Dietrich, E. and A. B. Markman (2003). Discrete Thoughts: Why cognition must use discrete representations. *Mind and Language*. v. 18, n. 1, pp. 95-119.

Discrete Thoughts:

Why cognition must use discrete representations

Eric Dietrich

Dept. of Philosophy

Binghamton Univ.

Dept. of Psychology Univ. of Texas

Arthur B. Markman

Binghamton, NY

Austin, TX

dietrich@binghamton.edu

markman@psy.utexas.edu

Abstract

Advocates of dynamic systems have suggested that higher mental processes are based on continuous representations. In order to evaluate this claim, we first define the concept of representation, and rigorously distinguish between discrete representations and continuous representations. We also explore two important bases of representational content. Then, we present seven arguments that discrete representations are necessary for any system that must discriminate between two or more states. It follows that higher mental processes require discrete representations. We also argue that discrete representations are more influenced by conceptual role than continuous representations. We end by arguing that the presence of discrete representations in cognitive systems entails that computationalism (i.e., the view that the mind is a computational device) is true, and that cognitive science should embrace representational pluralism.

Cognitive science is young. It has been less than 150 years since the scientific method began to be applied in earnest to cognitive phenomena. One mark of cognitive science's youth is that paradigm shifts -- wholesale changes in the way we conceptualize the scientific basis of the field -- are common. It has been less than 50 years since computationalism, the current dominant paradigm, became ascendant in cognitive science, replacing behaviorism's positivist grip on the field. Computationalism is the view that thinking involves algorithmic processes operating over representations that serve as the data structures of the mind (see Dietrich, 1990 and 1994, and Dietrich and Markman 2000a, for an extended discussion of computationalism).

Another mark of the youth of cognitive science is the existence of significant problems with its foundations. The science of computation is still being developed, so both cognitive science and computational theory are being worked out simultaneously. Furthermore, the foundational concept of representation also has a number of unsolved mysteries. For example, there is no agreed upon theory of representational content. Finally, there are many cognitive phenomena that are difficult to explain within the computational paradigm. Despite nearly 50 years of research, we do not yet have machines that can communicate, read, or classify with even the skills of an average five-year-old child.

The authors thank William Bechtel, Valerie Hardcastle, John Hummel, Brad Love, and Linda Smith for discussions of these issues and comments on earlier versions of this manuscript. We also thank two *Mind and Language* referees for good comments. Thanks are also due to Tiger Woods for his skill and competitiveness, and just for playing a darn good game of golf. This work was supported in part by NSF grant SBR-9905013 given to the second author.

Many sciences have significant unsolved problems, and these problems form the basis of the daily work of scientists in the field. Most of these problems are finally solved in ways that fit within a science's current paradigm. Only rarely do the solutions to these problems require a radical restructuring of the theoretical basis of a field. This success instills confidence in the paradigm. But, because of its youth, unsolved problems in cognitive science are a source of doubt that the computational paradigm is correct. There is, therefore, significant temptation to assume that computationalism is wrong, and that the unsolved problems are signs that a new paradigm is needed.

The main way to attack computationalism, at least the main way it has been attacked historically, is to show that cognition is incompatible with the notion of representation. Logically speaking, there are two ways to do this. One can try to show that cognition doesn't require any representations at all, and that representations just confuse the issues; or one can try to show that cognition does require representations but that the representations have properties that prevent them from being processed by algorithms. Despite the rhetoric (e.g, the title of Brooks' 1991 paper 'Intelligence without representation'), very few researchers have tried the first approach, and for a rather deep reason that was not widely appreciated. We have made an extensive argument that the minimal, core notion of a representations (Dietrich and Markman, 2000a; Markman and Dietrich, 2000a and b¹. (van Gelder is one who has actually tried to argue that cognition is possible without representations (van Gelder, 1995), but his view is now regarded as erroneous. See Bechtel, 1998).

The second approach, arguing that cognition requires representations that are not manipulable by algorithms is by far the most favored approach. A current example of this approach that has been proposed with some force is dynamic systems theory (e.g., Port and van

¹ The literature on the necessity of some kind of representation for the possibility of cognition is large. For some classical texts on the topic, see: Bruner, 1957; Fodor, 1975; Marr, 1982. For a good, modern discussion, see Pinker, 1997.

Gelder, 1995; Thelen and Smith, 1994). Dynamic systems are collections of nonlinear differential equations that express the way continuous quantities change over time. Advocates of this view argue that a cognitive science based on dynamical systems can dispense with the kinds of representations and processes that have been the norm under the computational view. The dynamic systems view differs from computationalism in the kinds of representations that are assumed to be necessary for cognitive processing. The dynamic systems approach posits that continuous, time-dependent representations are sufficient to explain cognitive processing. This view explicitly tries to conclude that discrete representations (e.g., symbols) are not necessary for understanding cognitive phenomena, and using discrete representations will lead cognitive science astray (see, for example, Bickhard and Terveen, 1995; Horgan and Tienson, 1996; Port and van Gelder, 1995; Thelen and Smith, 1994; van Gelder, 1995).

In this paper, we argue that this conclusion is false; cognitive systems *do* require discrete representations. To establish this point, we first define the concept of representation, and distinguish between discrete and continuous representations. Then, we present seven arguments supporting the claim that cognitive systems must have discrete representations. We further show that the presence of discrete representations in cognitive systems entails that computationalism is the proper paradigm for cognitive science. Finally, we suggest that cognitive science should embrace representational pluralism, because cognitive systems are likely to involve many different kinds of representations.

2. Defining discrete representations

In order to distinguish between discrete and continuous representations, we first give a definition for the concept of representation. This definition is quite broad, and is intended to encompass many different types of proposals that cognitive scientists have made about representation. We have defended this view of representation extensively elsewhere, and so we will focus here only

on points that will be relevant for the later discussion (see Markman, 1999, Markman & Dietrich, 2000, for more discussion).

2.1. A general definition of mental representation

A cognitive *agent* is an organism or artificial system that thinks. An agent takes input from the outside world and performs actions that affect the outside world. In addition, the cognitive agent has subcomponents that we refer to as *systems*. We define a system as any subcomponent within the cognitive agent that uses information in an attempt to satisfy its goals. Systems, on this definition, must have feedback loops (at least negative ones) because they need to determine whether or not their goals are satisfied. Some of an agent's systems take input from the agent's external environment, but for a cognitive agent of any complexity, i.e., one with many interacting systems, most often a system's input is the output of some other system within the agent. The output of a system may likewise be an action that affects the outside world or, frequently, information that is used as input by another system within the cognitive agent.

On this view, a representation is any internal state that mediates or plays a mediating role between a system's inputs and outputs in virtue of that state's semantic content. We define semantic content in terms of information causally responsible for the state, and in terms of the use to which that information is put. Hence, any state that mediates between the input to an agent and the actions performed by the agent is a representation (this is somewhat unintuitive because it is so inclusive, but it has enormous benefits, especially in unifying cognitive science; see below and Markman and Dietrich, 2000a).

This definition can be expressed as four necessary and jointly sufficient conditions on the existence of representations inside an agent: 1) there must be at least one component (i.e., a system) within the agent that has internal states that carry information and that govern the system's behavior, 2) there must be an environment external to the system (as noted, this need not be, and for most systems in the agent won't be, the environment external to the agent), 3) there must be informational relations between the system's environment and the system's

internal states that determine how the two correspond, 4) the system must have internal processes that use the information in the internal states to perform processes such as accomplishing goals and solving problems (see Markman and Dietrich, 2000a, for a more detailed discussion of these four conditions).

