Characteristics of Connectionist Knowledge Representation

Technical Report CS-91-147*

Bruce MacLennan

Computer Science Department
University of Tennessee, Knoxville

November 26, 1991

Abstract

Connectionism — the use of neural networks for knowledge representation and inference — has profound implications for the representation and processing of information because it provides a fundamentally new view of knowledge. However, its progress is impeded by the lack of a unifying theoretical construct corresponding to the idea of a calculus (or formal system) in traditional approaches to knowledge representation. Such a construct, called a simulacrum, is proposed here, and its basic properties are explored. We find that although exact classification is impossible, several other useful, robust kinds of classification are permitted. The representation of structured information and constituent structure are considered, and we find a basis for more flexible rule-like processing than that permitted by conventional methods. We discuss briefly logical issues such as decidability and computability and show that they require reformulation in this new context. Throughout we discuss the implications for artificial intelligence and cognitive science of this new theoretical framework.

1 The Importance of Connectionism

Our goal is to show that connectionist systems are not simply new techniques for clustering data nor methods for constructing associative memories but rather represent a fundamentally new concept of knowledge and of the representation

*Modified slightly for electronic distribution.

and processing of information. Traditional approaches to knowledge representation all make the *formalist assumption*: that knowledge may be represented by finite structures composed of discrete atomic symbols arranged in accord with a finite number of syntactic relations.[1]

This assumption is related to a long-standing belief that the discrete is more fundamental, primitive and basic than the continuous. This belief derives from certain *pragmatic invariants* of discrete symbol systems, namely, that they are *definite, reliable*, and *reproducible* (MacLennan, in press). By *definite* we mean that all tokens have a determinate type and all formulas have a determinate syntax; there are assumed to be no differences of opinion or difficulties of classification. By *reliable* we mean that small errors do not affect the use of a calculus, since in each step tokens can be replaced by "clean copies" — assuming that errors entering in a step are not so large as to change the type of a token or formula. By *reproducible* we mean that errors do not accumulate with successive copying of formulas, since each copy begins afresh with tokens that are clear instances of their type. These invariants hold only under ideal circumstances, but reality is often sufficiently close to the ideal that it is frequently useful to view certain systems as calculi (e.g., hand arithmetic algorithms, digital computers, formal reason, board games). These invariants also constitute the principal advantages of digital over analog technology.

Connectionist knowledge representation challenges the assumed priority of the discrete in several ways: First, it shows that continuous symbol systems satisfy a different set of pragmatic invariants, that in some contexts are more important than those of discrete symbol systems. Connectionist systems are characteristically *flexible, robust, adaptive*, and *responsive*.[2] By *flexible* we mean that connectionist representations are continuous, and — in simplest terms — permit adjustment between too much and too little. They tend not to be *brittle*, a common characteristic of calculi and traditional, discrete knowledge representations. By *robust* we mean that since connectionist processing is defined in terms of continuous functions, small errors tend to lead to small effects; connectionist systems do not respond discontinuously (and perhaps catastrophically) to errors. By *adaptive* we mean that, as a result of continuous representation, connectionist systems are able to change their behavior gradually, in contrast to calculi, which change their behavior discontinuously when rules are added or deleted. Finally, connectionist systems that operate continuously in time are *responsive*, because the solution to a problem is an asymptotic equilibrium, and so a partially completed computation may be usable in a time-critical situation. In traditional systems, partial inferences are often useless for action. For many applications flexibility, robustness, adaptivity, and responsiveness are more important than definiteness, reliability and reproducibility, and for these applications connectionist approaches may be more appropriate.

[1] The implications of this assumption are explored in detail by Pepper (1942, Ch. VIII).

[2] We are not claiming that every connectionist system has these properties, or that every connectionist researcher would agree to our characterization. We do claim, however, that the contribution of connectionism lies in exactly these properties.

A second way in which connectionist knowledge representation challenges traditional approaches is by showing that the continuous is just as fundamental, primitive and basic as the discrete, and in fact that either may be defined in terms of the other. We will further elaborate this theme later.

