*.

Simulation is by far the most common method of modeling used in cognitive science. It is so common, in fact, that it is quickly becoming imperative for any theory of cognition to be simulated successfully in order for it to be taken seriously. Furthermore, cognition is so complex that simulation models are often the only source of empirical verification for a theory, and help to guide the theory's evolution: “models help organize the data and motivate experiments; they suggest how data might fit together to yield an explanation of a phenomenon” (Churchland and Sejnowski, 1992, p. 6) . In other words, if the simulation behaves similarly to a natural cognizer, the underlying model and theory are considered to be good (Thagard, in press) .

1.1.3 Models and Theory

For our purposes, the relation between models and theory is one of mutual dependence. Although it is more common for a theory to be formulated prior to its testing via models, the results from testing the models may lead to clarification, or conversely, complete rejection of a theory. Thus, the theory is in some ways logically prior[4], but is not completely independent from its models: “Soar is a model of the human cognitive architecture - that is, it embodies the theory” (Newell, 1990, p. 207) . In cognitive science, and on the Popperian view of scientific explanation, without any model there is no support for the acceptance of a theory. Thus, without a model, the statement which may be called a theory is not falsifiable (*i.e.* “no elimination of error” (Popper, 1981, p. 83) , and should be correctly termed a *hypothesis*. Of course, in practice theories are sometimes held without strong empirical support, as we shall see. Other views of the status of theories, such as those of Putnam and Newell, lend theories a stronger resistance to controversy and dispute (Putnam, 1981; Newell, 1990) .

Unfortunately, it is often true that “what is called a model is also called a theory” (Achinstein, 1968, p. 212) . But this does not mean that what *is* a model

---

[4] Hesse correctly notes that "the notion of model is not dependent on prior development of a formal theory (in science)" (Hesse, 1972, p. 356), however, this does not put into question its logical priority.

is a theory. A model describes, through resemblance and often for the purposes of empirical verification, the structures and mechanisms posited by a theory. Clearly, one theory may be demonstrated by many models (Churchland and Sejnowski, 1992, p. 6) . By entwining model and theory it becomes impossible for secondary observers, with different, possibly superior resources, to test a researcher's theory. In other words, if a particular model is equated with its theory, there is no way to verify the theory separately from the model. Furthermore, a novel model could reveal hidden strengths or weaknesses of its theory. Therefore, model and theory are much more usefully viewed as distinct, though intimately related.

Occasionally, a confusion between model and theory arises from the haphazard use of the term *theoretical model*. An example of this can be found in Achinstein (1968), with the claim: "a theoretical model describes a type of object or system by attributing to it what might be called an inner structure, composition or mechanism, reference to which is intended to explain various properties exhibited by that object or system" (Achinstein, 1968, p. 212) . Similarly, synonymous use of the terms *model* and *picture* when referring to *theory* can be found in some relevant literature (Ubbink, 1961, p. 179) , though this usage is too imprecise for our purposes.

Rather, it is more useful to realize the distinction that the theory attributes an inner structure to a system and identifies its mechanism whereas the model plays the role of demonstrating how this mechanism behaves and of rigorously representing that inner structure: "A cognitive *theory* postulates a set of representational structures and a set of processes that operate on these structures. A computational *model* makes these structures and processes more precise by interpreting them..." (Thagard, in press, p. 2) . Without this distinction, verification of the theory becomes dependent on a particular model, which contradicts the logical priority of the theory. Since we are to be examining competing *theories* of cognitive science, this distinction is important in preventing us from rejecting a theory on the basis of one poor model.

So, more generally, a theory is a speculative explanation of the functionings and nature of a particular system. It attempts to explain how a system works, addressing the organization and relationships of the system's elements. For Koertge, explanation of this sort is "genuine" only if a theory gives an account of the system's causal mechanisms (Koertge, 1992, p. 91). However, he also notes that Woodward (1992) argues that many seemingly explanatory advances in science are not causal. So, for our purposes we will be non-restrictive in the criteria we place on an explanation, it may be causal or not.

Finally, theoretical unity has been heralded as a main aim of scientific inquiry. Such a belief has propelled physicists to seek the elusive Grand Unified Theory (GUT) of forces (Hawking, 1988). Not surprisingly, in cognitive science the search for unified theories has also arisen: “Psychology has arrived at the possibility of unified theories of cognition - theories that gain their power by positing a single system of mechanisms that operate together to produce the full range of human cognition” (Newell, 1990, p. 1). Interestingly, the unity criteria plays a large role in the connectionist rejection of classicism, the classicist rejection of connectionism, and the dynamicist claim to oust both of these views. All three camps cheerily point to evidence that the other theories will fall apart at some level of cognitive modeling and thus fail to meet the unity criteria. Nevertheless, representatives of all approaches to cognitive science and many philosophers of science have upheld the criteria of theoretical unity (Hesse, 1972; Churchland, 1989; Newell, 1990; Koertge, 1992; van Gelder and Port, in press).

1.1.4 Cognitive Theories

In cognitive science, testing a theory's usefulness most often results in, and depends on, its descriptions being implemented in a model through realistic simulation. In general, a simulation is expected to provide some sort of prediction. As per the unity criteria, the better the theory behind the model, the more related systems or phenomena it can simulate and predict.

These properties of cognitive theories are clearly elucidated in Churchland's definition of the ideal cognitive theory (1989, p. 18):

> [The theory should] yield excellent explanations/predictions of internal change and external behavior, at least in the short term. As for long term activity, the theory provides powerful and unified accounts of the learning process, the nature of mental illness, and variations in character and intelligence across the animal kingdom as well as across individual humans.

As of yet, no cognitive theories have come close to proving that they are ideal in this sense. However, that does not mean that the theories are to be flat out rejected: “Theories cumulate. They are refined and reformulated, corrected and expanded” (Newell, 1990, p. 14). This evident rejection of Popperian theoretical ideals in science is not uncommon amongst cognitive scientists. However, a theory which does not have any sort of empirical support will be short-lived in the field. Theories may be non-ideal but they may not be unsupported.

What this means in cognitive science is that rigorous, often computational, models must support a theory. As Thagard notes: “Cognitive theories by themselves are normally not precise enough to generate such quantitative predictions, but a model and program may fill the gap between theory and observation” (Thagard, in press, p. 3). Here again lies the reason that cognitive scientists are so interested in computational simulation. Only with such an implementation can the theory be tested and proven to behave as a natural cognizer. Thus, especially in cognitive science, the relationship between model and theory is one of strong mutual dependence. It is thus mandatory to examine implementations of classicist, connectionist, and dynamicist theory to wisely evaluate their success, limitations, and future prospects, as we do in later chapters.

1.2 Applying Models

It is imperative to make clear the importance of computational simulation in the validation of a cognitive theory: “for a system as complex as the mind, it is difficult to have much faith in a theory without having it instantiated in some operational form” (Newell, 1990, p. 194). More often than not, it is the simulations and models of a theory which are critically examined by competing approaches. The reason is simply that simulations tend to objectify part of the body of knowledge on which a particular theory relies. Thus, everyone has unbiased access to the information generated by and contained in a simulation; it is interpretations of this information that produce conflicting opinion.

A general method for validation of cognitive models is as follows. The computational model is initially made to behave similarly to natural cognizers in some particular way. Once a resemblance to natural state transitions is substantiated, the model will be applied to psychological tests for which it was not specifically designed. And possibly, the model can be used to predict the behavior of natural cognizers on newly designed tests. If a model provides accurate predictions *etc*., the theory behind the model is verified for the particular areas of cognition addressed by the psychological tests (see Newell, 1990).

Of course, such a method seldom covers a wide variety of cognitive behaviors. As Churchland and Sejnowski point out, "what goes into the model depends on what one is trying to explain" (Churchland and Sejnowski, 1992, p. 136). Thus, models are often quite tailored to explain particular cognitive phenomena, purposefully ignoring the unity criterion in their construction. However, this criterion does play a role in the evaluation of the model. Precisely, the more important properties of the source that are exactly demonstrated by the model, the better the model "one intends that a model capture the *salient* features of reality modeled" (Churchland and Sejnowski, 1992, p. 15, italics added). Of course, determination of what is *important* is often the subject of heated debate: "any given modeler tends to think his favored modeling level is *the* important level, that lower levels are properly ignorable, and modeling levels higher than his favored level is premature and unrewarding" (Churchland and Sejnowski, 1992, p. 137).

A further difficulty with this method of validation is that often a cognitive model may predict behavior, or mechanisms of behavior, for which it is extremely difficult to design psychological tests. Furthermore, the degree of "resemblance" to a natural cognizer exhibited by a model is often quite difficult to determine and gives rise to varied interpretations of the model's success. Thus, it is important for us to outline some criteria we can use to determine the *degree* of the validity of a particular computational model: criteria which are not *solely* dependent on the model's ability to provide results similar to known psychological tests.

As Thagard notes, a theory is potentially considered valid if the resulting model can be formulated in computational terms, as this "helps to show that the postulated representations and processes are computationally realizable." (Thagard, in press, p. 3). Evidently it is assumed that computational viability is necessary to model cognitive behavior. Of course, there is room for disagreement here. However, for practical purposes, if computational simulations are not used to implement the models of brain functioning, there is no way to model the system at all. Thus, this criterion has implications which gives it, at the very minimum, pragmatic value.

Once it has been established that a theory is computationally realizable, the ability of this realization to predict a system's behavior becomes our next main interest. The role of prediction, and its importance has already been established (see section 1.1.4), and will not be considered further here. Still, this criterion ties in with the additional criterion of explanation. “Scientific explanation is generally agreed to have two main aims:” claims Hesse, “successful prediction of the behavior of things in experimentally defined conditions, and theoretical representation of the causal structure of the world from which this behavior follows” (Hesse, 1972, p. 325). Since explanation is often considered the purpose of all scientific inquiry (Hesse, 1972; Cartwright, 1991; Koertge, 1992; *c.f.* Le Poidevin, 1991), it deserves further attention.

In cognitive science, as we have examined it so far, neither the theory nor the model can unilaterally account for explanation. Both a theory's representation and the model's prediction play an important role. Firstly, the predictive evaluation must be based on the similarity of behavior of the model (*i.e.* the simulation) to a natural cognizer under a variety of conditions. Secondly, an evaluation of the usefulness of a theory's representation of cognitive functioning can also only be based on the behavior of the model. This is simply because we do not have enough empirical information to simply rule out particular methods of representation (or lack of representation) of a theory. Therefore, all we can do is assume that the better the predictions provided by a model, the better the representations must be and the better the explanation that can be provided. So again, we see how much a theory relies on its implementation for validation.

Prediction and explanation must also lie in the compass of the further criterion of unity. As has been mentioned (see section 1.1.3), the *number* of salient properties of a system that can be predicted by a particular theory is of great consequence; the more, the better unified a theory. Some philosophers of science even choose to have the explanation criterion rely on the unity criterion: “For Kitcher the key ingredient in explanation, and thus the fundamental aim of scientific inquiry, is unification. There is no question that scientists place a high value on unified conceptual systems... Kitcher posits that the aim of science is not just knowledge, but *unified* knowledge. It is its unified character which makes science explanatory” (Koertge, 1992, p. 93). It is not *necessary* for us to ally ourselves with this particular stance. However, the importance of theoretical unity should be well taken. It is very impressive for a highly unified theory of cognition to survive much scrutiny because a unified theory, especially in the area of cognition, is more easily open to attack as it will necessarily be making claims about vastly differing instances of cognition; ranging from perception to logic, and from reflex to intent.

Our criteria for evaluating cognitive theories have so far focused mostly on theoretical considerations. However, there are significant secondary criteria which often lend credibility to a theory. First, the ability of a paradigm to stimulate further research can be an important indication of, if nothing else, how many other researchers approve of the theory. Of course, it would be delinquent to claim that science should be democratic. Nevertheless, the longer a theory generates interest and survives scrutiny, the more reasonable it seems.

Second, after Langer (1957), a good theory should pose new questions while addressing a core issue of the target of the theory (*i.e.* it should be projectible). This is a criterion similar to that of stimulating new inquiry, but can also be used independently. If core problems in a particular field are addressed (*e.g.* language, logic, object perception, recognition, *etc*. in cognitive science) by a particular theory, and as a consequence interesting questions can be posed in such a framework, the framework has proved itself at least philosophically interesting.

These latter criteria are far more subjective in nature and are overshadowed in importance by the former. However, they can be useful in understanding why theories occasionally outlast their usefulness (Newell, 1990, p. 14) and in finding merit in theories which might otherwise be cast aside. Nevertheless, they remain secondary to our theoretical considerations and cannot be expected to significantly influence our examination of the three competing cognitive theories.

1.3 Summary

In order to develop a reasonable evaluation of the symbolic, connectionist, and dynamicist approaches to cognitive science, it is important to understand distinctions between metaphor, analogy, and model and to outline criteria for evaluating cognitive theories.

Metaphor and analogy are synonymous and share, with models, the property of mapping one situation to another. However, models are more precise in their description of these relations. Theories are logically prior to models, in that they identify the mechanisms of behavior which are described by models. Thus models are implementations of a theory.

In the context of congitive science, good theories have four important criteria under which they should be evaluated. First, theories should be predictive. Second, a theory should provide some sort of explanation of the cognitive behavior. Third, the more important features of cognition that a theory describes, the more unified it becomes and thus the better the theory. Lastly, a cognitive theory should lend itself to computational simulation to allow for empirical validation of its models.

Chapter 2

Comparison of Cognitive Theories

Dynamicism arose from a perceived need for a new conception of cognition. Since many of the powerful criticisms which the symbolicist and connectionist paradigms leveled at one another remained unanswered, it seemed there must be a better approach to understanding cognition. But, more than this, there are a number of cognitive issues which dynamicists feel are inadequately addressed by either alternative approach. Dissatisfaction with the symbolicist *computational hypothesis* and the *connectionist hypothesis* on the basis of their different emphases and conceptions of time, architecture, computation and representation have lead dynamicists to forward their own *dynamicist hypothesis*.