In order to serve as a mediator, a representation's content has to be explicit. That is, content must be psychologically real in order to be causally efficacious, and causal efficacy requires explicitness. As we discuss below, content is a kind of information, and implicit information (e.g., the number four in 2 + 2) is not causally efficacious because no process can use it.

It is important to be careful here. We do not define 'explicit' in terms of being physically manifested, that is, in terms of what physics studies. Another way to put this point is to say that representations have to take part in *mental* causation, but the physical laws that govern the same instance of mental causation in different agents may be different. Helen's beauty caused a thousand ships to be launched, but there is nothing at the level of physics that separates instances of beautiful things from instances of non-beautiful things (some particle physicists do study beauty, but not the kind we refer to here). Therefore, we are not saying, nor are we committed to the view that the semantic content of a representation is its physical manifestation. That is, we are not committed to a strong reductionism whereby psychological explanations are replaceable by physical ones.

In the computational view of mind, mental causation is identified with a level of explanation involving representations and processes. As a simple example, consider a computer running the spell-checker in a word processor. The spell-checker compares strings such as 'teh' to a list of words, and finding 'the' queries the user that perhaps 'teh' is a misspelling of 'the'. This simple explanation is couched at the level of the functioning word processor running its spell-checker. That is, the word processor is reified and treated as a separate, stand-alone entity. The technical term for the spell-checker is a *virtual machine*. ('Virtual' does *not* mean 'not real.' The

reality of this level of explanation is readily apparent when one has to debug the spell-checker or alter it, say, by adding or deleting words from it (see Dietrich & Markman, 2000a)). The string 'teh' has to have a physical manifestation as a pattern of bits in a memory location, as does the spell-checker, and indeed, as does the entire word processor. But the level of explanation at which these bit patterns exist is far below the level of the spell-checker and is usually not relevant to explaining, augmenting, or debugging the spell-checker. Indeed, the physical instantiation of this word processor will vary widely depending on the particular computer on which it is run.

The conclusion to be drawn here is that the physical level is *not* the only level at which causation occurs. Physical causation is just one level of causation, and causation occurs in its own way in all the virtual machines that can be used as higher level descriptions of the physical machine (Marr, 1982; Pylyshyn, 1980). This in turn means that the psychological level of explanation is a real level and that psychological explanations are ineliminable. Hence, when we say that semantic content has to be explicit, we mean that it must be explicit in some virtual machine. Being explicit in a virtual machine entails having a physical manifestation, but the physical manifestation is not 'more real,' and is not to be equated with the content.

Given this foundation, we can discuss the way representations come to have content at the level of mental causation. The content of a representation derives from two separate sources. The first source is external and informational, by which we mean that there are informational relations between the representation and states in the system's external environment (which is, to repeat, not necessarily the environment external to the agent) that connect the representation with the external states it represents. We will often refer to this source of semantic content as *correspondence*.²

² When we use the term 'information' we mean externally provided information in the sense just defined. See Markman and Dietrich, 2000a, for our technical discussion of information in representation which borrows from and extends Dretske's theory, (1981).

The second source of content is internal and functional, by which we mean that there are various internal processes that operate on the representations that, in virtue of their operation, give rise to (and are in turn constrained by) functional, logical, and associative relations between the representations. Because representations are individuated by their contents and their structure, the functional role source of representational content entails that representations are in part individuated by what they do. Functional relations are rather numerous and usually implicit, so there must be some mechanism that reifies some of the functional relations and makes them an explicit a part of the representations. (For more on the functional source, see, e.g., Millikan, 1984; our view about the origins of representational content is similar to Bechtel's (1998)).³

The two sources typically work together to establish a representation's content. That is, both correspondence and functional role influence the content of a representation. Though the two sources work together, they generally do not contribute equally to the content of a particular representation. For example, the content of early sensorimotor representations is primarily fixed by external, informational relations, particularly for continuous representations such as vibrating ear drums and the like.⁴ In contrast, higher, more abstract, less modality-specific representations have their content fixed primarily by their functional roles. Some higher-level, abstract representations no doubt have their content fixed almost entirely by functional role. For example, concepts such as <u>infinity, deity</u>, <u>sacredness</u>, <u>profundity</u>, <u>the universe</u> depend primarily

³ In terms of temporal precedence, it is likely that the informational source precedes the functional, internal source, especially at the sensory surface of the agent. But, perhaps sometimes, a representation's functional role precedes its informational relations. This might happen, for example, if certain concepts (a kind of representation) are innate.

⁴ It might be thought that vibrating ear drums are not representations. They are not *mental* representations, to be sure, but they are representations. Not all representations which a cognitive system needs are mental ones. Seeing vibrating ear drums as representations is a small price to pay for the power and coherence of the resulting theory (in Markman and Dietrich, 2000a). The two main conclusions of this theory are that 1) many different representational types, methodologies, and explanatory models are going to be needed to explain something as complex as the mind, and 2) all of these representational types can be derived by adding various properties to the single, invariant notion of a *mediating state* (for details, see Markman and Dietrich, 2000a).

on their functional roles and little on their informational connections to the world because these concepts vastly outstrip their informational content.

It might be that at least some concepts are strictly functional because they are connected to absolutely nothing in the system's external environment. This possibility is unlikely, because it is hard to imagine a concept that is not connected to anything external to the system (Barsalou (1999) defends the importance of perceptual processing in very abstract concepts.) Even contradictory concepts such as <u>round-square</u> have some correspondence, because the constituents (<u>round</u> and <u>square</u>) have correspondences external to the system. Thus, the contradictory concept can be grounded in the correspondence of its constituents.

Our definition of a representation is not intended as a theory of representational content. We are specifying quite generally and abstractly the conditions that a system's internal states have to meet to be representations. At present, our proposal lacks the detail needed to classify it as a genuine theory. It is as if we had said that to have a fire one needs heat, fuel, and oxygen without saying how one goes about getting some heat, fuel, and oxygen, and without saying much about the three ingredients work together. To get a theory of content, we would have to solve a suite of problems that include the problem of reference (does the existence of an informational relation between a representation and something in the environment entail that the representation refers?), the problem of representational error (all representations allow misrepresentation, but how does a representation have content if the representation doesn't correspond to anything in the environment, i.e., if it doesn't refer?), the problem of holism (does changing a miniscule part of a representation's functional role change the meaning of the representation?), and the problem of explaining how reification of functional role works. However, we think that our definition of representation is robust and clear enough to work as a basis for our definition of discrete representations.