Thus one of the most important characteristics of connectionism is that is gives a fundamentally different way of viewing the representation and processing of knowledge. By forcing ourselves to view traditional problems from the continuous, connectionist viewpoint, as opposed to the familiar discrete, symbolic viewpoint, we will expand the universe of approaches to knowledge representation. By this "phenomenological variation" (Ihde, 1977) we will discover the range of approaches to knowledge representation and processing.

In addition to this, connectionist knowledge representation promises to illuminate information processing in the brain. In particular it will explain how human discrete (or approximately discrete) symbolic cognition can be implemented in terms of the slow, low-precision analog processing of the neurons. Connectionism also promises to provide a basis for *flexible information processing*. It is apparent that people can handle discrete symbols more flexibly than conventional computers. We hypothesize that this is because human symbolic cognition is implemented in terms of continuous subsymbolic processes, and so can partake of the flexibility of these processes when that is advantageous. We expect connectionism to explain the grounding of symbolic processes in the subsymbolic substrate (Smolensky, 1988).

2 Postulates of Simulacra

2.1 Introduction

Our understanding of connectionist knowledge representation is impeded by the lack of a general theoretical framework analogous to the idea of a *calculus* (or formal system) in traditional knowledge representation. In other words, every discrete, symbolic knowledge representation system is a calculus of some sort, and so by investigating the generic properties of calculi we come to understand the range of possible discrete, symbolic knowledge representations. This investigation has led to the important theorems of Turing, Gödel, Löwenheim, Skolem and others, which expose both the capabilities and limitations of discrete information representation and processing. A major goal of our research has been to identify a theoretical framework that similarly captures the basic assumptions of connectionism, nonpropositional-knowledge representation, and image-based cognition (MacLennan, 1988, in press). It is our hope that by outlining the range of possible connectionist knowledge representations, this theory will lead to an understanding of continuous information representation and processing that is comparable in depth to the theory of discrete formal systems.

In this paper we investigate a mathematical object called a *simulacrum*, which is proposed as the continuous analog of the discrete calculus. Just as a calculus defines a set of *formulas* (representations) and a set of *rules* for transforming them (processes), so a simulacrum defines a space of *images* and a space of *processes* for transforming the images.[3] Our goal is to characterize these spaces mathematically, but before we can do so, it is necessary to determine the *invariants* that we want the mathematics to capture. This process is described in detail elsewhere (MacLennan, in press), here we discuss only its results, a set of postulates for simulacra.

2.2 Postulates

We have identified a number of postulates that we believe to be satisfied by all or most simulacra. We will state each postulate formally, but also explain it in intuitive terms; a more detailed justification will be found elsewhere (MacLennan, in press).