In this chapter, I have chosen a champion of each of symbolicism, connectionism, and dynamicism. Newell, Churchland and Sejnowski, and van Gelder and Port, will respectively represent the views of the three approaches. Of course, there are many schools within each of symbolicism, connectionism, and (less so) dynamicism, not all of which will be represented by this choice of exemplars. However, each of these researchers has written a comprehensive text which has become accepted as representative of the views of their respective approach. Furthermore, by choosing a single representative of each field, we will be able to easily extract concise explanations of important issues from experts in the field of cognitive science. This will allow us to compare and contrast widely varying perspectives in a reasonably limited discussion.

This discussion will commence with an introduction to the dynamicist approach to modeling cognition. Subsequently, this view will be compared with those of the symbolicist and connectionist approaches. This critical comparison of these competing theories will use the most controversial aspects of cognition - namely time, architecture, representation and computation - to highlight the relative strengths and weaknesses of each of the approaches. The discussion will make it evident that dynamicists have a number of persuasive arguments against the symbolicist approach. However, many of these same arguments cannot be leveled effectively against connectionism. Thus, it remains difficult to distinguish connectionist and dynamicist views. However, it will be left to the third chapter to establish the exact place of the dynamicist theory of cognition.

2.1 An Introduction to Dynamicism

2.1.1 Dynamical Systems

Taken at its most literal, the class of *dynamical systems* include any systems which change through time. Clearly, such a definition is inadequate, since both connectionist nets and symbolicist algorithms are dynamic in this sense (Guinti (1991) as in van Gelder, 1993). Understandably, the dynamicist will wish to delineate a specific *type* of dynamical system that is appropriate to describing cognition. This is exactly what van Gelder accomplishes with his assertion of the *Dynamicist Hypothesis* (in press, p. 4):

> Natural cognitive systems are certain kinds of dynamical systems, and are best understood from the perspective of dynamics.

Consequently, the dynamicist must define those “certain kinds” of dynamical systems which are suitable to describing cognition. Clearly, this is the class of systems that will be of interest to our discussion of the dynamicist approach in the context of the connectionist and symbolicist approaches and will thus be the only class we will describe in any detail. The “certain kinds” of systems dynamicists are interested in are “those state-determined systems whose behavior is governed by differential equations... Dynamical systems in this strict sense always have variables that are evolving continuously and simultaneously and which at any point in time are mutually determining each other's evolution” (van Gelder and Port, in press, p. 5); in other words, systems governed by coupled nonlinear differential equations (van Gelder and Port, in press, p. 6).

More precisely, these systems must have a number of component behaviors which are intimately linked, described with respect to the independent variable of time, of low dimensionality, deterministic and generally complex. The coupled, or intimately linked nature of such a system of equations implies that changes to one component (most often reflected by changes in a system variable) immediately effect other parts of the system. Insisting on a description with respect to time gains its impetus from two sources. The first is the belief that the other approaches to cognition “leave time out of the picture” (van Gelder and

Port, in press, p. 2). The second is the observation that we, like many natural systems which are effectively described by dynamical systems, evolve through time. The low dimensionality of the systems is a feature which contrasts the dynamicist approach with the connectionist[5]. By noting that certain dynamical systems can capture very complex behavior with low dimensional (*i.e.* few variables) descriptions, dynamicists have insisted that complex cognitive behavior should be modeled via this property. This insistence has the effect of avoiding difficult analyses of high dimensional systems, as experienced by the connectionist, but also of making the choice of equations and variables very difficult and somewhat arbitrary. The deterministic nature of the behaviors of these systems is simply a result of their being described by deterministic mathematical tools. Lastly, complexity is a desirable trait of the appropriate class of systems since human behavior, the target of description, is quite complex (van Gelder and Port, in press).

Furthermore, the desired class of cognitive dynamical systems has a special relation with their environment. They are not easily distinguishable from their surroundings. Rather, a particular cognitive system is linked with its environment in the same manner in which it is linked to other parts of itself. In other words, since the environment is also a dynamical system and since it is affecting the cognitive system and the cognitive system is affecting it, they are coupled. Such *embeddedness* of the cognitive system makes a distinction between the system and the system's environment occasionally appear arbitrary. But this fact, the dynamicist would claim, is not only a good reflection of how things really are, it is a unique strength of the dynamicist approach (van Gelder and Port, in press, p. 23). Couplings amongst not only the equations describing a cognizing system, but also between that system and its environment will result in complex temporal behaviors. Conveniently, mathematicians and scientists have dealt with these behaviors for some time, using and developing the powerful tools of dynamical systems theory.

---

[5] Dimensionality is directly related to the number of parameters in a given system of equations, thus a low-dimensional system has fewer variables. Connectionist systems have the same dimensionality as number of nodes in the system. Considering even simple, limited connectionist models can have hundreds of nodes, it is not surprising that connectionist implementations are generally considered high-dimensional.

2.1.2 Dynamical Systems Concepts

Dynamical systems theory describes systems using essentially geometrical terminology. Concepts commonly employed by dynamicists include: state space, path or trajectory, topology, and attractor. The state space of a system is simply the space defined by the set of all possible states that the system could ever pass through. Thus, a trajectory or path plots a particular succession of states through the state space and is commonly equated with the behavior of the system. The topology of the state space describes the attractive properties of all points of the state space. Finally, an attractor is a point or path in the state space towards which the trajectory will tend when in the neighborhood of that attractor.

Using these, and related concepts, dynamicists can thus attempt to predict the behavior of a system if the set of governing equations, and a state on the trajectory is known. Thus, a dynamical system can be further categorized as “any system which evolves over time such that its state always depends in a rule-governed manner on its previous state” (van Gelder and Port, in press, p. 5)[6].

Difficulties arise for dynamical descriptions when either the set of equations is not perfectly known, or when the system is extremely sensitive to initial conditions and the state is not perfectly characterized. The former problem generally arises from a lack of empirical knowledge, but the latter difficulty is much more profound and is often characteristic of a system. The examination of this sensitivity to initial conditions of some dynamical systems has given rise to the field of chaos theory.

The advantages of dynamical descriptions are multitude. Even the disposition of some dynamical descriptions to exhibit chaotic behavior can be considered an advantage if that is truly how the systems they are describing, *i.e.* cognitive systems, behave. Two important aspects of dynamical descriptions which are perhaps most clearly advantageous are the ability to exhibit continuous *or* discrete behavior (catastrophe theory), and that it is the nature of dynamical systems descriptions to incorporate *real-time* considerations. These two aspects alone make dynamical systems theory a good candidate for effectively describing natural cognitive processes (van Gelder and Port, in press).

---

[6] Clearly such a characterization still includes connectionist and symbolic models in the same class.

2.1.3 Cognition and Dynamics

The power of dynamical systems theory has been demonstrated through its application to many natural phenomena ranging from turbulence and cells to weather patterns and entire ecosystems. Still the questions remains: Why should we apply these tools to a cognitive system? Why should we accept the claim that "cognitive phenomena, like so many other kinds of phenomena in the natural world, are the evolution over time of a self-contained system governed by differential equations" (van Gelder and Port, in press, p. 6)?

The van Gelder and Port answer is: "Since the nervous system, body and environment are all continuously evolving and simultaneously influencing each other, the cognitive system cannot be simply the encapsulated brain; rather, it is a single unified system embracing all three" (van Gelder and Port, in press, p. 9). Such observations of dynamical coupling lead van Gelder and Port to conclude that dynamical descriptions of cognition are not only sufficient, but also necessary for an understanding of mind: "...whenever confronted with the problem of explaining how a natural cognitive system might interact with another system which is essentially temporal, one finds that the relevant aspect of the cognitive system itself *must* be given a dynamical account" (van Gelder and Port, in press, p. 24, italics added). This strong commitment to a particular form of modeling has resulted in the dynamicists claiming to posit a "paradigm for the study of cognition" (van Gelder and Port, in press, p. 29); not, notably, an extension to either of connectionism or symbolicism, but a new *paradigm*. Thus, the dynamicists are insisting that there is an inherent value in understanding cognition as dynamical *instead of* connectionist or symbolicist (van Gelder and Port, in press).

Some critics will claim that a dynamical systems approach to cognition is simply not new (Simon and Newell, 1970, p. 273). Giunti (1991) showed that the symbolicist Turing Machine *is* a dynamical system (van Gelder, 1993). So, it could be concluded that there is nothing to gain from introducing a separate dynamicist method of studying cognition. However, Turing Machines and connectionist networks have also been shown to be computationally equivalent (Fodor and Pylyshyn, 1988, p. 10) yet the approaches are vastly disparate in their methods, strengths, and philosophical commitments. Similarly, though Turing Machines are dynamical in the strictest mathematical sense they are nonetheless serial and discrete. Hence, symbolicist models do not behave in the same ideally coupled, dynamical and continuous manner as dynamicist systems are expected to. They are not linked in the same way to their environment, and the types of processing and behavior exhibited is qualitatively different. Thus it is difficult, in any practical way, to see dynamicist and symbolicist models as truly equivalent (Simon and Newell, 1970).

However, Smolensky's (1988) claim that connectionism presents a dynamical systems approach to modeling cognition is different altogether. Connectionist nets *are* inherently coupled, nonlinear, parallel dynamical systems. These systems are self-organizing and evolve based on continuously varying input from their environment. Still, dynamicists claim that connectionist networks are limited in ways that a *truly* dynamical description is not. Nevertheless, the relationship between connectionism and dynamicism is undeniably more intimate than that between either of these approaches and symbolicism.

An important way in which dynamicists wish to distinguish themselves from any other approach is through their view of representation. To be a truly dynamical system, there should be *no* representation in the cognitive model. Symbolicist models are based on symbolic representations, so clearly they are inadequate. Connectionists represent concepts (via either distributed representation or local symbolic representation) in their networks, so the dynamicist concludes: “It is the processing of representations that qualifies simplified nets as computational (*i.e.* symbolic). In realistic nets, however, it is not the representations that are changed; it is the self-organizing *process* that changes via chemical modulation. Indeed, it no longer makes sense to talk of “representations”” (Globus, 1992, p. 302).

For dynamicists, a cognitive agent can best be described by a certain type of system of equations which display the same dynamical characteristics as that agent. It should be the goal of cognitive science to discover those dynamical properties since only a system which has the “right dynamical properties” will exhibit cognitive capacity. Furthermore, these properties can be captured with relatively[7] simple sets of equations. In other words, the aim of dynamicists is to “provide a *low-dimensional* model that provides a scientifically tractable

---

[7] This relative simplicity is with respect to modeling the dynamics of every cell in the nervous system, thus the resulting dynamical models can still be extremely complex.

description of the same qualitative dynamics as is exhibited by the high-dimensional system (the brain)" (van Gelder and Port, in press, p. 28).

2.2 Symbolicism and Dynamicism

2.2.1 Symbolicism

Though I have chosen Newell's implementation of Soar as an exemplar, work by Chomsky, Minsky and Anderson is also often cited as defining the symbolic approach to cognitive modeling (see van Gelder and Port, in press, p. 1). However, Newell and Simon (1976) are cited by van Gelder as having best defined the *computationalist hypothesis* with the following, which they refer to as the *Physical Symbol System Hypothesis* (see Newell, 1990, pp. 75-77; van Gelder and Port, in press, p. 4):

> Natural cognitive systems are intelligent in virtue of being physical symbol systems of the right kind.

Thus, Newell is a natural choice as a champion of the symbolicist approach. Newell's characterization of the approach is widely accepted and his assertions will be the focus of this discussion of symbolicism[8].

Much of Newell's work has been aimed at the actual construction of the right kind of physical symbol system. His particular model is called Soar. Soar is an attempt to produce at least the beginnings of a complete cognitive system as the implementation of a unified theory of cognition. Though Anderson's (1983) Act* is classified by Newell as a unified theory of cognition, Newell believes that Soar will expand extensively on the range of behaviors exhibited by Act* (Newell, 1990, p. 29). To quickly characterize Soar, it is best to turn to Newell's own description (1990, p. ix, 39):

> Soar has its roots in thirty years of research on the nature of artificial intelligence (AI) and human cognition. A number of basic concepts have been explored along the way, prominent among them being problem spaces, goal-subgoal hierarchies, weak methods for problem solving, and production systems...[Soar] has a symbol and goal hierarchy. It uses a production system as the foundation of the architecture, thus being a kind of recognize-act

---

8 The more controversial aspects of Newell's theory, such as chunked learning and memory characterization (Newell, 1990), will thus not be discussed here.

> system. It uses problem spaces everywhere to do all of its business, that is, to formulate all its tasks.

Soar is undeniably an impressive model. It has been tested on many tasks, including: AI toy problems (*e.g.* Blocks World, Missionaries, Tower of Hanoi, the Eight Puzzle), expert systems (*e.g.* medical diagnosis, computer system configurations), forms of parsing, version spaces (a scheme for learning concepts), syllogistic reasoning, algorithm design, and cryptarithmetic (Newell, 1990, p. 217). All of these tasks were tackled, many quite successfully, relying on the symbolic approach.