2.2. Discrete representations

Our goal in this paper is to defend the claim that cognitive systems have discrete representations, and the further methodological claim that theories of cognitive systems need to postulate and use discrete representations. Having given a definition of representation, we now turn to a definition of discrete (and continuous) representations. The standard way of defining discreteness is to consider a set of items and define discreteness over that set. A set is discrete if and only if it has gaps between its members. If the set has no gaps whatsoever between its members, it is continuous (and vice versa). For example, the set of rational numbers (numbers of the form p/q) is discrete. Though there are an infinite number of rationals and though the rationals are dense (between any two of them there is a third rational), the set has gaps: e.g., π and $\sqrt{2}$ are not in the set. Indeed all the irrationals are missing. The set of real numbers (the rationals together with the irrationals) is continuous because this set has no gaps -- it is missing no numbers. This treatment of discreteness and continuity codifies the intuition that discreteness means that there are several different individual entities that can be discriminated whereas continuity, in one important sense, means that there is just one unified entity. Though there are an infinity of real and rational numbers, the real number line is a unity, but the set of rational numbers is not.⁵

Given this, it seems natural to define 'discrete representation' this way: a system has discrete representations if it contains more than one representation, and the representations are bounded and uniquely identifiable. There is something right about this, of course, but there is something deeper to be said. If a system *categorizes* environmental inputs then it has discrete representations. A system categorizes inputs if it has internal states that impose *classes of sameness* on those inputs. This means that the system will be in the same state for different

⁵ Technically, sets aren't continuous; functions are. The property of sets related to continuity of functions is *connectedness*. We use the term 'continuous' in conjunction with representations for three reasons: 1) this is the term usually used in cognitive science, 2) the distinction between continuity and connectedness is not that important here because representations can be viewed as functions (from inputs to outputs), and because functions can be viewed as sets (of ordered pairs), and 3) even though we are eliding the distinction between 'continuous' and 'connectedness,' the central distinction for our purposes remains: some representations, but not others, are tightly coupled to gapless events in the environment (e.g., a sound wave) and hence have no gaps in themselves. Such representations we call continuous.

inputs: though the inputs themselves differ, the system is unable to discern the difference. This claim requires some discussion.

There is a connection between being able to categorize and being able to discriminate. (A system discriminates inputs if it somehow 'notices' that there is more than thing in the input stream). Categorizing inputs is identifying them or classifying them. Discriminating inputs is necessary for categorizing, but not sufficient. To categorize, enduring classes of sameness are needed. A system cannot discriminate between two external, environmental states with one, single continuously varying representation. With a continuous representation, the system can be in different states at different times (the sets of states and times having the cardinality of the reals), but it cannot distinguish among those states. To distinguish between two external states, S1 and S2, say, the continuous infinity of intermediate states between S1 and S2 have to somehow be elided. The only way to do that is if the system has two internal representations, R1 and R2, say, that chunk all the states in some neighborhood of S1 in with S1, and all the states in some other neighborhood of S2 in with S2. If a system cannot discriminate (and hence cannot categorize) inputs, but it still represents its environment, then the system has continuous representations.

This all suggests the following definition:

A system has discrete representations if and only if it can discriminate its inputs.

It follows from what we have said, and our definition, that if a system categorizes, then it has discrete representations 6 .

⁶ It has been suggested that perhaps our definition of discrete representation is too weak. Also, some have worried that our definition, coupled with our inclusive definition of representation, gives us our conclusion – that higher cognitive processing requires discrete representations – too easily, i.e., in a way that doesn't do justice to all the intuitions and discussion that has surrounded the notion of representation.

It is striking that cognitive science does not have a theory of representation. There are many theories out there, but none have the support of anything like the majority of researchers. How could a

Given our definition, we can now see the importance of the different contributions of correspondence and functional role to a mental representation's meaning. In a continuous representation, the primary contributor to meaning is correspondence. The representing relation that binds a continuous representation to what it is representing is a time-dependent physical law (i.e., the relation is *nomic*--see argument seven below). Such nomic binding is called *coupling*. Coupling means that the continuous representations and what they represent are tied together extremely tightly. Functional connections between a continuous representation inside the system can only moderate the effects of the coupling. For example, imagine a sensory device that vibrates in the presence of sound waves (like the cochlea of the human auditory system). Other representations might change the response properties of the device, perhaps by making it more rigid.

Discrete mental representations show more influence of the functional role source of information. A discrete representation will still have some informational relation that binds it to situations in the system's environment. However, this binding is not nomic and is not that tight. More importantly, a discrete representation will have many connections to other representations

notion that is so important be so poorly understood? There are two main reasons. One, which we mentioned at the beginning of our paper is that cognitive science is young and representation is exceedingly complex. We just haven't had enough time to figure it out fully. The second reason (and the one that we are most concerned to push in this paper) is that cognitive scientists (especially the philosophers) are not being pragmatic enough. We contend that it is far better to be inclusive about what counts as a representation and then let a theory of the different types of representation do the work of separating out the different types and levels of cognitive scientists are hoping that a theory of representation will enable them to draw a sharp distinction between those processes that are cognition and those that are not. But such a sharp distinction is not in the cards. Cognition is a graded notion, varying between species, individuals within a species, and even among the various mental or quasi-mental processes within a given individual. Our notion of representation has to be flexible enough to accommodate this fact.

So, our definition of representation is inclusive for a reason: it will be much more useful to distinguish between types of representation armed with an inclusive definition than it will be to forcibly draw a sharp boundary between representing and not representing, between cognizer and non-cognizer.

One can now see that our definition of discrete representation is not too weak. It does precisely the work that needs to be done in accordance with our claim about types of representation. It draws distinctions where they need to be drawn and it strongly highlights the importance of discrete representations in categorization and discrimination. Finally, our result is surprising: we have a shown, we believe, a necessary connection between mental capacities for discrimination and possessing discrete representations.

and representational elements in the system. Models incorporating discrete representations have posited connections that permit the spread of activation in semantic memory, and they have also suggested role-argument connections that bind discrete representations.

2.3. The Watt governor and the thermostat

A good way to illustrate the distinction between discrete representations and continuous ones is by contrasting the Watt steam engine governor and the standard thermostat. This device is designed to keep flywheels of steam engines rotating at a constant speed. The governor consists of a spindle with two arms attached to it, and on the end of each arm is a metal ball (see Figure 1). The governor is connected to the flywheel, usually via a pulley. The faster the flywheel rotates, the faster the governor rotates. As the governor rotates faster, the balls rise, due to centrifugal force. There is also a mechanical linkage between the rising and falling arms and a valve on a steam pipe that controls the amount of steam driving the flywheel. As the engine runs faster, the rising balls, via the linkage, cause the valve to close. The restricted valve decreases the amount of steam driving the flywheel, so it slows down, which in turn causes the governor to spin more slowly. The slower spinning balls drop, opening the valve, which causes more steam to flow, which causes the flywheel to accelerate, etc. In this way, a relatively constant speed for the flywheel can be maintained.

The arm angles are continuous representations of the speed of the flywheel, the speed of the piston in the steam engine, the amount of the steam flowing through the valve, etc.⁷ The process by which the governor controls the steam pressure in the engine is called *equilibration* because it maintains equilibrium in the steam engine via a tight feedback loop. Equilibration can be described as a path in a dynamic system's phase space, usually as a cycle around an attractor. Notice, the governor at no time discriminates between arm angles, nor between different degrees

⁷ The Watt governor has been touted as a feedback system that works without representations at all (Van Gelder, 1995, Thelen and Smith, 1994), but this claim has been shown to be wrong (Bechtel, 1998). The arm angles of the governor are representations, in fact, continuous representations. Note that the governor can misrepresent -- it can be fooled -- by artificially increasing or decreasing the amount of steam driving the flywheel, or by artificially increasing or decreasing the speed of the flywheel itself.

of how open the valve is on the steam pipe. The beauty of this device is that it maintains pressure equilibrium without having to discriminate. As the pressure continuously increases or decreases, the governor changes speed thereby continuously opening and closing the valve to which it is connected.

The Watt governor with its continuous representation can be contrasted with a thermostat, which uses both continuous and discrete representations. As shown in Figure 2, a thermostat usually has a curved bimetal strip. The strip changes its degree of curvature with temperature, because the two metals expand at different rates. If the faster expanding metal is placed on the outside of the curve, then when the temperature rises, the thermostat increases its curvature (because the outside metal expands faster than the inside one). The degree of curvature of the bimetal strip is a representation of temperature.