Just as a calculus is constructed around one or more *formula spaces* — rep-

[3] A *simulacrum* can be defined as that which resembles something else, a likeness, an image, a representation. By metonymy it refers to the entire system, as 'calculus' (pebble, token) does to the entire discrete system. We considered a number of Greek and Latin terms (analogous to 'calculus'); they are listed here with their definitions and our evaluations — *Simulacrum*: "that which resembles something in appearance, sound, etc., a likeness;" "a visual representation, image" (*Oxford Latin Dict.*, s.v.). *Idolum* or *eidolon*: "An apparition, ghost; (in Epicurean philosophy) an 'image'." *Imaguncula*: "A small image, statuette" (a diminutive, like 'calculus'; but doesn't sound good). *Imago*: "A representation in art of a person or thing, picture, likeness, image;" "An image emitted by an object and apprehended by the eyes" (Epic. phil., translation of Greek ϵ ἴδωλον, eidōlon); "A representation to the imagination, mental picture;" (etc., meanings are quite good; but doesn't sound good). *Species*: "Something presented to view, a spectacle, sight;" "An act of looking," etc.; "Visual appearance, look, aspect" (meanings are not right, and it is too confusing with English word). *Effigies*: "An artistic representation, statue, portrait, etc.;" "A copy, reproduction;" "Outward appearance, aspect, shape, or guise" (meanings are not right, it may be confusing with English, and doesn't sound good). We also considered terms used in philosophy to describe images that would be somewhat analogous in usage to 'calculus'. Lucretius uses simulacra (a direct translation of ϵ ἴδωλα, *eidōla*, and τύποι, *tupoi*), imagines, effigiae and figurae (Loeb edition, p. 278). McKeon (*Sel. Med. Phil. II*, pp. 460–1) discusses 'imago', 'idolum seu fantasma'. Aristotle uses φάντασμα (*phantasma*), but the English connotations of this are apt to be misleading. Another set of possibilities comes from words for 'diagram', which, like calculus, also connotes a representation used for information processing. *Forma*: "Visible form, appearance, aspect;" "Shape, outline; (spec.) the shape of a thing as essential to the performance of its functions, a geometrical figure;" "Arrangement, pattern, configuration, or conformation" (these all suggest discrete as well as continuous symbol systems, as would be expected from words like 'formal', 'formula', etc.); 'forma geometrics' is used for a geometric diagram. *Figura* has a similar set of meanings to 'forma' and similar disadvantages. *Descriptio*: "The drawing of a diagram, plan, or sim[ilar];" "A descriptive narrative (written or spoken…)" (this suggests the English word 'description', which in turn suggests a language-like representation). Words for 'map' have similar information processing connotations, and are analogously concrete like 'calculus'. *Tabula*: although this may mean 'map' in context, its more usual meanings — board for games, counting board, inscribed tablet, account book, records, etc. — suggest discrete symbol systems. There are no other common words for 'map'. Thus, the best choice seems to be to use 'simulacrum' by metonymy for the systems and to reserve 'imago' or 'image' for the elements of these systems.

resentational spaces of well-formed formulas — so a simulacrum is constructed around one or more *image spaces*. The first two postulates characterize image spaces (or continua), which are the basic representational resources of simulacra.[4] We explain briefly how these postulates follow from several reasonable assumptions about images spaces.

The first such assumption is that image spaces are characterized by differing *degrees* of similarity between the images. If we make the further assumption that these degrees are quantifiable and that the quantification satisfies the formal properties of a *metric*, then we can conclude that an image space is a *metric space*.[5] Our second assumption gets at the essence of a continuum, for we assume that every image in the space is reachable from every other by a continuous process of transformation. Without this assumption, some regions of the image space would be essentially unreachable from others; in topological terms, the image space would be *disconnected* and hence not a continuum.[6] In contrast, the discrete formula spaces of calculi are *totally disconnected* (Hausdorff, 1957, p. 175), and this lack of connection is the essence of discreteness.

More formally, for any two images x and y, we require there to be a continuous function p (representing the transformation process) such that at "time" t_o, $p(t_\mathrm{o}) = x$ and at time t_f, $p(t_\mathrm{f}) = y$. In topological terms, the points x and y are *path-connected*, since p is a *path* from x to y.[7] Thus we have:

Postulate 1 *Image spaces are path-connected metric spaces.*

This postulate has several consequences that would be expected to hold for continua and image spaces, and so these consequences increase our confidence that this postulate is correct. We state them without proof; first:

Proposition 1 *For any two given images, there is always a third image closer to the first given image. That is, given images x and y, we can always find an image z such that $d(x, z) < d(x, y)$.*

This is a simple consequence of the continuity of any path connecting the given images. Its practical consequence is that arbitrarily small adjustments of images are possible, which supports the adaptability of connectionist systems.

[4]We generally avoid using the obvious term 'continuum' for image spaces, because that term's usual meaning in topology — a nontrivial connected compact metric space — is not equivalent to an image space. However, where confusion with topological usage is unlikely, 'continuum' is a useful and suggestive synonym for 'image space'.