Symbolicists claim that “humans have programs, indeed, that they have analogous programs” (Newell, 1990, p. 248). However, such a claim cannot be simply interpreted as claiming that humans are just like the familiar serial digital computer. Rather, the symbolicist introduces the abstract characterization of knowledge-level systems and symbol systems (Newell, 1990, p. 51, 76). The former is a system which processes knowledge about the system’s goals, actions, environment and the relations between these items. Furthermore, the system has a single law of behavior which is: take actions to attain goals using all of the knowledge available to the system (Newell, 1990, p. 50). A symbol system is a form of a universal computational system (*e.g.* a Turing Machine) which Newell defines in terms of general characteristics of memory, symbols, operations, interpretation and capacities (Newell, 1990, p. 77). The purpose of these symbol systems is to “realize knowledge systems by implementing representation law so that the symbol structures encode the knowledge about the external world” (Newell, 1990, pp. 78-79). So, through application of these definitions, natural cognizers are defined more precisely: “humans are symbol systems that are at least modest approximations of knowledge systems” (Newell, 1990, p. 113). Soar, of course, is attempting to achieve this same distinction. Thus, it is necessary for Soar to be a symbol system, which all computers are, and to approximate a knowledge system to the same degree that humans do in order for it to be a complete account of cognition.

The reason cognition is cast in terms of knowledge-level systems is that this level can be described “before anything about the internal workings of the system is determined” (Newell, 1990, p. 50). Thus symbolicists wish to implement generic knowledge-level descriptions of behavior and they feel the most appropriate way to accomplish this task is to discover “the symbol-level mechanisms that permit a close approximation to the knowledge level" (Newell, 1990, p. 80). Obviously, such a description of cognition will likely conflict with the dynamicist assertion that cognition is embedded, non-representational and described by coupled nonlinear differential equations. Indeed, the conflict is not minor. By examining the most contentious aspects of descriptions of cognitive functioning - namely time, architecture, computation, and representation (see section 2.0) - the next four sections will further characterize the symbolicist view of cognition through the critical eye of the dynamicist.

2.2.2 Time

The purpose of cognitive science is to describe the behavior of natural cognizers. By definition natural cognitive systems operate in real time - in the real world. Therefore, a successful cognitive science model must (unanimously among the three approaches) explain cognitive processes which can occur in real time.

In an attack on the poor temporal nature of symbolic systems, van Gelder and Port claim that symbolicists "*leave time out of the picture*" (van Gelder and Port, in press, p. 2). However, Newell includes operation in real time as the third most important constraint that shapes the mind (Newell, 1990, p. 19). Thus, it is not the case the symbolicists ignore time, but it may be the case that they have difficulty meeting temporal constraints.

Newell uses neurological data to lend support to his assumption that any particular step in a cognitive algorithm operates on the time scale of approximately 10ms (Newell, 1990, p. 127). However, application of this constraint seems rather contrived[9] and more importantly is inconsistent. For instance, one Soar application employs a single production to encode whether or not a light is on (Newell, 1990, p. 275) and a second application uses a single production to encode: "If the problem space is the base-level-space, and the state has a box with nothing on top, and the state has input that has not been examined, then make the comprehend operator acceptable, and note that the input has been

---

[9] Time values and the number of productions per step seem to have been chosen to allow the reaction time models to fall within human limits found through psychological experimentation. The claim that Soar has somehow allowed rough predictions of human reaction time is very unconvincing. It is rather more likely that the modeler's analysis, experience with psychological results, and chosen constants allowed such rough predictions (Newell, 1990, pp. 274-282).

examined" (Newell, 1990, p. 167). It seems unrealistic that both of these productions should fire on the same order of magnitude, *i.e.* approximately 10ms.

Though the dynamicist would likely concur with Newell's classification of different time scales in the nervous system (Newell, 1990, pp. 121-131; *c.f.* van Gelder and Port, in press, p. 20), they are incorporated into the approaches very differently. For the dynamicist, time is incorporated via the very mathematical tools used to model the cognitive system. Thus, a single dynamical description can have "both '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" (van Gelder and Port, in press, p. 21). On the other had, symbolicists attach qualitatively different explanations at different time scales. Thus, it is reasonable to ignore lower time scale in an explanation of a higher one - for the dynamicist, this is impossible (Newell, 1990, p. 154).

The power of dynamical systems theory to model real time phenomena is so impressive that van Gelder and Port exploit this property to support a major objection to symbolicism: "Cognitive processes always unfold in real time. Dynamical models specify in detail how processes unfold in real time. Computational models, by contrast, specify only the mere sequence of states that the system passes through. Consequently, dynamical models are inherently superior to computational models" (van Gelder and Port, in press, p. 14).

Though it may not be completely futile for the symbolicist to attempt to incorporate realistic time constraints into his or her model, it is undeniably more natural for this constraint to be satisfied by dynamical models. Currently, the symbolicist cannot confidently assert how time in their model of cognitive processes relates to time in the natural cognizer. For, as Newell himself notes: "minor changes in assumptions move the total time accounting in substantial ways that have strong consequences for which model fits the data" (Newell, 1990, p. 294).

2.2.3 Architecture

An architecture provides a boundary, of sorts, which separates the structure of the system from its content. For Newell, this means that "behavior is determined by variable content being processed according to the fixed processing structure, which is the architecture" (Newell, 1990, p. 82). In contrast, the dynamicist would insist that both the structure and content are variable, and that it is wrong to assume this fixedness in the face of a dynamical changing environment and agent. However, it seems biologically well-founded to assert that the human architecture is “a hierarchy of multiple system levels and that it cannot be otherwise structured” (Newell, 1990, p. 117; Churchland and Sejnowski, 1992; Churchland, 1993). The more important concerns that arise are how those system levels interact and which levels are more or less important for effectively modeling cognition.

For the symbolicist, there is a definite gap between “central cognition” and perception and motor behavior (see figure 2.1). Taken together, these levels combine to form the “total cognitive system.” Newell realizes that cognitive scientists “need to describe this structure as well, for it is critical to a unified theory of cognition” (Newell, 1990, p. 194).. Nonetheless, Newell has admittedly focused on “central architecture” and points out that “Soar has not yet dealt with any of the central phenomena of perception or motor behavior” (Newell, 1990, p. 304). However, the symbolicist does not see this as a major drawback to a theory of cognition, because it is generally accepted that “*mind* (is) the control system that guides the behaving organism in its complex interactions with the dynamic real world” (Newell, 1990, p. 43). Thus, central cognition supervises the other aspects of cognition and seeing as it is at the top of the hierarchy, it is the most interesting (perhaps) place to begin an examination of cognitive behavior.

Figure 2.1: The symbolicist cognitive architecture.

For the symbolicist, there is a strict division between central architecture and the peripheral cognitive activities: so for perception, “the arriving perceptual elements and the encoding activity are outside central cognition, that is, inaccessible to it” (Newell, 1990, p. 257); or in the case of motor control, “cognitive intention produces a command to release motor action” (Newell, 1990, p. 259). Of course, these commands are not always *conscious* commands (*e.g.* the command to beat the heart), but they *are* symbols which are encoded by central cognition and then decoded and executed once they have been passed to the motor system. These sharp distinctions between perceptual, central and motor systems, and the multiple encodings and decodings necessary[10] for cognitive functioning

[10]"The original external situation is *encoded* into an internal situation. The external transformation is also *encoded* into an internal transformation. Then the internal transformation is *applied* to the internal situation to obtain a new internal situation. Finally the new internal situation is *decoded* to an external situation.” (Newell, 1990, p. 59).

are a source of great controversy. Newell himself realizes the debate over the viability of this architecture is neither easily resolved: “unfortunately, the nature of the command language to the motor system is exactly where obscurity is deepest...” (Newell, 1990, p. 259), nor trivial: “interestingly, the major barrier to a more complete explanation is the poorly defined motor system in Soar” (Newell, 1990, p. 301).

Not surprisingly, these “barriers” are reiterated and elaborated by competing cognitive approaches in order to discredit symbolicism. Van Gelder and Port heavily criticize the symbolicists view of “the cognitive system as a box inside a body, which in turn is inside a physical environment” (van Gelder and Port, in press, p. 8). It seems to the dynamicist that distinctions between the environment and the agent, and the distinctions between the agent's perception, motor control and central cognition, are very artificial. Alternatively, the dynamicist sees all aspects of cognition and the environment as mutually influencing each other continuously, without parameter passing, or encoding and decoding of representations or structures. For the dynamicist, symbolicism has been misguided in an important way. Though a particular research project may focus on central cognition, there must be a commitment to an approach which takes into account the total cognitive system. Furthermore, that approach must unequivocally deny that there exist divisions between perception, central cognition, and motor behavior.

A dynamicist might question the grounds for insisting upon this distinction. Though it seems evident that there are certain areas of the brain which are dedicated to general categories of processing (Donald, 1991, p. 46) it is not clear that these areas can be divided into the three categories that the symbolicists have chosen. What is even less likely is the idea of a high-level command language which would activate complex coordinated movements.

More importantly, it seems that Newell's distinction between perceptual and central cognition is invalid. Kossyln (1980) notes that rotating mental objects activates the visual buffer in the brain. Newell identifies imagery as being part of central cognition, so we would have to include the visual buffer as part of central cognition. However, the optic nerve terminates in the visual buffer. Thus, such a distinction would include almost the whole of the brain (i.e. everything except the perceptual and motor nerves terminating in the brain) in central cognition. Such an extreme biological division would have difficulties accounting for encoding and decoding in Newell's sense, as the perceptual and motor systems become mere periphery. Thus, the distinction becomes useless.

Why have such distinctions been posited? The symbolicist may claim that there is clearly a behavioral difference between these aspects of cognition, and there should be a functional one. However, such a claim is not supported by biological evidence. Newell himself realizes the importance of biological data in developing cognitive models: “information processing can no longer be divorced from brain structure because our knowledge of the biological band has begun to take on a scientific solidity that breeds confidence about the broad outlines of what is going on in the temporal space that reaches from milliseconds to behavior” (Newell, 1990, p. 483, *c.f.* Fodor, 1988 #2442, p. 62-64). However, Newell seems very selective in the biological data he uses in developing Soar. The division of the total cognitive system into perceptual, motor and central sub-systems is not biologically supported. If brain structure does not clearly support such a division, an information processing model which does may be misleading.

2.2.4 Computation

The symbolicist architecture outlined above clearly needs to identify what acts as the supervisory controller of the cognitive agent. In answering this question, an analogy is often drawn between central cognition and a computer. Hence, symbolicism is sometimes referred to as the “computational approach” to cognition, and is characterized by van Gelder and Port as “that approach (which) takes the mind to be a special kind of *computer*. The computational approach identifies the mental computer with the brain. Sensory organs deliver up to the cognitive system representations of the state of its environment. The system computes a specification of an appropriate action. The body then carries this action out” (van Gelder and Port, in press, p. 1), a characterization completely in agreement with the architecture outlined by Newell (see section 2.2.3).

From the family of universal computers, Newell clearly identifies the type of computer deemed necessary to exhibit cognition; one that will “provide a means to build representation systems” (Newell, 1990, p. 68). These representational systems are, for Newell, strictly symbol systems (Newell, 1990, p. 76). This conjecture has been labeled the *Physical Symbol System Hypothesis* by Newell and Simon (see section 2.2.1). This very strong position, has been attacked by the dynamicists in their attempt to validate their own approach to cognitive theorizing.

As van Gelder and Port note, this hypothesis creates a problem for the symbolicists in that they are unable to account easily for some aspects of embeddedness: “describing cognition in computational terms automatically creates a theoretical gap between cognitive systems and their surrounds, a gap which must then somehow be bridged” (van Gelder and Port, in press, p. 23). This is not a trivial gap to bridge. For the symbolicist, doing so means creating the motor and perceptual systems, which is a task not assured of any measure of success (see section 2.2.3). Furthermore, embeddedness is essentially a problem of behaving in time. Thus, being committed to an “atemporal” approach will cause symbolicists to present “*ad hoc*, biologically implausible” solutions (van Gelder and Port, in press, pp. 23-24).

The dynamicist answer to the embeddedness problem is not to create it in the first place. Thus, as already noted, environment, perception, motor systems, and central cognition are intimately linked and simultaneously influence each other's behavior: “*Everything is simultaneously affecting everything else*” (van Gelder and Port, in press, p. 19). These systems are not distinct in the same way as they are for the symbolicist. However, it must be noted that the functional differences of certain areas of the brain are not incompatible with the dynamicist approach. Rather, dynamicists have a great deal to say about morphogenesis - or the formation and differentiation of tissues and organs - and its influence on biological structure (Goodwin, 1994). So, it is not brain structure that dynamicists object to, it is rather the discontinuities that symbolicists have introduced through their commitment to particular types of architecture and computation.

Similarly, by insisting on supporting the hypothesis that cognitive processes can only be tractably described through symbolic computation, the symbolicist “ignores fine-grained temporal issues” (van Gelder and Port, in press, p. 17). The symbolicist even *admits* limited success (see section 2.2.2) in incorporating temporal issues into their models. So, theoretically, “*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” (van Gelder and Port, in press, p. 17) by virtue of its commitment to symbolic representational computation.

2.2.5 Representation

By defining cognitive systems as *necessarily* representational, and then claiming "*symbolic* could be defined to be essentially synonymous with *representational*" (Newell, 1990, p. 72), Newell has noted the close relationship between computational and representational commitments. From the assumption that any form of representation naturally denotes a symbolic representation, it is obvious why symbolicists in general rely on logic-like structures to represent knowledge (Newell, 1990, p. 54). Newell claims that "what a logic lets us do is represent the knowledge of X as a finite set of expressions plus a process for generating the infinite set of other expressions that comprise X's total knowledge" (Newell, 1990, p. 55).

Thus, in an implementation like Soar or Act* for that matter, all knowledge is represented in the form of a production system - a list of if-then statements. The dynamicist would note that it is important to realize that assuming "the representational medium is a list structure" (Newell, 1990, p. 61), is not based on biological analysis; and cannot be supported by such analysis. Thus, it is rather dubious that the symbolicist has empirical support for claiming: "Putatively, *all* that the lower-level systems of the biological band do is support the computational mechanisms, which in turn operate to enforce the representational laws" (Newell, 1990, p. 149). Of course, such assumptions are based on functional observations on the part of the symbolicist[11] - humans do, after all, process symbols (*e.g.* in our use of language, logic, mathematics, *etc*.).