In order to affect the temperature of a room, the bimetal strip has a vial at the end of it that is filled with mercury (or some other conducting liquid). The vial has electrical contacts on each end (shown as dark circles on the vial). When the bimetal strip is highly curved (as in the example in Figure 2), the conducting liquid is pooled at the end of the vial, and the electrical circuit (which controls the heater) is open. When the strip is less curved (representing a cooler temperature), the vial is oriented so that the liquid closes the contacts, turning on the heater. Thus, the vial has two discrete representational states. The continuously varying information in the bimetal strip has been converted into an *on* and an *off* position.

It is important to notice that the thermostat uses continuous change and continuous representations, but this continuity is strictly inessential to controlling the air temperature. In the thermostat, the discrete representations are crucial to the functioning of the system. That is, discrete representations (when they exist) can *supervene* on continuous representations. (Properties of type A supervene on properties of type B if any two possible situations identical with respect to their B properties are also identical with respect to their A properties.) In the thermostat, the discrete representations are crucial to explaining the discriminating behavior of

the system. The bimetal strip in the thermostat continuously varies as a result of continuous changes in the air temperature. The amount of bending in the strip continuously represents the air temperature. It is only when the strip has bent far enough to close the circuit, however, that the bending matters. A digital thermometer would work just as well. This device would sample the air temperature every so often (time would have to be measured discretely, perhaps by counting vibrations of something), and then compare the measured temperature against a previous set temperature. When the measured temperature fell below the set one, this thermometer would turn on the heater. Hence, though certain components of a standard thermostat are subject to continuous change, this continuity is at most only pragmatically responsible for the functioning of the thermostat. Thus, just because a system uses continuous representations doesn't mean that those representations (or any other continuous representations) are crucial to the functioning of the system.

To summarize, a system has discrete representations if and only if it is able to use those representations to discriminate at least two distinct conditions in its input (typically system input), and categorize those inputs accordingly. A continuous representation is a single continuous quantity. Continuous representations derive most of their content from their correspondence to the system input. In contrast, discrete representations typically get an important part of their content from functional role relations to other representations in the system. Finally, a system that has discrete representations may also have some continuous representations in it (for more complex organisms, this is probably the typical case). In this case, the discrete representations may supervene on its continuous ones.

3. Why cognition requires discrete representations.

In this section, we give seven arguments that cognition requires discrete representations. These arguments, summarized in Table 1, are organized around two themes. Arguments 1-4 focus on the importance of discrimination in cognitive systems and on the importance of having discrete

units as components of cognitive representations. Arguments 5-7 focus on the relative contribution of correspondence and functional role in providing content in mental representations. These arguments turn on the fact that functional role is a more important source of content for discrete representations than for continuous representations.

3.1. Arguments based on discrimination.

The first argument for discrete representations is simply the observation that cognitive systems do in fact discriminate among states in the world. Cognitive agents reason and judge (that is why cognitive science posited representations in the first place). Agents reason and make judgments about individual items in their environment. Individual items are selected out of the environment by discriminating them from their surround. Hence cognitive systems have to discriminate. As discussed above, however, it is not possible to discriminate among states using a continuous representation. Because it co-varies in a time-dependent fashion with its input, a continuous representation treats the entire input stream as a unified whole that varies continuously. Hence, it can't sunder the represented input into two or more states. The time-dependent nature of continuous representations is often touted as a virtue (e.g., van Gelder & Port, 1995), but in fact, it prevents continuous representations from being able to discriminate among states. In short, coupling, the very thing that makes continuous representations useful, also prevents them from being able to discriminate, and hence categorize.

Another way to put this argument is to note that cognitive systems need to be able to *refer* to individual entities or items in their environment (for example, in order to avoid predators, find mates, find food, and to communicate with others). Reference requires discrimination, and thus it requires discrete representations.

We are not saying that continuous systems cannot act differently in different circumstances. They can. The physical law that relates input to the continuous representation will typically be nonlinear and give very different outputs as a function of the value of the relevant quantity in the input. But acting differently in different situations is *not* the same as

discriminating between two situations. The thermostat's bimetal strip behaves one way when the air temperature is increasing and another when the air temperature is decreasing, but it is not thereby discriminating between the two situations; only the binary, mercury switch can do that. And that switch has discrete states.

The second argument parallels the idea that cognitive systems need to be able to refer to individual entities in the world. Cognitive systems also need to be able to pick out specific properties of represented items, properties like 'has wheels' and 'has a gas-powered engine'. Doing this requires discrete representations.

To see this, consider, for example, that people are good at accessing the commonalities and differences that emerge from comparisons. However, continuous representations are unable to provide access to specific commonalities and differences. Spaces are commonly used as continuous representations in cognitive science. The use of a space as a representation is supported by techniques like multidimensional scaling (Shepard, 1962) and also by techniques for creating high-dimensional semantic spaces from corpora of text (Burgess & Lund, 1997, 2000; Landauer & Dumais, 1997). In these models, concepts are represented as points or vectors in space. Comparisons of pairs of concepts involves finding the distance between points or vectors. The output of these comparisons is a scalar measure of similarity. But scalar distances between concepts clearly are not the same thing as the commonalities and differences between the concepts. To get these, one needs to treat concepts as composed of discrete subparts that are recoverable during comparisons and which allow discrimination between properties.

Many comparisons require more than an overall measure of similarity. For example, Landauer & Dumais (1997) demonstrated that a model generated by finding proximities among words in a large corpus of text was able to perform a synonyms test (the Test of English as a Foreign Language (TOEFL)) at a level equivalent to that of a non-native speaker of English by taking the vector for the target word, and then finding the choice item that was most similar to the target. This model could not perform an antonyms test, however, because antonyms are also

highly similar words. Finding antonyms requires aligning the representations of a pair, finding a salient dimension of difference, and ensuring that the items have the opposite value on that dimension. Thus, finding antonyms requires accessing the commonalities and differences of a pair, which cannot be done with a continuous representation (see Markman, 1999, and Medin, Goldstone, & Markman, 1995, for more discussion of this point).

Research on structural alignment in similarity comparisons also demonstrates the way people access commonalities and differences of similarity comparisons (Gentner & Markman, 1997). For example, when people are asked to list the commonalities and differences of pairs of similar words or pictures, they are able to list many commonalities and also many differences related to those commonalities, called *alignable differences* (Markman & Gentner, 1993, 1996). When people list the commonalities and differences of dissimilar pairs, they are able to list few commonalities and alignable differences. The ability to access commonalities and differences, indeed, the ability to make all sorts of complex, robust comparisons between concepts, is incompatible with spatial representations (and continuous representations in general).

The third argument builds on the second. Not only do cognitive systems compare concepts (and representations of all sorts), but they often need to combine them too. By 'combine' we mean that information in one representation is joined with another. For example, representations are combined in the process of concept formation. Concepts are constructed from lower level representations of nonidentical instances. To learn the concept <u>apple</u>, for example, one builds it out of representations of the apples one comes into contact with. Representation combination is also done in a variety of other cognitive processes such as analogical reasoning and reminding. Even animals without a symbolic, natural language need to combine representations from time to time because, of course, they form concepts.

Combining information isn't possible unless one uses discrete representations. Basically, when continuous representations are combined, all that results is another continuous

representation where all the original information is lost due to blurring. We illuminate this point with a tiny portion of the history of genetics.