[5]These assumptions are not entirely unproblematic; in particular it is not obvious that the triangle postulate, $d(x, y) + d(y, z) \geq d(x, z)$, can be taken for granted, and for some purposes a semimetric space may be the most we can safely assume (MacLennan, 1988).

[6]Needless to say, separate image spaces need not be connected with one another.

[7]We are being a bit informal here so that the intuitive justification of the postulates will be clear. First, note that we are using the term "continuum" informally, not in any of its (several) technical uses in topology. Second, there are spaces that are connected (in the topological sense), but not path-connected. We do not consider settled the question whether image spaces should be assumed to be path-connected or merely connected.

The second implication of the postulate, which is proved in Hausdorff (1957, p. 175), increases our confidence that image spaces are "big enough" to be continua:

Proposition 2 *Image spaces have at least the cardinality of the real numbers.*

Our second postulate for image spaces is less intuitive and more problematic than the first:

Postulate 2 *Image spaces are separable and complete.*

This means that we can construct convergent sequences of approximations to images, and that the limits of these sequences belong to the image spaces. The operational relevance of this postulate is not entirely clear, but it is mathematically convenient and a property expected of continua.

Additional support for this postulate is provided by a theorem of Urysohn, which shows that any metric space with a countable base is homeomorphic (i.e., topologically equivalent) to a subset of the Hilbert space $\mathcal{L}_2$ (Nemytskii and Stepanov, 1989, p. 324). (Note, however, that the homeomorphism need not preserve the metric.) Since any separable metric space has a countable base, any image space (by Postulate 2) is homeomorphic to a subset of the Hilbert space $\mathcal{L}_2$. For these reasons we claim:

Proposition 3 *Image spaces are Hilbert spaces.*

Thus, for most purposes we can treat image spaces as Hilbert spaces, which is very convenient, because Hilbert spaces have many useful properties and have been extensively studied. In addition to being a common framework for understanding vision and audition, Hilbert spaces seem to provide an attractive theoretical framework for neural computation (MacLennan, 1987, 1990; Pribram, 1991).

The formulas of a calculus are characterized by *syntactic relations*, which determine the types of formulas and part-whole relations between formulas, upon which depend in turn the rules of derivation. These syntactic relations can be expressed as maps between discrete spaces, either from the space of WFFs (well-formed formulas) to itself, or from the space of WFFs to a space of types. Analogously, the syntactic relations in simulacra are defined by maps, in this case, by continuous functions between image spaces. (The functions must be continuous if we are to preserve the pragmatic invariants.) Thus we have the postulate:

Postulate 3 *Maps between image spaces are continuous.*

We will see that this intuitive postulate has a very important consequence: exact classification is impossible in simulacra. That is, the "logic" of simulacra is inherently fuzzy.

Calculi normally provide *rules of derivation* that define discrete-time processes on the space of formulas. Analogously, simulacra normally provide continuous-time processes on the space of images. We require that these processes satisfy:

Postulate 4 *Formal processes in simulacra are continuous functions of time and process state.*

Finally we consider the semantics of simulacra. The preceding postulates deal with the *formal* properties of simulacra, those which depend on form but not meaning, and thus these postulates provide a framework for understanding *uninterpreted simulacra* (*continuous formal systems*). This is important for both cognitive science and artificial intelligence, since it shows us what can be computed without appeal to homunculi for interpreting the symbols. On the other hand, since our goal is to understand knowledge representation, we must also consider the relation between images and their meanings. This is addressed by our last postulate.