Dynamicists tend towards a different extreme. Rather than any form of representation[12], the dynamicist wishes to rely on covariation of the system with its environment. Therefore, there is no explicit representation of anything in the system, it simply alters its configuration to reflect a change in the environment. Thus the system is influenced by the environment and the system dynamically self-organizes to reflect such influences. If our brain is coupled to the real world

---

[11]Van Gelder and Port note: "The prevalence of *innateness hypotheses* in the computational framework (Chomsky, 1988; Fodor, 1975) testifies to the inability of that framework to account for the origins of the mechanisms it postulates" (van Gelder and Port, in press, p. 21).

[12]"dynamical sub-systems...do not manipulate symbolic representations; they evolve through numerical state-spaces. They do not communicate with other modules by passing representations; they interact by being coupled (i.e., by simultaneously influencing each other's behavior)" (van Gelder and Port, in press, p. 9).

in a dynamical way, the Cartesian and symbolicist views of the brain as a self-contained representer (Newell, 1990) are inaccurate.

However, it is not easy to convincingly deny that representation plays an important role in cognition. It seems obvious that human cognizers use representation in their dealings with the world around them. For example, people seem to have the ability to rotate, and examine objects in their head. It seems they are manipulating a representation (Kosslyn, 1980). More striking perhaps is the abundant use of auditory and visual symbols used by human cognizers everyday to communicate with one another. However, dynamicists are committed to not having sub-systems manipulate representations or communicate with other sub-systems by passing representations. Exactly where the ever-present communicative representations arise from in the dynamicist approach is uncertain. It will evidently be a significant challenge for dynamicists to give a full account of human cognition, without naturally accounting for the representational aspects of thinking.

2.3 Connectionism and Dynamicism

2.3.1 Connectionism

A simple, yet accurate, definition of connectionism is provided by Newell: "connectionism is a commitment to a particular neural-like computational technology" (Newell, 1990, p. 484). In *The Computational Brain*, Churchland and Sejnowski note that their commitment to this computational technology results in the implicit hypothesis that "emergent properties are high-level effects that depend on lower-level phenomena in some systematic way" (Churchland and Sejnowski, 1992, p. 2). However, to avoid various difficulties with pure reductionism, they point out that "microlevel data are *necessary* not *sufficient*" (Churchland and Sejnowski, 1992, p. 5); a contention that likely all approaches to cognitive modeling would support. Though microlevel data is important to neural modeling, it is often the complex effects arising from the dynamics of neural networks which is of primary importance to the connectionist. Thus, the connectionist commitment is to a neural-network *type* of architecture to achieve these complex effects: "Interaction of neurons in networks is required for complex effects, but it is dynamical, not a simple wind-up doll affair" (Churchland and Sejnowski, 1992, p. 4). These commitments are effectively encapsulated by Smolensky's version of the *connectionist hypothesis*[13] (1988, p. 7):

> The intuitive processor is a subconceptual connectionist dynamical system that does not admit a complete, formal, and precise conceptual-level description.

To unravel the mysteries of natural cognizers, connectionists tend to rely on biological analysis strategies more than their symbolic counterparts: "a central part of the basic strategy for figuring out how a novel device works is reverse engineering" (Churchland and Sejnowski, 1992, p. 48). This tactic is unique among the three approaches, and is often touted as a strength of connectionism: "by contrast (to reverse engineering), a purely *a priori* approach, based entirely on reasonable principles of engineering design, may lead us down a blind alley" (Churchland and Sejnowski, 1992, p. 8).

After an analysis of the basic unit of the nervous system, the neuron, connectionists generally focus on the network level of neuronal interaction. For precisely this reason, the cognitive capacities that are most commonly, and most convincingly, modeled by connectionists are also low level: "visual motion detection rather than planning, bending in the leech rather than chess playing in the human" (Churchland and Sejnowski, 1992, p. 12). However, these models are also more easily compared to natural cognizers than are symbolic models because neural models have a more direct physical correlation to the natural cognizer.

Importantly, observable correlations between the natural cognizer and the model are reasonably independent of the modeler's potential biases or misunderstandings about the system being modeled. In other words, a connectionist net can be trained "innocent of engineering prejudices" (Churchland and Sejnowski, 1992, p. 359) and can reveal to the modeler how a task could be performed in a real neural network[14] rather than the modeler directing the model to a particular solution.

---

[13]Smolensky actually refers to this hypothesis as the *subsymbolic hypothesis* to address the distinction between local and distributed connectionist commitments. However, Churchland and Sejnowksi reject local connectionist networks as biologically unrealistic (Churchland and Sejnowski, 1992, p. 179-182) and so this form of the hypothesis is suitable for our purposes.

[14]For example in vestibulo-ocular reflex models, "the model networks, like their bone fide counterparts, discover Sherrington's law of reciprocal innervation" (Churchland and Sejnowski, 1992, p. 361). Churchland and Sejnowski observe that such a "match between model and actual networks, where the model network was not designed to have that property, but rather

Commonly, the *backpropagation* algorithm is used to train networks in this objective manner. Backpropagation has been used to construct networks that are able to “perform sensorimotor coordination, direct gaze accurately despite changes in head position, recognize subtle similarities among sonar returns, pronounce well-articulated speech from printed text, recognize three-dimensional shapes independently of the angle of illumination, parse sentences into grammatical types, recognize voiced phonemes, predict the folding of protein molecules, correctly recognize colors across changes in illumination, and so forth” (Churchland, 1989, p. xv). Some such successes have not been effectively handled by the symbolic approach[15]. However, there are also many symbolic successes which have not yet been effectively handled by connectionism.

The symbolic and connectionist approaches currently focus on different levels of cognition - though both wish to deliver a unified cognitive theory of all levels. Various arguments against the ability of symbolicism to deliver such a theory have been provided by the dynamicist critique of symbolicism (see sections 2.3.2-2.3.6). However, these same arguments do not similarly apply to connectionism. In fact, distinctions between the connectionist approach and the dynamicist approach seem to be very subtle at best - possibly non existent. Connectionists like Churchland and Sejnowski often discuss their models with dynamical systems terminology. They refer to state spaces (Churchland and Sejnowski, 1992, p. 64), point attractors (Churchland and Sejnowski, 1992, p. 88), cyclical attractors (Churchland and Sejnowski, 1992, p. 123) and trajectories (Churchland and Sejnowski, 1992, p. 64) in describing the behavior of connectionist systems.

Furthermore, any dynamical system behavior can be produced by a neural network (*e.g.* chaos, catastrophe, *etc*.). Thus, complex dynamics in no way intimidates connectionists, but is rather also seen by them as being essential to cognition: “Our brains are dynamical, not incidentally or in passing, but essentially, inevitably, and to their very core” (Churchland and Sejnowski, 1992,

developed it during the course of training” (Churchland and Sejnowski, 1992, p. 349) is extremely compelling because of the objectivity lent to such a solution.

[15]“...a simple three-layer net trained by a backpropagation learning algorithm can solve, and solved very successfully, a part of shape-from-shading tasks. This is an instructive demonstration, at least because conventional AI attempts to solve this problem have not delivered a working solution. Yet for a network, learning to estimate local shape from shading by generalizing from examples was simplicity itself” (Churchland and Sejnowski, 1992, p. 188).

p. 178). Not only is dynamics understood as being essential to higher cognition, but also to modeling the fundamental building blocks used by connectionists: “in real neurons, the dynamics are everywhere critical” (Churchland and Sejnowski, 1992, p. 174). In other words, connectionists embrace dynamics: “Using the dynamical framework, we can begin to bring nonlinear networks to heel; that is, to understand their capabilities, and most important, to give us insight into how best to design networks to solve particular computational problems” (Churchland and Sejnowski, 1992, p. 89).

Dynamical connectionist networks like those of Anastasio (Churchland and Sejnowski, 1992, p. 363) and Lisberger and Sejnowski (Churchland and Sejnowski, 1992, p. 369), are not yet as common as they perhaps should be, but they are undeniably connectionist models. However, the dynamicist, wishing to carve out a unique place in cognitive science wishes to subordinate connectionism to dynamicism. Though dynamicists know that the dynamics of connectionist networks are undeniable - “Indeed, neural networks, which are themselves typically continuous nonlinear dynamical systems, constitute an excellent medium for dynamical modeling” (van Gelder and Port, in press, p. 26) - they maintain that dynamicism is a distinct paradigm. The next four sections, under the familiar headings of time, architecture, computation and representation, discuss the reasons why.

2.3.2 Time

The timing of neural networks is, like any dynamical system, integral to the equations describing the system. Connectionists in no way need to contrive to include time in a model of cognition as van Gelder and Port have shown symbolicists often do (see section 2.3.2). Rather, network models naturally incorporate time constraints. Hence Churchland and Sejnowski claim: “A theme that will be sounded and resounded throughout this book concerns time and the necessity for network models to reflect the fundamental and essential temporal nature of actual nervous systems” (Churchland and Sejnowski, 1992, p. 117).

Connectionist models are construed as having the potential to be naturally temporal in the exact manner dynamicists claim a cognitive model must be. A connectionist network is, after all, “a dynamical system, meaning its inputs and internal states are varying with time; it is basically engaged in spatiotemporal vector coding and time-dependent matrix transformations” (Churchland and

Sejnowski, 1992, p. 338). Thus, networks which model natural cognitive processes, like the vestibulo-ocular reflex, function in the same temporal manner as their natural counterparts (Churchland and Sejnowski, 1992, p. 353). It is the process of reverse engineering which allows connectionist models to so readily exhibit the same temporal behaviors as natural systems: “evolution has evidently stumbled upon solutions to space-time representation using, among other things, time constants for various neural activities...” (Churchland and Sejnowski, 1992, p. 306); connectionists are using a very similar architecture to find, not surprisingly, those same solutions.

Furthermore, connectionist models naturally produce behaviors on any time scale, just as dynamical systems models do (*e.g.* Traub et al. as in Churchland and Sejnowski, 1992, p. 287). In other words, slow dynamical processes, as well as fast dynamical process are naturally incorporated into connectionist networks: “Speed of response is now the *only* thing that distinguishes activation from weight modification, and it does so relatively, not absolutely” (Churchland and Sejnowski, 1992, p. 177).

Of course, there are connectionist networks which do not initially address time constraints. It could even be argued that the majority of connectionist models are of this sort. Of course, these are the networks which dynamicists see as inadequate (van Gelder and Port, in press, p. 26) - with good reason. However, in many instances, as with the case of LeechNet I (Lockery et al. as in Churchland and Sejnowski, 1992, p. 347), a non-dynamical model can be modified to become a dynamical model - LeechNet II[16]. Importantly, many of the characteristics of the non-dynamical model are consistent with the dynamical model: “both LeechNet I and II attests to the fact that the simpler LeechNet I had yielded some valid results despite its dearth of dynamics” (Churchland and Sejnowski, 1992, p. 347). This indicates that the non-dynamical model is quite valuable as a stepping stone to a more complete model. Furthermore, initial non-dynamical models tend to be much simpler than their dynamical counter-parts, allowing for a simpler

[16]The modifications to create a dynamical model are explained as follows: "The basic groundplan is the same as in LeechNet I, but the model now includes chemical synapses with temporal delays, electrical synapses between the motoneurons, and an electrical time constant for each unit. The chemical synapses were modeled by two compartments called s-units, one with a fast time constant and the other with a slow time constant, chosen to match the temporal dynamics observed in these synapses" (Churchland and Sejnowski, 1992, p. 347).

initial theoretical understanding of the natural system - an option not available to the dynamicist.

2.3.3 Architecture

The connectionist commitment to reverse engineering the natural cognizer has led connectionists to claim that it is imperative to "understand the architecture of the brain itself" (Churchland and Sejnowski, 1992, p. 17) - not to simply create random, or intuitively appealing neuronal models (Churchland and Sejnowski, 1992, p. 17):

> The brain's computational style and the principles governing its function are not manifest to casual inspection. Nor can they be just inferred from behavior, detailed though the behavioral descriptions may be, for the behavior is compatible with a huge number of very different computational hypotheses, only one of which may be true of the brain...there is no substitute for conjuring the ideas in the context of observations about real nervous systems: from the properties of neurons and the way neurons are interconnected.

The basic unit of a connectionist architecture is the neuron. These units are arranged in highly interconnected networks in an attempt to reflect the brain's basic architecture. However, these networks are seldom as complex as the networks in the brain and are generally considered in isolation from the rest of cognitive machinery. Considering networks in isolation does not mean that connectionists are in any way committed to the sort of modularity that symbolicists are[17]. Rather, networks are studied in isolation to provide a certain degree of simplicity for analysis. In many cases, these "mini-nets" are interconnected to elicit new behaviors, once their properties under isolation are well understood (Churchland and Sejnowski, 1992, p. 17).

Like dynamicists, connectionist architectures tend not to separate content and structure as symbolicists do. Rather, a network's architecture affects the content by changing the type of coding available to the network, and the content of the network (*i.e.* the knowledge represented) can affect the structure of the

---

[17]Though there are specialization of areas in the brain, there is nothing analogous to Soar's complete separation of higher cognition from perception and motor systems: "Some pathways seem specialized for certain tasks, such as perception of motion, yet a single area, such as MT, cannot be fingered as *the center* for motion perception, at least for the reason that other areas also have neurons that selectively respond to motion and where lesions result in deficits in motion perception" (Churchland and Sejnowski, 1992, p. 317).

network by strengthening or eliminating connections (Churchland and Sejnowski, 1992, p. 92):

> ...notice that in the Boltzmann machine, matters of computation, of algorithm, and of implementation are not really separable. It is the very physical configuration of the input that directly encodes the computational problem, and the algorithm is nothing other than the very process whereby the physical system settles into the solution. This contrasts rather vividly with the standard separation of hardware and algorithm in digital computers.