Following the rediscovery in 1900 of Gregor Mendel's work on genetics, Darwin's theory of evolution via natural selection (first published in 1859) should have been immediately accepted, for Mendel's theory solved a long-standing problem about inheritance.⁸ During Darwin's lifetime, the mechanisms of inheritance were poorly understood. The going theory, embraced by Darwin, was that inheritance worked by a sort of blending of whatever it was that was responsible for offspring having the coloring, form, and other properties that they do. This was called *blending inheritance*. The theory of blending inheritance assumed that the stuff responsible for inheritance was an infinitely divisible, continuous fluid, perhaps even blood. In 1867, Fleeming Jenkin, a Scottish engineer, raised an objection to Darwin's theory of natural selection by pointing out that such blending inheritance would, in a very short time, result in a completely homogeneous population, and that natural selection would then have nothing to work on. For example, mixing many different colors of paint in a pail quickly leads to a darkish gray sludge from which the original colors cannot be extracted. Of course, strictly speaking, Jenkin's objection was only a reductio on the blending theory, but this theory was considered so much a part of Darwin's theory of natural selection, that any trouble for the blending theory was transferred to the Darwin's theory, too.

In spite of Jenkin's objection, the theory of blending inheritance remained, mainly because there was no obvious replacement. Ironically, Mendel's particulate theory of inheritance, which would turn out to be the rebuttal to Jenkin's objection, was published at about the same time as Jenkin's objection. Mendel's work was completely ignored until 1900.

⁸ We say Darwin's theory should have been accepted, but it was not. For many years, it was believed that Mendel's genetic theory and Darwin's theory of natural selection were in competition. Indeed, most viewed Darwin's theory as losing the battle. Then in the 1930s, R. A. Fisher, J. B. S. Haldane, and Sewell Wright showed that not only were Mendelian genetics and natural selection compatible, they needed one another: indeed, genetics offered a partial *reductive explanation* of evolution via natural selection.

Mendel's first law, or the law of segregation, states that hereditary characteristics are determined by a finite number of discrete particles in the cells of organisms, and passed on to future generations in the germ cells of organisms. (These particles are now known as genes. Mendel knew nothing of genes as we know them, nor did he know anything about chromosomes, mitosis, or meiosis, all of which were discovered after Mendel.) Mendel determined that it is the relative frequency of these particles that are either present or not that matters in inherited characteristics. Tall pea plants when cross-mated with short pea plants produce all tall plants in the first generation. However, individuals of this first filial generation, when mated with each other, produce, in the second filial generation, tall and short pea plants at a ratio of about 3:1. The particles responsible for shortness are still present in the first filial generation, but somehow just not active (this phenomenon is called 'dominance'). It is clear, as it was to Mendel, that mating in no way dilutes or alters what is responsible for height. Mendel generalized this to all inherited characteristics. He then reasoned that the only way this can happen is if discrete particles are responsible for inheritance.

Some traits do appear to blend. For example, the colors of white and red four-o'clocks seem to blend to produce pink flowers. But, if the pink flowers are mated with each other, red, white, and pink flowers are produced, indicating that the particles for red and white are still present— undiluted and unaltered—in the pink flowers. The pink color, it is now known, is due to a phenomenon called 'incomplete dominance.' There is no blending going on.

The crucial property of genes that solved the problem of inheritance plaguing the theory of evolution was that genes are *discrete*. Each gene carries a fixed amount of information for constructing some protein. And, that information has to be discrete if it is to be combined, recovered, and used over and over again to do the job it is supposed to do. Further, the operations that are carried out on genes are discrete operations of the sort computational systems perform. For example, genes are copied and interpreted (to build proteins).

Thus, it turns out that nature has solved the problem of combining information at least once before.⁹ Discrete genes were needed for a deep reason: only discrete genes can be combined in such a way that new organisms are produced which are different enough for natural selection to have something to operate on, but are alike enough to be able to successfully mate and producing viable offspring. (This way of putting the point is anthropomorphic, but not dangerously so.¹⁰)

We think the analogy between the genetic case and the cognitive case is quite robust. Discrete representations are also needed for a deep reason: only discrete (cognitive) representations can be combined in such a way that cognitive processes have different thoughts to operate on, but where the thoughts are similar enough to insure the coherence required for rationality.

There is another, short variant to this argument. It is well-known that concepts combine. The concept <u>brown</u> and the concept <u>cow</u> can combine to form <u>brown cow</u>. The concepts <u>snow</u> and <u>bicycle</u> can combine to form <u>snow bicycle</u> (whatever that might be) (see Smith & Osherson, 1984, and Wisniewski, 1997, for discussions of the psychology of conceptual combination). Yet, the original concepts once combined, are nevertheless recoverable. If concepts are represented in some continuous, fluid substrate then this ought not to be possible. When fluids combine, the original constituent fluids are lost, and recovering them is impossible (or at least very difficult). The fact that constituent concepts are easy to extract from combinations suggests that they do not combine in the manner that fluids do. Just as the blending theory of inheritance foundered on its prediction that heritable properties would lose their distinctiveness through blending (which did not occur), so too the view that most cognition (including higher cognition) is based on

⁹ And we are not the first to make the point that natural selection requires discretely represented information. Dawkins makes this point also. See his 1987, p. 115.

¹⁰ Some readers have questioned whether or not genes are really viewed as discrete in modern biology. The answer is: usually. Sometimes biologist do think of genes as dynamic, continuous structures, but this is only in special circumstances. In all the standard cases, genes are viewed as discrete entities.

continuous representations founders on its difficulty in explaining how concepts combine while keeping the constituents extractable.¹¹

The fourth argument turns on the fact that role-argument structures appear to be important for many cognitive processes in order to enable them to combine efficiently (Fodor & Pylyshyn, 1988; Markman, 1999; Norman & Rumelhart, 1975). Representations of complex entities, say, a coffee cup, are made up of simpler elements in specific relations. So, a representation is structured if and only if it is composed of parts; nothing that is homogenous is structured, and vice versa. One way of writing out structured representations is using predicate calculus notation with predicates that take arguments, for example:

cause [love (John, Mary), kiss (John, Mary)]

Continuous representations typically are not homogenous, yet they do not have parts. Continuous representations have *regions* that grade smoothly (even if sometimes rapidly) into one another. Regions aren't parts, because they cannot be extracted from the representation and replaced with other parts. Potentially, some other representation system could describe the regions of a continuous representation as parts, but that would require a discrete representational system that supervenes on the continuous representation. For example, in the thermostat example, the discrete representation (i.e., the mercury vial) supervenes on the continuous one (the bimetal strip).

¹¹ It might be objected that one's concepts <u>brown</u> and <u>cow</u> do not in fact combine, but rather that copies of them combine, and once these copies combine they do lose their ability to be uncombined or extracted. This objection doesn't work. In the first place, there is no evidence that concepts are copied for any purpose. Secondly, there is no evidence that copies of concepts are the ones used in combining concepts. Concept composition is an important cognitive process, but it doesn't appear to work only on copies. In most cases of conceptual change, the change appears to be permanent, which wouldn't happen if the change happened to copies (Dietrich and Markman, 2000b). Further support for the use of original concepts in combination comes from research on analogy. There is now rather good evidence that when two concepts interact to produce an analogy, the two concepts are changed somewhat (see Dietrich, 2000, and Gentner and Wolff, 2000). Most of the large changes happen after the analogy is made , and they clearly happen to the original concepts, since after the analogy, inferences made from those concepts take into account the information transferred in the analogy.