Just as the *domain of interpretation* of a calculus is necessarily discrete (this is the import of the Löwenheim-Skolem Theorem), so the domain of interpretation of a simulacrum is inherently continuous, that is, it is an image space. Further, just as the interpretation function of a calculus is required to be *compositional*, that is, to respect the constituent structure of the formulas, so we impose an analogous requirement for simulacra, namely, that their interpretations be continuous (which would, in any case, be expected from Postulate 3). Thus we have:

Postulate 5 *Interpretations of simulacra are continuous.*

It will be observed that not all connectionist and neural network systems satisfy all of these postulates. For example, neural activity may be drawn from a discrete space (typically binary or bipolar), or network activity may be updated at discrete time intervals. Nevertheless, we claim that these postulates capture the essence of connectionist knowledge representation and connectionist information processing. In any case, they are compatible with information processing in the brain, since neural processes are continuous.[8]

2.3 Idealization

In this connection it is important to observe that a simulacrum is an *ideal*, as is a calculus. That is, when we interpret a real, physical system as a calculus, we make certain *idealizing assumptions* that are not literally true. For example, we may assume that the voltage across a device can assume only two values (representing 0 and 1), whereas in fact the voltage must change continuously between these extremes, and may vary somewhat in each state. For another example, observe that when we interpret a system as a calculus we assume that the tokens can be unambiguously separated from the background and that they can be unambiguously classified as to syntactic type. The difficulties that may arise in real-world signal detection and pattern classification reveal the idealistic nature of these assumptions. There is nothing wrong with making

[8]This includes even apparently discrete phenomena such as the generation of action potential, since they result from the interaction of continuous physical laws.

these assumptions, so long as we acknowledge that we are doing so, and take care to ensure that the *real* system in question is sufficiently close to the *ideal*.

In an exactly analogous fashion, a simulacrum is an idealization of real, physical embodiments of continuous systems, and this idealization is unproblematic so long as the real and ideal are sufficiently close. Thus it matters little whether the image space is in fact complete, or whether the maps and processes are truly continuous, so long as these are good approximations. For example, it may be perfectly proper to treat an artificial neural net as a simulacrum even though activity levels are represented by rational or floating point numbers (which form a discrete space), or even though the state is updated at discrete time intervals.

3 Classification

3.1 Disallowed Connectionist Categories

It is a simple theorem of topology that there cannot be a continuous map from a connected space, such as an image space, into a nontrivial discrete space, such as $\{0, 1\}$. This result is of critical importance for the theory of connectionist information processing, for it means that connectionist categories are inherently fuzzy: their *classification functions* must pass continuously between 0 and 1.[9] The transition may be as sharp as we like, but it cannot be discontinuous. Thus we have:

Proposition 4 *Simulacra cannot classify exactly.*

We believe that this property of simulacra is biologically and psychologically very realistic, but several objections can be made to it:

Objection 1: It is factually incorrect, since humans and other animals do in fact distinguish disjoint classes. For example, suppose we say to a person, "I will show you a shape. If it looks more like a circle, then raise your hand; if more like a square, then don't." It will be argued that even though there is a continuum between circles and squares, the subject will either raise his hand or not, and thus exhibit exact classification (a map into a discrete set). We answer this objection as follows. Let C (for circle) represent the decision to raise the hand, and S (for square) the decision not to do so. Then certain images will be in the basin of attraction of C and others in the basin of attraction of S. But the basins of attraction are disjoint open sets (Hirsch and Smale, 1974, p. 190), and since a connected space cannot be the union of disjoint open sets, there must be some images that do not lead to either outcome. These images must either (1) lead to other equilibria, or (2) lead to periodic behavior, or (3) be stationary points, or (4) be unstable. In behavioral terms, the subject will

[9] By *classification function* we mean any map that reduces an image space to a linear continuum, such as $[0, 1]$. This is equivalent to the *characteristic function* of a set in the case where the function is identically 1 on the set and identically 0 on its complement.

either hover in a state of indecision or will act in an unpermitted way (e.g., run away). Of course, a real subject would not be permanently paralyzed by indecision (i.e., be in a stationary state), since changes in the context will alter the dynamics of the system and allow his state to change. Practically, the set of images leading to such anomalous behaviors is usually very small, but the mathematics tells us that it is nonempty.[10]