The integration of algorithm and hardware is nowhere more prevalent than in the way memory is handled by networks. The information processing structures (*i.e.* the networks) are themselves modified with the addition of any experience. It is in this manner that memories are stored - they are not filed away in some previously blank space as is the case with a production system like Soar (Churchland and Sejnowski, 1992, p. 317).

The ability to self-organize, which the connectionist architecture is particularly adept at, is one of its greatest strengths: “Perhaps the cognitive power most universally held to be provided by connectionist networks is the ability to become organized in task-relevant ways just by performing tasks, without requiring programming by an intelligent agent” (Newell, 1990, p. 485). Though symbolicist systems like Soar can behave in a similar way, they do not yet do so continuously - that aspect of cognition has been “put off to [Soar's] future agenda” (Newell, 1990, p. 488). Natural cognizers seem to self-organize continuously (Thelen and Smith, 1994, p. 203), so it is a distinct advantage for connectionist models to be able to easily incorporate continuous organization. As well, the dynamicist can rightfully claim that this ability is naturally within the scope of dynamicism for a connectionist network *is* a dynamical system.

In general, connectionists and dynamicists share a similar view of cognitive system architecture. It is coupled to the content of the system, and it can change continuously with time. However, connectionist systems are relatively high-dimensional systems (*i.e.* of many variables) as compared with a typical dynamicist system. Because connectionist architectures have a basic unit, the final system can be broken down into interactions between these various fundamental components - this is not the case for the dynamicist. Rather, “non-connectionists, 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" (van Gelder and Port, in press, p. 27).

Of course, the dynamicist will, hopefully, have a lower-dimensional set of equations describing the same phenomenon. However, important sacrifices are made for this simplicity. Only behavioral comparisons can be used to validate a model which has had parameters "determined by hand," (van Gelder and Port, in press, p. 27) there are no direct structural comparisons which can be made as with a connectionist model. Furthermore, it is extremely difficult to know how to set the system's parameters. Even if someone provided a dynamicist with the right set of equations, there is little to guide the choice of parameter values. Unlike the connectionist networks which can be trained from real data, and compared directly to a real network, the dynamicist must rely on intuition - a capricious faculty at best (see section 3.1.2 for further discussion).

2.3.4 Computation

Connectionists subscribe to the view that what a neural network does is best understood as "computing with nets" (Churchland and Sejnowski, 1992, p. 76). Thus, it seems straightforward to classify connectionists as holding a computational view of cognition; a view which seems to offend the dynamicist: "almost all computational approaches attempt to superimpose on this multiple, simultaneous, interactive behavior a sequential, step-by-step structure" (van Gelder and Port, in press, p. 20). However, this criticism is clearly not true of connectionists.

The problem is, of course, one of definition: What does it mean to be computational? Van Gelder and Port seem to use the terms *computational* and *symbolicist* interchangeably. However, Churchland and Sejnowski often refer to nervous systems and connectionist systems as *computational* simply because they have a different set of criteria for inclusion into the class of *computers*.

While van Gelder and Port simply assert: "Natural cognitive systems aren't computers" (van Gelder and Port, in press, p. 1), Churchland and Sejnowski claim that: "once we understand more about what sort of computers *nervous systems* are, and how they do whatever it is they do, we shall have an enlarged and deeper understanding of what it is to compute and represent" (Churchland and Sejnowski, 1992, p. 61). It seems that the dynamicists are using the terms *computer* and *computational* to refer solely to a *serial digital computer* where the connectionists are using the terms to refer to "a physical device with physical states and causal interactions resulting in transitions between those states" (Churchland and Sejnowski, 1992, p. 66). It is clear, then, why there is confusion on the use of the term "computer". Notably though, there seems to be agreement that nervous systems are not much akin to silicon computers "what is being modeled by a computer is itself a kind of computer, albeit one quite unlike the serial, digital machines on which computer science cut its teeth" (Churchland and Sejnowski, 1992, p. 7).

However, the dynamicists do attempt to lump connectionists into the same computational category as symbolicists, in order to have them fall prey to the same theoretical difficulties: "Much standard connectionist work is just a variation on computationalism, substituting activation patterns for symbols" (van Gelder and Port, in press, p. 2). However, this seems a rather naive statement on the part of the dynamicists. There is seldom an attempt to simply "substitute" patterns of activation for symbols. The type of computations being done are so different that such a substitution would defeat the purpose of using connectionist networks in the first place (Smolensky, 1988). Furthermore, symbolicist symbols and connectionist patterns of activation behave very differently. Though both are meant to represent something, activation patterns do so dynamically, being updated through a continuous variation of vectors in activation space, whereas symbols are more static and are generally replaced by a new symbol - there is not the same sense of continuity (Newell, 1990; Churchland and Sejnowski, 1992).

For van Gelder and Port to claim that such "substituting" is "standard" practice in the connectionist community also seems curious. The Churchland and Sejnowski (1992) book never once refers to such practices. Furthermore, Newell never even hints at such a devastating claim, on the contrary he notes: "another cognitive power (of connectionist systems) that seems significant is finding the solution to very large systems of weak or soft constraints" (Newell, 1990, p. 485). A power which is unique to connectionism, due to the difference in their approaches to handling representation and computation.

This attempt at discrediting the computational perspective of connectionism seems misguided and somewhat half-hearted. The only similar attack presented by the dynamicists against a computational view of cognition is the conclusion that: "Whatever computational properties a system might or might not have are basically irrelevant" (van Gelder and Port, in press, p. 7). Van

Gelder and Port arrive at this conclusion by claiming it is the dynamical properties of a system that count, not the computational ones. However, it seems evident that if a certain type of computation can support the appropriate dynamical properties, whereas another sort of computation cannot, computational properties are not irrelevant but can be used to choose the more appropriate implementation of the dynamics of natural cognizers.

2.3.5 Representation

There have been two main sorts of representation used in connectionist networks since their conception; local coding and distributed representation. However, in recent years local coding has fallen out of favor with those connectionists who wish to model more neurologically realistic networks. Distributed representation has been known under a number of different names. Churchland and Sejnowski refer to distributed representation as “vector coding”[18] but it has also been dubbed “state space representing” and “multidimensional representing” (Churchland and Sejnowski, 1992, p. 163).

Simply put, a connectionist network represents the state of the system at any particular time by the combined activity levels of all units in the network. Each of these units is normally interpreted as a state variable (and hence an axis in the state space), and thus there are as many dimensions to the state space as there are units in the network[19]. The units are connected to one another in various ways, most commonly with an excitatory link of some kind. Each unit generates its output based on the levels of input received from any units connected to it. After integrating the inputs the unit passes the result through a transfer function. Most often this function will be nonlinear, but there are no theoretical constraints on possible functions, as long as they are defined over the range of inputs. Such a system is nonlinear, differentially describable, complex, dynamical and self-organizing.

---

[18]Churchland and Sejnowski also cite evidence that vector averaging is used as a form of representation by nervous systems (Churchland and Sejnowski, 1992, p. 234). However, both vector averaging and vector coding are distributed forms of representation, only the encoding process differs.

[19]Under this interpretation, the brain would have a dimensionality of approximately one hundred million. So, even a connectionist model of a small area of the brain will be a very high dimensional model.

As an example, a typical evolution equation of a nonlinear connectionist network is:

$$\frac{dy_i}{dt} = -\tau_i y_i(t) \frac{1}{1+e^{-\left(\sum_j w_{ij} y_j + I_i(t) + \theta_i\right)}} \qquad (1)^{20}$$

Equation (1) shows that the change in activity level of any given unit (y_i) in the network depends on its own activity, and a function of the sum of its current input from other units, input from outside the network (the environment) and a bias, which acts to improve settling of the system. The dimensionality of the system is determined by the number of such equations (*i.e.* the maximum value of *i*).

Connectionist processing is the transformation of the resulting patterns of activity across y_i and can be thought of as a path (or trajectory) traced through a high-order (*i.e.* of size *i*) state space. Hence, the behavior of the network is *exactly* described by the system's trajectory through its state space; as in any typical dynamical system. In other words, the tools provided by dynamical systems theory are directly applicable to a description of the representational processing of connectionist networks. Clearly, these networks have a concrete representation of their environment which is captured by the vectors coded in the nodes of the network. These vectors are high-dimensional and quite possibly difficult to translate into a set of simple characteristics of what is being represented (Churchland and Sejnowski, 1992). Nonetheless they *are* representations.

Connectionist representations are clearly rather different from the physical symbol representations advocated by Newell. For connectionists, a concept, idea, experience, object, *etc*. is coded by a large population of neurons; it is *not* captured by *a* node or a symbol: “a configuration of weights embodies the knowledge needed to answer a variety of different questions and stored knowledge contributes to occurent representations” (Churchland and Sejnowski, 1992, p. 168). This manner of representation seems to be extremely effective for categorization and “best-fit” problems. However, there is continuing difficulty in representing relations between concepts rather than simply the concepts themselves (Bechtel, in press). However, symbolicists do not have the same difficulty in capturing such relations, as they simply define them.

---

[20]As in van Gelder (1993).

One of the defining papers on connectionist distributed representations is Smolensky's (1988) *On the Proper Treatment of Connectionism*. Smolensky explicitly claims that connectionism provides a dynamical systems alternative to symbolicism (see the *connectionist hypothesis*, section 2.3.1). He notes that connectionist nets are inherently coupled, nonlinear, parallel dynamical systems and that these systems are self-organizing and evolve based on continuously varying input from their environment. Still, the dynamicists argue that connectionist networks are limited in ways that a dynamicist description is not.

The fact that connectionists admit that they represent *at all* is reason enough for van Gelder (unpublished) to reject their position. To be a truly dynamical system, there *should be no representation* in the cognitive model. And clearly, connectionists represent concepts in their networks[21]. Dynamicists understand the connectionist attempt to abstractly represent the world as a different guise of the symbolicist's misguided commitment to symbolic representation. Van Gelder concludes: “it is the concept of *representation* which is insufficiently sophisticated” (van Gelder, 1993, p. 6) for understanding cognition. So, any cognitive approach which incorporates representations into its models is doomed to failure. Thus, only dynamical models are capable of demonstrating how cognitive systems defy description in representational terms.

Here lies the difference between the connectionist and dynamicist views of representation. For the connectionist: “To a first approximation, state spaces in nervous systems are only abstractly, not literally, in register” (Churchland and Sejnowski, 1992, p. 337). But for the dynamicist the state space *is*, if anything can be said to be, what captures (or represents, though not in any usual sense of the term) the agent's relationship with its environment. This view, however, raises the question that was briefly considered in section 2.2.5: It seems that we, as natural cognitive agents, manipulate representations; how is this possible if we do not represent things to ourselves?

There has been moderate success dynamically modeling natural cognizers without any use of representation - but only cognizers with the sophistication of cockroaches (Brooks, 1991). However, it is unknown whether such systems will scale up to be able to handle problems which are cast in representational terms.

---

[21]It is generally agreed that both connectionists and symbolicists are representationalists (Fodor and Pylyshyn, 1988, p. 7).

2.4 Summary

Dynamicists believe that natural cognitive systems are certain kinds of dynamical systems, and that cognition is best understood from the perspective of dynamical systems theory. They have identified what they feel should be the reigning paradigm in cognitive science, and have a mandate to prove that the dynamicist conception of cognition is the correct one to the exclusion of symbolicism and connectionism.

Symbolicism is committed to a view of mind as a type of computer. The cognitive computer uses symbolic representation in its processing of information concerning its environment. The total cognitive system is broken into three components, central cognition, perception, and motor control, with central cognition being the most important aspect of the total cognitive system.

Dynamicists reject the computational hypothesis on the basis of its difficulty accounting for time, embeddedness, and the intimate relationship between all parts of the total cognitive system. For these reasons, dynamicists believe their approach will give rise to fundamentally superior models of cognition. Biological evidence and the symbolicists' practical difficulties lend support to many of the dynamicists criticisms. However, it is unclear how dynamicists will account for the representational aspects of cognition.

In contrast, differentiating between connectionist networks and dynamical systems models is no easy task. Frequently, dynamicists realize that connectionist networks are often “continuous nonlinear dynamical systems” (van Gelder and Port, in press, p. 26). Smolensky, as early as 1988, outlined the many ways in which a connectionist network is a dynamical system. What he encapsulated was the essence of dynamical systems in their relation to cognition and connectionism. Churchland and Sejnowski have gone further, discussing limit cycles, state spaces, and many other dynamical properties of nervous systems and have even included purely dynamical analyses (Churchland and Sejnowski, 1992, p. 396) in their connectionist discussions of natural cognitive systems.

Nevertheless, dynamicists wish to subordinate and subsume connectionism under their cognitive approach (van Gelder and Port, in press, p. 27). Dynamicists have rejected the connectionist commitment to computationalism, representationalism, and high dimensional dynamical descriptions.

Chapter 3

Dynamical Challenges

On the basis of the discussion of the preceding chapter, we can now examine dynamicism in its cognitive science context. Applying the analysis of models, metaphors, analogies and theories from chapter one to the current state of dynamicism, we are now able to determine its success as a scientific theory.

After briefly outlining the most valuable contributions of dynamicism to the study of cognition, we will examine some of the more pressing difficulties with accepting the dynamicist insistence that dynamicism is a new paradigm for the study of mind. Subsequently, we will be able determine an appropriate place for dynamical systems theory in the context of existing cognitive theories.

3.1 Strengths of Dynamicism

The most compelling characteristic of dynamicist models is their inherent ability to effectively model complex temporal behavior. It is unanimously agreed that the temporal features of natural cognizers must be adequately accounted for in a good cognitive model. Not only does dynamical systems theory address the temporal aspect of cognition, it makes this aspect *the most important*. The reasons for choosing to do this are obvious to the dynamicists: we exist in time; we act in time; and we cognize in time - real time. Therefore, the dynamical systems equations, which have been applied successfully in other fields to predict complex temporal behaviors, should be applied to the complex temporal behavior of cognitive agents. Whether or not we choose to subscribe to dynamicist commitment to a particular type of model, they convincingly argue that we cannot remove our models of cognition from time and still expect good models. Natural cognition indeed seems to be inherently dynamical in nature.