The curve in Figure 3 can be used to illustrate this point. Assume this curve describes the behavior of a continuous representational subsystem *R*. The graph can be construed as having two regions: ray A and ray B, related by a smooth transition, curve S, between them. So of course *R* can be construed the same way. In order to construe *R* in this way, another representational system (call it *T*) is required that takes the behavior of *R* as input and yields as output discrete representations whenever it (*T*) determines that *R* is in the regions defined by ray A, ray B. If *T* also construes the region S as the relation between rays A and B, then we can imagine *T* creating a structured representation derived from *R*. If *T* represents S with the symbol Σ , and ray A and ray B as the symbols 0 and 1 respectively, then *T* can yield a tree of the form (Σ 0 1) as a construal of *R*. In this example, a structured representation is constructed from parts associated with the relevant regions of *R*, but *R* itself does not have parts and hence is not structured. Therefore, if a representation is structured it has discrete parts, and those parts are discrete representations.

3.2. Arguments based on functional role

Arguments 5-7 focus on the relative contribution of correspondence and functional role in continuous and discrete representations. The fifth argument is based on the observation that humans are able to identify and use thousands of concepts that are tightly interconnected. Such connections define the functional roles of the various representations. But, the representations connected in this way have to be discrete. It is meaningless to talk about connecting a nondiscrete, continuous whole. Only discrete representations can *have* functional roles with respect to each other. Since there is significant evidence that people's concepts are highly interconnected and that concepts derive their meanings in part from these connections, it must be that the concepts are discrete.

The observation that people's categories are organized into hierarchies leads to the same conclusion. For example, the same object can be categorized as a <u>station wagon</u>, <u>car</u> and a <u>vehicle</u> (Brown, 1958; Rosch, Mervis, Gray, Johnson, & Boyes-Braem, 1976). People recognize the

semantic relationships among categories. They are also able to use these levels of categorization flexibly, although they tend to treat the middle level of abstraction (e.g., <u>car</u>) as cognitively privileged. It is difficult to see how people could keep these various levels of abstraction distinct using only a continuous representation, particularly one that shows little influence of connections among concepts. Hence, the basic facts of conceptual organization argue strongly for discrete representations.

The sixth argument begins by noting that we are capable of thinking so efficiently about so many things in our world because the mind is able to abstract from its environment. We don't have to have a concept for each individual table we have come across in our lives; we can instead group them all under the concept 'table.' Continuous representations are not abstract however. They are typically highly specific, because they correspond to some specific continuously varying input . In contrast, discrete representations are often abstract. In fact, creating abstractions produces discrete representations.

Not much is known about how concepts are constructed from the input we receive, indeed there is little agreement even on the best way to define concepts. There is general agreement, however, that concepts are abstractions. That is, they abstract away from information in the environment. It is this property that enables concepts to allow the cognitive system to treat discriminably different items as equivalent. Of course, the abstraction process is not well understood either, but there is general agreement that abstractions leave out information from the environment. Given just these agreed-upon facts, abstraction must involve (at least) the following steps:

- Step 1. Information (e.g., sound waves) is received continuously from the agent's environment.
- Step 2. At the initial stages of processing, continuous representations encode this information. Information about the physical stimulus (e.g., the sound wave) is lost when the representation is not sensitive to that information.

For example, the human auditory system loses information about highpitched sounds, though the canine auditory system does not.

Step 3. Information is extracted (by various other systems) from these continuous representations to produce representations at higher levels, e.g., the representations of words. This extracted information is not simply lifted out of the continuous representations like extracting a cupful of water from a mountain stream. Rather, the extraction process abstracts from the information in the continuous representations. This abstracted information leaves out details from the continuous representations. Hence, as higher representations are produced, additional information is lost. Losing information by interrupting a continuous stream produces a set of unconnected representations. These representations are used to discriminate conditions in the environment (e.g., to chunk a soundwave into words). Hence, abstraction produces discrete representations.

Another way to put this picture of the role of abstraction is to say that the benefit of losing information from continuous representations is the production of a set of discriminating, potentially referring, discrete representations that are combinable. And finally, the reason argument six is included under the functional role category is that abstractions have to get a significant part of their semantics from functional role. Because abstractions are not tightly connected to information in the environment, they cannot get much semantics from environmental information. Functional role is all that is left.

Now we are ready for our seventh and final argument. For some representations, the connections between them and what they represent can be described by physical laws. The arm angles in the Watt governor are connected to how open the steam valve is, the speed of the flywheel, etc. This connection can be described by a set of differential equations. We call these representations *nomic*. (We will also use this term when talking about the descriptions as well as the connections themselves.) Not all representations are nomic; some are nonnomic. A

nonnomic representation is one where there is no physical law governing or describing the relation between the representation and what it represents. A representation (concept) of a wornout climbing rope is an example of nonnomic representation. There are no physical laws governing the relation between having that representation and the climbing rope (c.f. Fodor, 1986).

Nonnomic representations are extremely useful. It is important to be able to think about things where the representation and what is represented are not connected by a law of physics. Indeed, most human concepts (as opposed to sensations) involve representational relations that are not physical laws. Hence, almost all higher-level cognitive representations are nonnomic.

Claim: only discrete representations are nonnomic.

This claim can be proven by contradiction, by considering whether there could be nonnomic continuous representations.

For all continuous representations, it is their *co*-variation with something in the input stream that makes them the representations that they are; their variations are how the representation's content is manifested or made explicit. Furthermore, as noted previously, the crucial aspect of continuous representations is that they are *coupled* to what they represent. Indeed this coupling is what the relevant physical law describes. Coupling guarantees that the variations in a continuous representation vary in accordance with the representation's precipitating conditions. If there was a continuous representation that was nonnomic, then it could vary independently of the conditions in the input to which it is supposedly coupled. Hence it wouldn't represent anything, and so it would not be a representation. Thus, there can't be any nonnomic continuous representations. Hence all nonnomic representations must be discrete.¹²

¹² Fodor (1986) argues that no system with only nomic internal states has mental representations. Hence for Fodor, nomic internal states are not mental representations (even in humans). Whether this position is

As this argument demonstrates, the correspondence part of the semantics of a continuous representation is sufficient for its activation. Thus, a physical law can govern the activity of a continuous representation because the presence of something in the environment (i.e., external to the system) is sufficient to activate the corresponding (continuous) representation.

In contrast, discrete representations are activated both by the presence of information in the environment, and also by their functional connections to other representational elements. For example, within a semantic network, a representation is activated as a result of the spread of activation from neighboring concepts. This functional role source of activation is governed by psychological factors, not physical ones. Thus, discrete representations have a source of activation that is not present for continuous representations.

4. Conclusions.

We have established that cognition requires discrete representations. This conclusion is based on the above seven arguments (summarized in Table 1). These arguments center around two main properties of discrete representations. First, discrete representations permit a system to discriminate among states in the environment, but continuous representations do not. Second, discrete representations allow a greater influence of functional role information than do continuous representations.

Accepting that discrete representations are central to cognitive processing has an interesting consequence. It follows from the presence of discrete representations in cognitive

compatible with our own depends on what 'mental' means. A restricted definition applying only to higherlevel cognition is fine, for then the nomic representations, while crucial to cognition are not themselves representations used by higher-level cognitive processes. A broader definition of 'mental' that covers all of cognition and the states on which cognition depends, is not compatible with our view. Unfortunately, Fodor does not make clear what he means by a mental representation.

processing that computationalism is the best paradigm for cognitive science. The argument for this point is relatively straightforward; it proceeds in two steps. Any system that uses discrete representations, is finite, and has deterministic transitions between the states can be rendered as an algorithm. We have argued that the mind uses discrete representation, and the mind is clearly finite (it stores only a finite amount of information, and it works at a finite speed). It is also plausible that the mind is deterministic (it's high variability due to the fact that its inputs are never exactly the same). Thus, the mind can be described as a computational system¹³.