Furthermore, a cognitive model must not only model a particular level of temporal behavior, but all ranges of temporal behavior, from the neurological milliseconds to the cognitive seconds and beyond (see sections 2.0 and 2.2.2). Both short term, and long term behaviors can be effectively captured by systems of coupled nonlinear differential equations. Again, this inherent property of the tools espoused by dynamicists gives them reason to claim their approach is better suited to modeling the mind.

The ability of dynamical systems to account for multiple levels of behavior should not be underestimated. It is of great value for cognitive scientists to be able to examine both global and local behaviors with the concepts provided by a single approach. Dynamicists are able to inherently employ concepts such as state space, trajectory, minima and attractors in describing the overall behavior of a dynamical system, while simultaneously being able to examine the behaviors of small portions of the system. The advantages of this ability has been demonstrated through its application to other areas of science such as physics, biology, and much of engineering. The simple fact that there is this wide range of fields to which dynamical systems theory has been successfully applied lends it scientific credibility (van Gelder and Port, in press, p. 14). Furthermore, the fact that dynamicism is presented as a *unified* theory of cognition, with no obvious reasons not to consider it as such, makes it a candidate for a *good* theory (see section 1.1.3).

Again stemming directly from the dynamicist commitment to a particular type of mathematical model, they are able to naturally account for behavioral continuity. Though the question of whether or not all of behavior is continuous or discrete is a matter of great debate among psychologists (Miller, 1988; Molenaar, 1990), dynamical systems models possess the ability to describe either instance. So, relying on the assumption that behavior is "pervaded by *both* continuities and discrete transitions" (van Gelder and Port, in press, p. 14) as seems reasonable (Schweicker and Boggs, 1984; Luck and Hillyard, 1990; Egeth and Dagenbach, 1991; Churchland and Sejnowski, 1992), dynamicism is in a very strong position to provide good cognitive models based on its architectural commitments. Similarly, there is no embeddedness problem for the dynamicists because of their commitment to a homogeneous cognitive architecture.

Finally, the fact that dynamical systems theory provides a novel set of metaphors for thinking about cognition is paramount. Black's contention that science must start with metaphor (see section 1.1.1) makes us realize the importance of addressing new metaphors. These metaphors may provide us with a perspective on cognition that is instrumental in understanding some of the problems of cognitive science. Though we must be careful not to rely solely on metaphors for explanation (see section 3.1.2) they can be important conceptual aids.

At the foundation of all these strengths (except perhaps the last), lies the ability of dynamical systems theory to handle time. Dynamicists have unceasingly pointed to their temporal commitment as the most important (van Gelder and Port, in press, p. 14). Unfortunately, it is not clear that dynamicists have a monopoly on temporal models. If they do not, it will be difficult to distinguish dynamicism from its competitors to the extent that dynamicism should properly be referred to as a new paradigm.

3.2 Difficulties with Dynamicism

3.2.1 Introduction

Before launching into a discussion of the difficulties dynamical systems models face in their application to cognition, it is useful to briefly examine an example of such a model. The olfactory bulb model by Skarda and Freeman is one of the few well-developed models that dynamicists refer to as an example. Many authors, including van Gelder, Globus, Barton, and Newman have cited this work as strong evidence for the value of dynamical systems modeling of cognition. Upon further examination, it becomes clear that this model is subject to some theoretical difficulties, which are elucidated and expanded in sections 3.2.2-3.2.5, and it is not even obvious that this dynamicist exemplar is indeed a dynamicist model.

In Skarda and Freeman's (1987) article *How brains make chaos in order to make sense of the world*, a dynamical model for the olfactory bulb in rabbits was outlined and tested to some degree. They advanced a detailed model of the neural processing underlying the ability to smell. This model relies on a complex dynamical system description which may alternate between chaotic activity and more orderly trajectories corresponding to a learning phase or a specific scent respectively. They hypothesized that chaotic neural activity serves as an essential ground state for the neural perceptual apparatus. They concluded that there is evidence of the existence of important sensory information in the spatial dimension of EEG (electroencephalogram) activity and thus there is a need for new physiological metaphors and techniques of analysis.

Skarda and Freeman observed that their model generated output that was statistically indistinguishable from the background EEGs of resting animals. This output was achieved by setting a number of feedback gains and distributed delays "in accordance with our understanding of the anatomy and physiology of the larger system" (Skarda and Freeman, 1987, p. 166) in the set of differential equations that had been chosen to model the olfactory bulb. It is important to note that the behavior of the system can be greatly affected by the choice of certain parameters, especially if the system is potentially chaotic. Such difficulties in choosing system's parameters are faced by connectionist modelers as well, though it can be argued that these difficulties may more greatly affect dynamicist models (see section 3.2.4).

A further criticism of this model can be directed at its predictive or correlative properties. Although the model accounts quite well for a number of observed properties, "it does not correspond with the actual EEG patterns in the olfactory lobe" (Barton, 1994, p. 10). The consequences of this inaccuracy seem quite severe. For, if both the model, and what is *being* modeled are indeed chaotic systems (*i.e.* very sensitive to initial conditions), but they are not the *same* chaotic system and if there is any inaccuracies in their initial conditions[22], then the divergence of the state spaces of the model and the real system will be enormous within a short time frame. Consequently, the model will not be robust and will be difficult to use. Again, this problem is not a difficulty unique to dynamical systems models, rather it is related to chaotic systems in general, but such concerns must be addressed explicitly by dynamicist models.

Most tellingly, the authors themselves saw their paper and model as showing that "the brain may indeed use computational mechanisms like those found in connectionist models" (Skarda and Freeman, 1987, p. 161). Furthermore, they realized that: "Our model supports the line of research pursued by proponents of connectionist or parallel distributed processing (PDP) models in cognitive science" (Skarda and Freeman, 1987, p. 170). Dynamicists, however, wish to rest a new cognitive paradigm on such a model. Ironically, the model itself looks very much like a connectionist network, with the less typical addition of inhibition and far more complex transfer functions at each node. These facts make it rather curious that it is touted as a paradigmatically important dynamical

---

[22]Which there are theoretically guaranteed to be given that the systems are chaotic (Gleick, 1987).

systems model. The model's similarities with connectionism make it quite difficult to accept the assertion that this type of dynamical model is the seed of a new paradigm in cognitive modeling.

Establishing dynamical systems modeling as a new paradigm rests partially on its establishing a unique place in cognitive modeling, but more importantly in proving that it simply provides good models. However, rather than focusing solely on the shortcomings of specific models, the next sections will highlight some of the more important theoretical problems with a purely dynamical description of cognition.

3.2.2 Representation

Dynamicists have distinguished themselves from both symbolicists and connectionists with the assertion that representation is *not* necessary to adequately explain cognition. Of course, such a position is reminiscent of the unsuccessful behaviorist project.

In the late 1950s, there was extensive debate over the behaviorist contention that representation had no place in understanding cognition. One of the best known refutations of this position was given by Chomsky in his 1959 review of B. F. Skinner's book *Verbal Behavior*. Subsequently, behaviorism fell out of favor as it was further shown that the behaviorist approach was inadequate for explicating even basic animal learning (Thagard, 1992, p. 231). The reasons for the behaviorist failure was its fundamental rejection of representation in natural cognizers.

Dynamicists have forwarded a similar rejection of representation as important to cognition (see sections 2.1.3, 2.2.5, 2.3.5). Consequently, they fall prey to the same criticism forwarded over three decades ago. The early work of researchers like Johnson-Laird, Miller, Simon and Newell firmly established a general commitment to representation in cognitive science inquiries (Thagard, 1992, p. 233). There have been no reasons given by dynamicists which would fundamentally disturb this commitment.

Finally, it seems quite clear that the average human cognizer is continually manipulating representations. Everyday, human interaction involves speech, reading, writing (*i.e.* language) and possibly counting, yet the dynamicists claim that we do not use representations in our thinking. It is not clear how dynamicists would have such behavior reproduced in a non-representational system. Though dynamicists can remind us of the impressive behaviors exhibited by Brooks' (1991) dynamical robots, it is uncertain that insect-like reactions of these sorts will scale to the complex interactions of mammalian cognition.

3.2.3 Dynamical Systems Metaphors

The intuitive appeal of a dynamical systems theory description of many systems' behaviors is quite difficult to resist. It simply makes sense to think of systems as being *attracted* to certain states (*e.g.* some people seem to be disposed to being happy). However, can such metaphorical descriptions of complex systems actually provide us with new insights, integrate previously unrelated facts, or in some other way lead to a deeper understanding of these systems? In other words, can dynamical descriptions be more than metaphorical in nature?

In order to answer this question in the affirmative, we must be able to show the potential for new predictions and explanations (see section 1.1.1). Even though "dynamical concepts and theory are seductive, we may mistake translation for explanation" (Robertson, Cohen et al., 1993, p. 119). We cannot allow ourselves to accept new concepts which do not deepen our understanding of the system being modeled.

Philosopher of science Mary Hesse has noted that theoretical models often rely on a sort of analogy to the already familiar (1988, p. 356):

> [Theoretical models] provide explanation in terms of something already familiar and intelligible. This is true of all attempts to reduce relatively obscure phenomena to more familiar mechanisms or to picturable non mechanical systems...Basically, the theoretical model exploits some other system (such as a mechanism or a familiar mathematical or empirical theory from another domain) that is already well known and understood in order to explain the less well-established system under investigation.

Clearly, this tack is the one that dynamicists have taken. They are attempting to address the obscure and poorly understood phenomena of cognition in terms of the more familiar mathematical theory of dynamical systems, which has been successfully applied to complex mechanical and general mathematical systems.

An excellent example of this sort of analogical application of dynamical systems theory and concepts can be found in Abraham, Abraham, and Shaw (1994). In an attempt to better understand human behavior, these clinical psychologists have applied dynamical modeling techniques and concepts to the problems they encounter. This application is another attempt to deal with the general problem of "relating patterns of activity over time - as many psychological states surely are - to instantaneous states" (Morton, 1988, p. 141). In effect, psychologists wish to understand what effect a person's experience has on their current mental condition (Abraham, Abraham, and Shaw, 1994).

Often, clinical psychologists applying dynamics to their field ignore the differences between their field and the rigorous ones from which dynamical systems theory arose: "One way that the distinction between fields is set aside is when authors use rigorous terminology from nonlinear dynamics to refer to psychological variables that are multidimensional and difficult to quantify" (Barton, 1994, p. 12). For example, some psychologists have equated the dynamical concept of chaos with overwhelming anxiety, others with creativity, and still others with destructiveness (Barton, 1994). These applications of chaos are clearly more metaphorical than rigorous, and bear little resemblance to the definitions used in precise dynamical systems theory models.

In particular, Abraham, Abraham, and Shaw's (1994) dynamical descriptions of behavior apply the concepts of dynamical systems theory to a Jungian analysis of human behavior. However, applying these concepts in such a metaphorical manner simply seems to relate the phenomena in a new way. There is no rigor added to their model simply because the chosen metaphor is mathematical. Barton duly notes that in the paper describing one such dynamical model of a Jungian hypothesis, Abraham et al. "imply a level of measurement precision we don't have in clinical psychology" (Barton, 1994, p. 12).

There is no *explanation* in these clinical psychological applications of dynamics systems theory to the phenomenology or intentionality of cognition. These "models" are simply metaphorical descriptions, they advance no new insights in clinical psychology. They do not reveal any details about what is being modeled. There are no consistent and explicit mappings between dynamical systems theory and human behavior. We have clearly not been presented with anything resembling Beardsley's desirable "controlled metaphor" (see section 1.1.2).

From the stand point of cognitive models, there is not a lot of value in such descriptions. We cannot generate a rigorous model, nor produce computational simulations from metaphor, so we are not able to discover if the models are predictive. This is a serious failure for any scientific model (see section 1.2.1, 1.1.3). Of course, it is possible to haphazardly generate a model which produces data that seems appropriate, but since we have no explicit map between the concepts of clinical psychology and those of dynamical systems theory, the data is meaningful only in its mathematical context, not in a cognitive one.

Some dynamicists have presented us with more precise models. Robertson et al. (1990), outlined a model for CM (cyclicity in spontaneous motor activity in the human neonate) using a dynamical approach. It seems that such quantifiable physiological behavior should lend itself more readily to a dynamical description than perhaps clinical psychology would, allowing the psychophysiologist to avoid the poor conceptual mappings of clinical psychologists.

Robertson et al., after "filtering" the observed state space, obtained a dynamicist model with desirably few degrees of freedom which seemed to be able to model the stochastic process of CM. However, upon further investigation, the only conclusions that could be drawn were: "We clearly know very little about the biological substrate of CM" (Robertson, Cohen et al., 1993, p. 147). In the end, there is no completed dynamicist model presented, though models which do not work are discounted. So, they have employed dynamicist models to constrain the solution, but not to provide new insights. In their closing remarks, they note: "We are therefore a long way from the goal of building a dynamical model of CM in which the state variables and parameters have a clear correspondence with psychobiological and environmental factors" (Robertson, Cohen et al., 1993, p. 147). In other words, a truly dynamicist model is still a future consideration.

Even in the most rigorous of dynamical models, such as the Skarda and Freeman model previously discussed, extending dynamical concepts beyond the metaphorical proves difficult: "Given this broad picture of the dynamics of this neural system we can sketch a metaphorical picture of its multiple stable states in terms of a phase portrait" (Skarda and Freeman, 1987, p. 166). Furthermore, "in its psychological dimension our model is extremely limited, being competent to simulate only preattentive cognition" (Skarda and Freeman, 1987, p. 170).

The concepts of dynamical systems theory provide an interesting method of thinking about cognitive systems, but they have not yet been shown to be successfully transferable to rigorous definitions of human behavior or cognition. The "haziness" of clinical psychology does not allow for quantification of mechanisms in dynamical systems theory terms. Furthermore, even some physiological processes do not seem to lend themselves to precise quantitative dynamicist descriptions that are able to provide the predictive or explanative powers expected of good models (see section 1.2.1).

3.2.4 Dimensionality and Parameter Estimation

One reason that it has been so difficult for dynamicists to provide predictive models is that they have been unable to meet the challenge of identifying and quantifying the parameters sufficiently for a dynamicist model. It is extremely difficult, if not impossible, to simply examine a system and pick out the appropriate parameters to be used in a model. This discontinuity between model and reality is a common problem in investigating natural nonlinear systems: "Not only are investigators rarely able to completely characterize all the variables that affect a complex system, but they must isolate a system well enough to cut through what Morrison (1991) called a "sea of noise"" (Barton, 1994, p. 10).

When it is difficult to know the factors involved and to find the signal of interest in ambient noise, as is often the case, it is common practice for dynamicists to define *lumped parameters*. Barton attributes this practice to a confusion of techniques between levels of analysis. Whatever the reason, the result is that sundry real-world variables are represented by only *one* lumped model variable. Consequently, the meaning of this new lumped variable becomes very difficult to discern. If the meaning of a variable cannot be determined, it becomes next to impossible to test a model, or to verify hypotheses derived from observing the behavior of the model.

However, dynamicists wish 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 (the brain)"(van Gelder and Port, in press, p. 28). The only feasible way to generate low-dimensional models of admittedly high-dimensional systems is to used lumped parameters. So it seems we must insist that dynamicists reconsider their criteria of accepting only low-dimensional models as being valid models of cognition.

Perhaps the best solution to this problem is to not use lumped parameter models. Unfortunately, this solution gives way to an important new problem. If we are attempting to describe cognition with a number of parameters, let us say *n* of them, then we have an *n*-tuple that we can use (we hope) to completely characterize the behavior of our system. Of course, these *n* parameters are contained in coupled, nonlinear, differential equations. Let us think for a moment about the complexity of the system we are attempting to model; the human brain.

The number of parameters *affecting* this system seems almost infinite. Remember, we must account for not only the system itself, but all environmental factors. The brain contains approximately one trillion connections (Pinel, 1993). The environment, which must be coupled to our brain consists not only of other provably chaotic systems like weather, ocean currents, and species populations but billions of other brains, let alone the artificial systems we interact with every day, or the planets, moons, stars, *etc*. It could be claimed that expecting the dynamicist to characterize such a system is quite unreasonable. Perhaps, so we will, for argument sake, assume the number of parameters is clearly large, but also finite.

So, we have an *n*-tuple of parameters for the equations describing our complex system, where *n* is some large finite number. Is a coupled nonlinear differential equation description of human behavior of this size of value? Using even the most advanced numerical methods, and the most powerful computers, such a problem would probably be unsolvable. So, let us assume infinite computing power. Let us assume that we can solve our system of equations. The next important question that arises is: What can we use for initial conditions?

Is it feasible for us to be able to measure *n* starting conditions for our model, with any kind of precision? Let's assume once more, that we have *n* initial conditions of sufficient accuracy. What kind of answer can we expect? Of course, we will get an *n* dimensional trajectory through an *n* dimensional state space. How can we possibly interpret such output? It becomes more and more difficult to continue justifying further assumptions. Thus, at this point, it seems that an interpretation of such a trajectory becomes, if not impossible, meaningless. There is absolutely no way to uncouple the parameters, or find out exactly what it means for the system to move through the state space.

By adopting a purely dynamicist approach, it becomes impossible to identify the underlying mechanisms that affect behavior. With connectionism there is always a reasonably simple unit (the neurons or nodes) to which behavior can ultimately be referred. With symbolicism there are the fundamental symbols to which we can appeal. In both of these instances, understanding global behavior is achieved through small steps, modeling progressively more complex behavior and allowing a *backtrace* when necessary to explain a behavior. With dynamical equations, on the other hand, no such progression can be made. The model is general to such an extent as to lose its ability to explain from where the behaviors it is producing are coming.

3.2.5 Summary and Discussion

Unfortunately, many of the major shortcomings of the dynamical systems approach become clear in an attempt to move from the metaphorical explanation to an empirical one (see section 3.2.3). The limited practical success of purely dynamicist models sheds some doubt on the plausibility of rigorous dynamicist cognitive models. The difficulties of translating compelling metaphorical descriptions into computational models stem from a number of characteristics of the "certain kind" of dynamical systems that dynamicists, like van Gelder, wish to forward as a new cognitive paradigm (see section 2.1.1).

Having restricted dynamicist descriptions to systems of differential equations which ideally rely on lumped parameters, (see sections 3.2.4) the dynamicist has created serious problems in attempting to apply these models to cognition. The first of these is the difficulty of formulating the system of equations meant to describe a cognitive behavior. There are so few constraints available to the modeler that thousands of different systems could be created to model the same phenomena.

It is extremely difficult to choose the best of these systems. Often, the best that can be hoped for is a ruling out of certain models, but the number of models left often makes it difficult to draw definite conclusions (see section 3.2.3). The reason lies in the impossibility of providing good parameter estimates and the sensitivity of many of these systems to initial conditions. These difficulties, combined with the further difficulty of deriving a complete system of equations, makes for such an unconstrained problem that knowing when one has a good dynamicist model is extremely difficult.

Even in the instance when one has a seemingly good model of the brain it is, by definition, lacking. For, if dynamicists claim that a cognitive system must be embedded to the extent that the environment is as essential to modeling behavior as the rest of the system (see section 2.1.1), then any adequate dynamicist model would have to equally emphasize the environmental and cephalic aspects of the model. In other words, if dynamicists do not define system boundaries for their cognitive models, they have magnified the problems involved in modeling cognition immeasurably - *nothing* can be neglected. Even a model which is successfully reproducing an aspect of cognitive behavior is non paradigmatic if it does not account for the dynamics of the environment. Of course, it is compelling to claim that the environment plays an important role in cognition, but it is far less compelling to insist that researchers are unable to delineate reasonable systems boundaries which will allow for effective cognitive modeling.

A necessary result of rejecting such boundaries is the dynamicist commitment to nonrepresentational models. Without clear boundaries between the environment and the cognizer, there is no reason to posit any form of representation taking place internal to the cognizer. Though dynamicists present interesting instances when it seems representation may be inappropriate (*e.g.* motor control, habitual behavior, *etc*.) it is difficult to understand how dynamicists intend to explain the ubiquitous use of representation by human cognizers while maintaining a complete rejection of representation. This project failed with the behaviorists, and it is not clear why it should succeed now.

Furthermore, the dynamicist rejection of representation in dynamical systems models makes it extremely difficult to justify or even offer an interpretation of the behavior of the system. Additionally, the dynamicist emphasis on lumped parameter models makes interpretation of a system's behavior even more confounding. There does not seem to be any evident method for determining the meaning of a particular parameter in a model. For example, if we observe oscillator-like behavior in the movement of the limbs of human infants, we are likely to choose a dynamical model that we know has oscillator-like behavior. However, the parameters in the model do not seem to have a biological meaning, they merely exhibit the same end behavior as the system being modeled.

Such difficulties of interpretation arise for dynamicists in general because once a dynamical behavior is evident in a system, the differential equations known to exhibit a similar type of behavior are generated and the parameters are adjusted so the model seems to be approximating the observed behavior, but it is unclear how these parameters and equations relate to what they are modeling. Both equations and parameters have been chosen on the basis of known mathematical behavior, not for their biological import. In other words, the dynamicist has identified a theoretical behavioral similarity in two systems, used one to approximate the other, but is then unable to re-examine the initial conceptual mapping in a more realistic context, to provide an explanation.

In general it seems that dynamicists have a difficult time keeping arbitrariness from permeating their models. There are no clear ways of justifying parameter settings, choosing equations, interpreting data, or creating system boundaries. Despite the power and intuitive appeal of dynamical systems theory, application of this field of mathematics to cognitive modeling is neither trivial nor obviously preferable to connectionism and symbolicism, as dynamicists would have us believe.

3.3 The Place of Dynamicism

3.3.1 Dynamicism as a Cognitive Theory

We have now laid the foundations[23] for a discussion of the dynamicist approach in terms of its general theoretical viability and particularly in its relation to cognitive science. As will be recalled from chapter 1, the important characteristics of a good cognitive theory include its ability to stimulate research, its theoretical unity, its accuracy of prediction and explanation, its ability to elicit application through good models, and finally, being able to motivate computational simulation. These criteria will be used to motivate a comprehensive evaluation of the usefulness of the dynamicist approach.

To begin with, the ability of dynamical systems theory to stimulate cognitive research is undeniable. There have been a number of new books written with the dynamicist approach in mind (Robertson, Cohen et al., 1993; Thelen and

[23]In particular, sections 3.1 and 3.2 and subsections 1.1.3, 1.1.4 and 1.2.1.

Smith, 1994; van Gelder and Port, in press); the increase in dynamicist contributions to cognitive science literature is indisputable. Furthermore, connectionism has been strongly influenced by dynamical systems theory. It is difficult to tell whether the tendency to emphasize dynamical connectionist networks would be as strong as it is without the formulation of the dynamicist approach to cognition. In any case, the theory has forced cognitive scientists to examine the extent to which dynamics must be addressed in cognitive models and to attempt to delineate the certain classes of dynamical models which are most applicable to cognition. In this way, dynamicists have effectively challenged traditional stances on some core issues of cognitive science as outlined in chapter 2.

In doing so, they have posed a theory which is comprehensive and unified. Dynamicists propose to explain all of cognition - a cognitive theory cannot claim to be more unified. Though dynamicists to date have often limited the scope of their models to locomotor issues, their objective is unequivocally to explain all aspects of cognition (Thelen and Smith, 1994). Clearly, the dynamical systems approach to cognition is intended as a unified theory, one which is expected to rival the existing connectionist and symbolicist theories of cognitive science (van Gelder and Port, in press).

In an attempt to displace the existing paradigms, dynamicists have posed provocative questions and leveled some valid criticism of foundational commitments of these theories. However, their positive contribution to the field is more difficult to laud. Explanative and predictive powers have not been evident so far in the few examples of dynamicist models that are available for examination. It could be argued that an explanation of cognition can be offered through dynamicism, but it would be difficult to accept this as a *valid* explanation if there is no empirical evidence to back up the theory. If dynamicism is to be a non-empirical cognitive theory, it will encounter the same difficulties as folk psychology (Churchland, 1989) and cannot reasonably be considered in the same category as symbolicism and connectionism. However, this is clearly not the aim of dynamicists.

Clearly, lack of empirical exemplars in such a relatively new approach is no reason to dismiss it. However, it was also demonstrated that predictive and explanative difficulties do not seem to arise simply because of a deficiency of dynamicist exemplars (see sections 3.2.4 and 3.2.5). The limited predictive success of dynamical systems models and the theoretical difficulties with prediction pose a great threat to the survival of dynamicism. Any scientific theory simply must provide explanation and prediction (see section 1.1.3 and 1.1.4). Without these vital aspects, it is impossible for any scientific movement to establish itself as a new paradigm.

In the case of theories of cognitive science, much of this predictive power stems directly from the ability of a theory to spawn good models (see section 1.2.1). As has been discussed, dynamicism has not had a history of positing good models. Of the few good models which dynamicists claims as their own, some, like the Skarda and Freeman (1987) model, are classified by the modelers themselves as a different approach - most often connectionism. Without strictly dynamicist models, quantitative data is simply not available to support the theoretical commitments of dynamicists. Without supporting data, it is difficult to accept a theory, especially one which wishes to supplant those *with* such supporting empirical evidence.

Difficulties in the mathematical formulation of a model to generate the appropriate behavior and thus data are at least threefold. First, the equations chosen to mimic the behavior of the system being modeled may be only superficially mappable to the target system. Second, attempting to posit low dimensional models causes the parameters to become lumped parameters and hence difficult to precisely interpret. Lastly, having to place equal emphasis on environmental and internal dynamics seems to expand the scope of the problem to the extent that it becomes unmanageable or oversimplified (see sections 3.2.4 and 3.2.5).

Even having formulated a mathematical model, a dynamicist may run into severe implementation difficulties (see section 3.2.5). Of course, in general it is not prudent to reject a theory because resulting models cannot be simulated on available technology. However, in cognitive science it is questionable how useful such theories are. The importance of computational simulations of models is paramount in establishing their cognitive viability (see section 1.2.1). Thus, presenting intractable, non-interpretable, exceedingly complex, difficult or impossible to simulate models will circumscribe the development of a cognitive approach. This is especially true if the construction of the model is suspect in the first place. A model and its simulation play a major role in validating a cognitive theory by showing that it is practically applicable to the real world. Without such validation, a cognitive theory is not very compelling. Thus, it is difficult to accept the dynamicist approach as presenting a compelling new cognitive paradigm, especially since dynamicists are proffering their theory as one which should supercede existing approaches.

3.3.2 The Place of Dynamicism

Ascertaining the relation between the dynamicist approach to cognition and those of symbolicism and connectionism is not a straightforward task. Determining this relation is one of the primary aims of this thesis. Consequently, we will now attempt to elucidate this relation, drawing heavily on past discussions. As a result we will be in an excellent position to determine a suitable place for dynamical systems theory in cognitive science research.