So far, the argument has been apodictic. The second step, unfortunately, is not; it is a methodological variant of inference to the best explanation. Because the mind can be described as a computational system and because we understand computational systems rather well (but not completely), it follows (with some force) that computation is an extremely useful paradigm for theorizing about the mind, for theorizing about thinking. We ought to use the best, most powerful, explanatory paradigm we have available. Right now, that is computationalism. Hence, until something better comes along (and we don't anything will) computationalism ought to be the paradigm of choice among cognitive scientists.

Why have explanations of cognitive capacities proved so elusive if the field has found the proper paradigm? One important reason that many computationally-based explanations have fallen short is that these explanations often do not use true representations. As discussed above, in order to be a representation, there must be a set of internal states that bears some informational relationship to some input. In most computational models, however, the 'representations' are created by the user. These data structures have a *user semantics* because only the user knows what they mean. The model itself is given empty symbols (see also Bickhard & Terveen, 1995).

The field of situated action (sometimes called embodied cognition) has explicitly recognized this limitation in many models (Brooks, 1999; Clancey, 1997; Pfeifer & Scheier, 1999; Suchman, 1987). This approach criticizes mainstream research in cognitive science for failing to

¹³ Thanks to Dave Chalmers for discussion of this first step of the argument.

attend to the connection between the system and its environment. This research demonstrates that some problems that are quite difficult to solve using abstract representations can be greatly eased by taking the environment into account. For example, Pfeifer and Scheier (1999) describe a robot that classifies cylinders as large or small. The robot does not have a sophisticated visual system. Instead, it has simple procedures for wall following that enable it to circle around the cylinders. It also has sensors that detect the speed of its wheels. The ratio of the speed of the outer wheel to the speed of the inner wheel is larger for small cylinders than it is for large ones. By attending to this aspect of the environment, the robot is able to classify the cylinders successfully.

In many cases, it is possible to make claims about what can be computed with a representation simply by knowing the form of the representation. For example, Tversky (1977) demonstrated that models using spatial representations were not capable of accounting for patterns of human similarity judgments by demonstrating systematic violations of the metric axioms that define a space. On the basis of these computational concerns, it was reasonable to call the use of spatial representations into question (although Krumhansl, 1978, and Nosofsky, 1986, were able to extend the assumptions of spatial models to account for many violations of the metric axioms).

It is more difficult to argue for the sufficiency of a set of representational assumptions without taking into account how the representations come to have content. For example, the classic past tense learning model of Rumelhart and McClelland (1986) learned to associate an input pattern that could be interpreted (by the user) as representing the present tense form of a word with another pattern that could be interpreted as the past tense form of the same word. Many critiques of this model were generated that focused on the capacities of distributed representations (Fodor & Pylyshyn, 1988; Lachter & Bever, 1988). A significant limitation of this model, however, is that the representations were not sufficient to perform any linguistic operations, nor did the input to the model actually represent the words (see also Pinker &

Prince,1988). We do not mean to pick on this model specifically. It is simply a convenient example of a widespread problem in cognitive modeling.

There are a number of alternative approaches to modeling that can be adopted. Advocates of situated action argue that modeling can only be done with systems (or pseudoagents -- since such 'agents' don't, in fact, think) that have sensors and effectors tied to the environment. While this approach is fruitful, we believe that cognitive modeling can be done without first modeling perceptual and motor processes. One promising approach is to perform large-scale cognitive simulations and to use the same sets of representational assumptions across many models. For example, Forbus (2001) argues that computational models of analogical reasoning must be tested across many large problems rather than making isolated representational and processing assumptions for each individual situation. This approach deserves further exploration.

Finally, we believe the distinction between continuous and discrete representations is a case study for the importance of representational (and to that extent, an explanatory) pluralism in cognitive science. By advocating a paradigm shift, researchers using dynamic systems are implicitly assuming that there is only a single type of representation that is used by the mind (e.g., Thelen & Smith, 1994; van Gelder, 1998). What the discussion in this paper demonstrates is that discrete representations are necessary for cognitive processing. This is not an argument against the use of continuous representations. Indeed, in cases where internal states need to vary continuously with some quantity in the world (such as in the steam engine governor or the bimetal strip in a thermostat) continuous representations are quite useful. It is unlikely that a device as complicated as the mind uses only one type of representation (see Markman, 1999, and Markman & Dietrich, 2000b, for further discussion of this point).

In conclusion, cognitive science need not find a new paradigm. Computationalism is not only a reasonable approach to the study of mind, it is the best one. While cognitive scientists can take heart that they are operating within the right framework, some concessions must be made.

31

In particular, there are three ways that we as cognitive scientists must change the way we use representation. First, we must recognize that there will be many types of representations in cognitive systems, and should not take evidence of the sufficiency of a set of representational assumptions for one type of cognitive phenomenon as a demonstration that these representational assumptions are valid for all cognitive phenomena. Second, in order to balance this explanatory pluralism and to avoid generating a model for each individual task, we must make an effort to unify our representational assumptions across models in order to create representations that are sufficient to carry out the range of cognitive processes for which the representations have been posited. Finally, we must make an effort to use representations that actually represent, as opposed to representations that have only a user semantics. By following these guidelines, cognitive science will be poised to make headway on the difficult problems that remain to be solved.

References.

Barsalou, L. (1999). Perceptual Symbol Systems. Behavioral and Brain Sciences, 22(4), 577-660.

- Bechtel, W. (1998). Representations and cognitive explanations: Assessing the dynamicists' challenge in cognitive science. *Cognitive Science*, 22, 295-318.
- Bickhard, M.H., & Terveen, L. (1995). *Foundational issues in artificial intelligence and cognitive science: Impasse and solution*. Amsterdam: Elsevier Science Publishers.
- Brooks, R.A. (1999). Cambrian Intelligence. Cambridge, MA: The MIT Press.
- Bruner, J. S. (1957). Going beyond the information given. In J. S. Bruner (Ed.) *Contemporary Approaches to Cognition*. Cambridge, MA: Harvard Univ. Press.
- Burgess, C., & Lund, K. (1997). Modeling parsing constraints with high-dimensional context space. *Language and Cognitive Processes*, 12, 177-210.
- Burgess, C., & Lund, K. (2000). The dynamics of meaning in memory. In E. Dietrich & A. B. Markman (Eds.), *Cognitive Dynamics* (pp. 117-156). Mahwah, NJ: Lawrence Erlbaum Associates.
- Clancey, W.J. (1997). Situated cognition: On human knowledge and computer representations. New York: Cambridge University Press.
- Dawkins, R. (1987). The Blind Watchmaker. New York: W. W. Norton and Co.
- Dietrich, E. (1990). Computationalism. Social Epistemology, 4 (2),135-154.
- Dietrich, E. (1994). Thinking computers and the problem of intentionality. In E. Dietrich (ed.)
 Thinking Computers and Virtual Persons: Essays on the intentionality of machines. (pp. 1-34).
 San Diego, CA: Academic Press.
- Dietrich, E. (2000). Analogy and conceptual change, or You can't step into the same mind twice. In E. Dietrich and A. Markman (eds.) *Cognitive dynamics* (pp. 265-294). Mahwah, NJ: Lawrence Erlbaum Associates.
- Dietrich, E. and Markman, A.B. (2000a). Cognitive Dynamics: Computation and representation regained. In E. Dietrich and A.B Markman (Eds.) *Cognitive dynamics* (pp. 1-29). Mahwah, NJ: Lawrence Erlbaum Associates.