Dynamicists tend to be quite succinct in presenting their opinion of what this relation should be and are not shy about their project to replace current cognitive approaches: “we propose here a radical departure from current cognitive theory.” (Thelen and Smith, 1994, p. xix). The dynamicist wish to dissolve connectionism and symbolicism has given them reason to critically assess the theoretical commitments of both paradigms. A detailed comparison of the three accounts has been provided in chapter 2. We found that dynamicism has effectively challenged the symbolicist approach in its poor integration of temporal considerations in cognition and has, at the very minimum, put into question aspects of the architectural, representational and computational commitments of symbolicism. However, similar challenges were ineffective when dynamicists turned to scrutinize connectionism. Rather, it seemed that the dynamicists are unable to substantially differentiate themselves from connectionists. When they manage to do so, they encounter their greatest theoretical challenges (see sections 3.2.4 and 3.2.5). This claim is quite devastating to the dynamicist commitment to establishing a new paradigm and thus warrants further discussion.

Van Gelder urges us to accept that connectionist networks are too limited a way to think about dynamical systems. He claims that “many if not most connectionists do not customarily conceptualize cognitive activity as state-space evolution in dynamical systems, and few apply dynamical systems concepts and techniques in studying their networks” (van Gelder, 1993, p. 21). However, there are a great number of influential connectionists, including the Churchlands, Pollack, Meade, Traub, Hopfield, Smolensky and many others who have addressed connectionist networks in exactly this manner.

There does not seem to be any lack of examples of the application of dynamical systems descriptions to networks. Kauffman (1993) discusses massively parallel Boolean networks in terms of order, complexity, chaos, attractors, *etc*. In fact it seems the only viable way to discuss such large (*i.e.* 100 000 unit) networks is by appealing to the overall dynamics of the system and thoroughly apply dynamical systems concepts, descriptions and analysis (Kauffman, 1993, p. 210).

Van Gelder insists that dynamical descriptions of connectionist networks is where connectionists should be headed, with which many connectionists would clearly concur, but he goes on to conclude that connectionism "is little more than an ill-fated attempt to find a half-way house between the two worldviews [*i.e.* dynamicism and symbolicism]" (van Gelder and Port, in press, p. 27). Rather, it seems connectionism may be the only viable solution to a unified cognitive theory, since cognition seems to be neither solely representational/symbolic nor nonrepresentational/dynamical. Connectionism is able to naturally incorporate both dynamical and representational commitments into one theory. In any case, all that van Gelder has really accomplished is to cast a dynamical systems theory description of cognition into the role of a normative hypothesis for connectionism - he has no basis for claiming to have identified a new paradigm.

The fundamental disagreement between connectionists and dynamicists seems to be whether or not connectionist networks are satisfactory for describing the class of dynamical systems which describe human cognition. By claiming that connectionist networks are "too narrow" in scope, van Gelder wishes to increase the generality of the dynamical description. However, differentiating between connectionist networks and dynamical systems is no easy task. Frequently, dynamicists realize: "indeed, neural networks, which are themselves typically continuous nonlinear dynamical systems, constitute an excellent medium for dynamical modeling" (van Gelder and Port, in press, p. 26). In Smolensky's paper *On the Proper Treatment of Connectionism*, he has outlined some of the many ways in which a connectionist network is a dynamical system (1988, p. 6):

> The state of the intuitive processor at any moment is precisely defined by a vector of numerical values (one for each unit). The dynamics of the intuitive processor are governed by a differential

> equation. The numerical parameters in this equation constitute the processor's program or knowledge. In learning systems, these parameters change according to another differential equation.

As Smolensky explicitly noted, a connectionist network represents the state of the system at any particular time by the activity of all units in the network. These units are naturally interpretable as axes of a state space. Their behaviors can be effectively described in dynamical systems theory terms. Such systems *are* nonlinear, differentially describable, self-organizing and dynamical as they trace a path through their high order state space (see section 2.3.5).

The behavior of these networks is *exactly* describable by the state space and the system's trajectory; as in any typical dynamical system. In other words, the tools provided by dynamical systems theory are directly applicable to the description of the behavior of connectionist networks. Examples of strange attractors, chaos, catastrophe, *etc*. are all found in connectionist networks, and such concepts have been used to analyze these networks.

Critiques of connectionism from dynamicists do not seem to present any sort of united front. Some dynamicists note the lack of realism in some networks (Globus, 1992). Others reject connectionism not because of a "failure in principle" but because of "a failure of spirit" (Thelen and Smith, 1994, p.41).[24] Still others reject connectionism as being too committed to symbolicist ideas: ideas like *representation* (van Gelder and Port, in press, p. 27).

The lack of realism in networks is often intentional due to the limitations of current computational power. Networks as complex as those found in the human brain are infeasible to simulate on present-day computers. The complexity of real networks does not represent a qualitatively distinct functioning, rather just the end-goal of current connectionist models. Thus, claims consonant with: "simplified silicon nets can be thought of as computing but biologically realistic nets are non computational" (Globus, 1992, p. 300) are severely misleading. The chemical modulation of neurotransmitter synthesis, release, transport, *etc*. is simply a more complicated *process*, not a qualitatively different method of functioning. As Globus (1992) later admits, connectionist networks "severely

---

[24]This sort of criticism is again a normative one. It does not serve to support a rejection of connectionism in favour of dynamicism, it serves only to note considerations which *should* be addressed by connectionists; it seems quite clear connectionists have the ability to address the considerations raised by Thelen and Smith (1994).

stretch" the concept of computation in the direction of dynamical systems theory. There is no dynamical behavior that cannot be reproduced by a connectionist network. Thus, it is a very safe assumption that the distinctions between real and simplified networks which Globus advances are ones which will, with time and improved processing power, become obsolete. From a theoretical standpoint such a distinction is uninteresting.

In the connectionists' favor, a recent paper by Meade (1994) has proposed a closer relationship between dynamical systems models and connectionism than would be comfortable for any dynamicist. In this paper, Meade uses a feed forward artificial neural network to solve a system of three coupled nonlinear second and third order ordinary differential equations (Meade and Fernandez, 1994) . The techniques that he has developed are also applicable to partial differential equations. And, with the introduction of recurrent networks, Meade is currently working towards using "ANN [artificial neural network] elements to solve any series of ODEs [ordinary differential equations] that have been solved by computational methods" (Meade, personal comment). Thus, any time dependent system of differential equations, in other words any of the class of dynamical systems that dynamicists wish to call their own, can be modeled non-trivially by connectionist networks. With this sort of evidence being presented, it is difficult to accept the dynamicist claim that they are introducing a new paradigm.

3.4 Conclusion

It is undeniable that brains are dynamical systems. Cognizers are situated agents, exhibiting complex temporal behaviors. The dynamicist description emphasizes our ongoing, real-time interaction with the world. For these reasons, it seems that dynamical systems theory has far greater appeal for describing some aspects of cognition than classical computationalism. In describing these cognitive dynamics, how exactly should we apply dynamical systems theory? In resolving this question, dynamicists are unconvincing in their claim that they have identified a need for a paradigmatic shift.

Rather, connectionists have a more compelling claim that they have been effectively applying dynamical systems theory for years. There are indeed many connectionists who do not apply dynamical systems concepts in describing the behavior of their models. However, this fact only leaves the dynamicists with a normative claim, nothing stronger.

Connectionist networks are dynamical systems. They stand up extremely well to the criticisms of dynamists - though they are currently simple approximations of real neural networks for technical reasons. Again it seems that dynamical explanations should be a normative requirement of connectionists, allowing for an integration of dynamical tools into the connectionist paradigm.

There is simply no reason to reject connectionism in favor of dynamicism. The dynamicists have a lot to say about the direction connectionist research should follow, but there are no clear distinctions between possible behaviors of connectionist and dynamicist models. The importance of dynamics should not be underestimated by the connectionists. Dynamical systems theory can contribute invaluably to the description, discussion and analysis of cognitive models. Possibly more connectionists should realize “Our brains are dynamical, not incidentally or in passing, but essentially, inevitably, and to their very core” (Churchland and Sejnowski, 1992, p. 187). But even with such realizations to rely on, the dynamicists fail to create a “new paradigm”.

It is difficult to accept that dynamical models can effectively stand as their own class of cognitive models. The difficulties which arise at the proposed level of generality seem insurmountable, no matter the resources available. Currently, dynamicist models end up being metaphorical description or a connectionist implementation of a complex system. They seem to offer exciting new ways of understanding these systems and of thinking intuitively about human behavior. However, as a rigorous descriptive model of either, the purely dynamical description falls disappointingly short. At most, dynamicists offer new metaphors and interesting discussion, but shaky models. However, at the very least they offer a compelling normative direction for cognitive science.

References

Abraham, Abraham, and Shaw (1994). Dynamical Systems for Psychology.

Achinstein, P. (1968). Concepts of science. Baltimore, Johns Hopkins Press.

Anderson, J. R. (1983). The architecture of cognition. Cambridge, MA, Harvard University Press.

Barton, S. (1994). "Chaos, self-organization, and psychology." American Psychologist **49**(1): 5-14.

Beardsley, M. (1972). Metaphor. The encyclopedia of philosophy. New York, MacMillan Publishing Co. & The Free Press. 284-289.

Bechtel, W. (in press). "Natural deduction in connectionist systems." Synthese.

Brooks, Rodney A. (1991). "Intelligence without reason." MIT AI Lab Memo 1293, April 1991.

Cartwright, N. (1991). "Fables and models I." Proceedings of the Aristotelian Society Supplement **65**: 55-68.

Churchland, P. (1989). A neurocomputational perspective. Cambridge, MA, MIT Press.

Churchland, P. S. (1993). Presidential address. American Psycological Association,

Churchland, P. S. and T. Sejnowski (1992). The computational brain. Cambridge, MA, MIT Press.

Donald, M. (1991). Origins of the modern mind: Three stages in the evolution of culture and cognition. Cambridge, Harvard University Press.

Egeth, H. E. and D. Dagenbach (1991). "Parallel versus serial processing in visual search: Further evidence from subadditive effects of visual quality." Journal of Experimental Psychology: Human Perception and Performance **17**: 551-560.

Fodor, J. and Z. Pylyshyn (1988). "Connectionism and cognitive architecture: A critical analysis." Cognition **28**: 3-71.

Gleick, James (1987). Chaos: making a new science. New York: Viking.

Globus, G. G. (1992). "Toward a noncomputational cognitive neuroscience." Journal of Cognitive Neuroscience **4**(4): 299-310.

Goodwin, B. (1994). How the leopard changed its spots: The evolution of complexity. United States, Charles Scribner's Sons.

Harnard, S. (1992). Connecting object to symbol in modelling cognition. Connectionism in Context. London, Springer-Verlag.

Hawking, S. (1988). A brief history of time. London, Bantam Press.

Hesse, M. (1972). Models and analogies in science. The encyclopedia of philosophy. New York, MacMillan Publishing Co. & The Free Press. 354-359.

Holyoak, K. and P. Thagard (1995). Mental leaps: Analogy in creative thought. Cambridge, MA, MIT Press.

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

Koertge, N. (1992). "Explanation and its problems." British Journal of the Philosophy of Science **43**: 85-98.

Kosslyn, S. (1980). Image and mind. Cambridge, Harvard University Press.

Langer, Susanne K. (1957). Philosophy in a new key: A study in the symbolism of reason, rite and art. Cambridge, Harvard University Press.

Le Poidevin, R. (1991). "Fables and Models II." Proceedings of the Aristotelian Society Supplement **65**: 69-82.

Luck, S. J. and S. A. Hillyard (1990). "Electrophysiological evidence for parallel and serial processing during visual search." Perception and Psychophysics **48**: 603-617.

Meade, A. J. and A. A. Fernandez (1994). "Solution of nonlinear ordinary differential equations by feedforward neural networks." Mathematical and Computer Modelling **To Appear**:

Miller, J. O. (1988). "Discrete and continuous models of human information processing: Theoretical distincitons and empirical results." Acta Psychologica **67**: 191-257.

Molenaar, P. C. M. (1990). "Neural netowrk simulation of a discrete model of continuous effects of irrelevant stimuli." Acta Psychologica **74**: 237-258.

Morton, A. (1988). "The chaology of mind." **48**(3): 135-142.

Newell, A. (1990). Unified theories of cognition. Cambridge, MA, Harvard University Press.

Pinel, J. P. (1993). Biopsychology. Allyn & Bacon Inc.

Popper, K. (1981). The rationality of scientific revolutions. Scientific Revolutions. New York, Oxford University Press.

Putnam, H. (1981). The 'corroboration' of theories. Scientific Revolutions. New York, Oxford University Press.

Robertson, S. S., A. H. Cohen, G. Mayer-Kress. (1993). Behavioural Chaos: Beyond the Metaphor. A dynamic systems approach to development: Applications. Cambridge, MIT Press. 120-150.

Schweicker, R. and G. J. Boggs (1984). "Models of central capacity and concurrency." Journal of Mathematical Psychology **28**: 223-281.

Simon, H. A. and A. Newell (1970). Information-processing in computer and man. Perspectives on the computer revolution. Englewood Cliffs, Prentice-Hall, Inc.

Skarda, C. A. and W. J. Freeman (1987). "How brains make chaos in order to make sense of the world." Behavioral and Brain Sciences **10**: 161-195.

Smolensky, P. (1988). "On the proper treatment of connectionism." Behavioral and Brain Sciences **11**(1): 1-23.

Thagard, P. (1992). Conceptual revolutions. Princeton, Princeton University Press.

Thagard, P. (in press). "Modelling conceptual revolutions." Dialogue.

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

Ubbink, J. B. (1961). Model description and knowledge. The concept and the role of the model in mathematics and natural and social sciences. Holland, D. Reidel Publishing.

van Gelder, T. (1993). "What might cognition be if not computation?" Cognitive Sciences Indiana University Research Report 75 :

van Gelder, T. and R. Port (in press). It's about time: An overview of the dynamical approach to cognition. Mind as motion: Explorations in the dynamics of cognition. Cambridge, MA, MIT Press.