Dietrich, E. and Markman, A.B. (eds.) (2000b). *Cognitive Dynamics*. Mahwah, NJ: Lawrence Erlbaum Associates.

Dretske, F. (1981). Knowledge and the flow of information. Cambridge, MA: MIT.

Fodor, J. (1975). The Language of Thought. New York: Crowell.

- Fodor, J. (1986). Why paramecia don't have mental representations. In *Midwest Studies in Philosophy, vol. X, Studies in the philosophy of mind.* P. French, T. Uehling, and H. Wettstein (eds.). Minneapolis, MN: Univ. of Minnesota Press, 3- 23.
- Fodor, J. A., & Pylyshyn, Z. W. (1988). Connectionism and cognitive architecture: A critical analysis. *Cognition*, 28, 3-71.
- Forbus, K.D. (2001). Analogy in the large. In D. Gentner, K.J. Holyoak, & B. Kokinov. *The analogical mind: Perspectives from cognitive science*. Cambridge, MA: The MIT Press.
- Gentner, D., & Markman, A. B. (1997). Structural alignment in analogy and similarity. *American Psychologist*, 52(1), 45-56.

Gentner, D. and Wolff, P. (2000). Metaphor and knowledge change. In E. Dietrich & A.B. Markman (Eds.) *Cognitive dynamics* (pp. 295-342). Mahwah, NJ: Lawrence Erlbaum Associates.

- Horgan, T. and Tienson, J. (1996). Connectionism and the philosophy of psychology. Cambridge: MIT.
- Kelso, J.A.S. (1995). *Dynamic patterns: The self-organization of brain and behavior*. Cambridge, MA: The MIT Press.
- Kelso, J.A.S. and Scholz, J. (1985). Cooperative phenomena in biological motion. In H. Jaken (ed.)
 Complex Systems: Operational approaches in neurobiology, physics, and computers. (pp. 124-149) Heidelberg: Springer-Verlag.
- Krumhansl, C.L. (1978). Concerning the applicability of geometric models to similarity data: The interrelationship between similarity and spatial density. *Psychological Review*, 85(5), 445-463.
- Lachter, J., & Bever, T.G. (1988). The relationship between linguistic structure and associative theories of language learning: A constructive critique of some connecitionist learning models. *Cognition*, 28, 195-247.

- Landauer, T. K., & Dumais, S. T. (1997). A solution to Plato's problem: The Latent Semantic Analysis theory of acquisition, induction, and representation of knowledge. *Psychological Review*, 104(2), 211-240.
- Markman, A.B. (1999). *Knowledge representation*. Mahwah, NJ: Lawrence Erlbaum Associates.
- Markman, A.B., & E. Dietrich, (2000a). In defense of representation. *Cognitive Psychology*, 40, 138-171.
- Markman, A. and E. Dietrich (2000b). Extending the classical view of representation. *Trends in Cognitive Science* v. 4, n. 12, pp. 470-475.
- Markman, A.B., & Gentner, D. (1993). Splitting the differences: A structural alignment view of similarity. *Journal of Memory and Language*, 32(4), 517-535.
- Markman, A.B. & Gentner, D. (1996). Commonalities and differences in similarity comparisons. *Memory and Cognition*, 24(2), 235-249.
- Marr, D. (1982). Vision. New York: W.H. Freeman and Company.
- Medin, D. L., Goldstone, R. L., & Markman, A. B. (1995). Comparison and choice: Relations between similarity processing and decision processing. *Psychonomic Bulletin and Review*, 2(1), 1-19.
- Medin, D.L., Lynch, E.B., & Solomon, K.O. (2000). Are their kinds of concepts? Annual Review of Psychology, 51, 121-147.
- Meyer, D.E., & Schvaneveldt, R.W. (1971). Facilitation in recognizing pairs of words: Evidence of a dependence between retrieval operations. *Journal of Experimental Psychology*, 90(2), 227-234.
- Millikan, R. (1984). Language, thought, and other biological categories. Cambridge, MA: MIT.
- Neely, J.H. (1976). Semantic priming and retrieval from lexical memory: Evidence for facilitatory and inhibitory processes. *Memory and Cognition*, 4, 648-654.
- Norman, D. A., & Rumelhart, D. E. (1975). *Explorations in Cognition*. San Francisco: W.H. Freeman.
- Nosofsky, R.M. (1986). Attention, similarity, and the identification-categorization relationship. *Journal of Experimental Psychology: General*, 115, 39-57.

- Palmer, S. E. (1978). Fundamental aspects of cognitive representation. In E. Rosch & B. B. Lloyd (Eds.), *Cognition and Categorization* (pp. 259-302). Hillsdale, NJ: Lawrence Erlbaum Associates.
- Pfeifer, R., & Scheier, C. (1999). Understanding intelligence. Cambridge, MA: The MIT Press.
- Pinker, S. (1997). How the mind works. New York: Norton.
- Pinker, S., & Prince, A. (1988). On language and connectionism: Analysis of a parallel distributed processing model of language acquisition. *Cognition*, 28, 73-193.
- Pylyshyn, Z.W. (1980). Computation and cognition: Issues in the foundations of cognitive science. *Behavioral and Brain Sciences*, 3, 111-169.
- Rosch, E., Mervix, C.B., Gray, W.D., Johnson, D.M., Boyes-Braem, P. (1976). Basic objects in natural categories. *Cognitive Psychology*, 8, 382-439.
- Rumelhart, D.E., & McClelland, J.L. (1986). On learning past tenses of English verbs. In J. L.
 McClelland & D.E. Rumelhart (Eds.) *Parallel Distribute Processing: Volume 2* (pp. 216-271).
 Cambridge, MA: The MIT Press.
- Shepard, R. N. (1962). The analysis of proximities: Multidimensional scaling with an unknown distance function, I. *Psychometrika*, 27(2), 125-140
- Smith, E.E., & Osherson, D.N. (1984). Conceptual combination with prototype concepts. *Cognitive Science*, 8, 337-361.
- Suchman, L.A. (1987). *Plans and situated actions: The problem of human-machine communication*. New York: Cambridge University Press.
- Thelen, E., & Smith, L. B. (1994). <u>A Dynamic Systems Approach to the Development of Cognition and Action</u>. Cambridge, MA: The MIT Press.
- Tversky, A. (1977). Features of similarity. *Psychological Review*, 84(4), 327-352.
- van Gelder, T. (1995). What might cognition be, if not computation? *Journal of Philosophy*, 91, 345-381.
- van Gelder, T. (1998). The dynamical hypothesis in cognitive science. *Behavioral and Brain Sciences*, 21(5), 615-666.

Wisniewski, E.J. (1997). When concepts combine. Psychonomic Bulletin and Review, 4, 167-183.

Table 1. Seven arguments that cognitive systems require some discrete representations.

- 1. Cognitive systems must discriminate among states in the represented world.
- 2. Cognitive systems are able to access specific properties of representations.
- 3. Cognitive systems must be able to combine representations.
- 4. Cognitive systems must have some compositional structure.
- 5. There are strong functional role connections among concepts in cognitive systems.
- 6. Cognitive systems contain abstractions.
- 7. Cognitive systems require some non-nomic representations.

Figure 1. The Watt steam engine governor



Figure 2. A diagram of a simple thermostat.



Figure 3. A continuous representation offered as an objection to the claim that only discrete representations can discriminate environmental conditions. Time moves from left to right.