Objection 2: The proposition is irrelevant, since it allows arbitrarily sharp class borders, and so the categories are, from a practical standpoint, discrete. This objection misses the point. *Pragmatically*, a discontinuous border is no different from a continuous but arbitrarily sharp border: just as a calculus may approximate arbitrarily closely a continuous transition, but still moves in discrete steps, so a simulacrum may approximate arbitrarily closely a discontinuous transition, though it must still move in a continuous path. *Mathematically*, on the other hand, the difference may be very important: just as certain fundamental results (such as the Löwenheim-Skolem Theorem and perhaps Gödel's Theorem) follow from the inherent discreteness of calculi, so equally fundamental results may follow from the inherent continuity of simulacra. Such results would help distinguish fundamental properties of *any* representation of information, from those that are artifacts of our idealized models: calculi and simulacra. For example, if both models were subject to something like Gödel's incompleteness result, then we would have to take that to be a fundamental limitation of all information representation and processing; if not, then we might conclude that it is an artifact of an idealized model of discreteness, and of no fundamental importance for biological cognition (which is not mathematically discrete).

3.2 Allowed Connectionist Categories

Having found that simulacra do not permit exact categories, we must ask what kinds they do permit. Several kinds, generally corresponding to topological separation axioms, are easy to identify, and the resulting connectionist categories seem biologically and psychologically plausible. For example, for each open set in a metric space there is a continuous function that is positive on the set and nonpositive on its complement. This gives us a kind of fuzzy classification, since as we approach the border the value of the function must approach 0 (representing indeterminate classification). Thus we claim:

Proposition 5 *Any open set in an image space has a continuous classification function that is positive on just that set, but this function must approach 0 at the boundary of the set.*

Another kind of classification is permitted by Urysohn's Lemma, which says that for any disjoint closed subsets A, B of a normal topological space, there

[10] In many cases, such as gradient systems, the anomalous states will be so improbable that they constitute a set of measure 0.

is a continuous function that is identically 0 on on A and identically 1 on B.[11] Since metric spaces are normal, this result applies to image spaces, and gives a useful kind of classification. For example, if A is the set of images of things that are definitely artichokes and B is the images of definite bananas, then we are allowed a classification function that distinguished definite artichokes from definite bananas. On the other hand, since an image space cannot be the union of two or more disjoint closed sets (since that is the definition of a *disconnected* space), there must be other images in the space that belong to neither A nor B. For these images the classification function may return values between 0 and 1, and thus classify them indefinitely. In general terms:

Proposition 6 *For two disjoint closed subsets of an image space there is a continuous classification function that is identically 0 on one set and identically 1 on the other, but the classification function must take on each value between 0 and 1 on some image not in either set.*

Again we emphasize that the two closed sets could *nearly* exhaust the image space, and thus very nearly correspond to A and non-A, but the mathematics requires that there be an indefinite residuum between the definite categories.

Finally, for completely regular spaces (such as a metric spaces, and hence image spaces) we have this theorem: For every point and every open set containing that point, there is a continuous classification function that is 1 at that point and identically 0 outside the set. Thus we have the following classification result:

Proposition 7 *For every open set in an image space and for every fixed image (the* exemplar*) in that set, there is a continuous classification function that is 1 on the exemplar and identically 0 outside the set. However, the classification function must take on each value between 0 and 1 on some image in the set.*

We may interpret this kind of category as follows: We have an exemplar that is a definite case of the class (so the classification function returns 1), and there are images that are definitely excluded from the class (so the classification function returns 0), but for other images the classification may be indefinite (between 0 and 1). Clearly this theorem also permits a category to be structured around any finite number of exemplars.[12]

4 Structured Information

Fodor and Pylyshyn (1988) identify two characteristics that distinguish "classical" (i.e., discrete, symbolic) cognitive theories from connectionist theories: (1)

[11]Clearly, there is nothing special about the values 0 and 1; any two distinct points in a space would do.

[12]For an example, we may take Berlioz' *Symphonie Fantastique* as an clear exemplar of "romantic symphonies." Conversely, while there are obviously innumerable musical compositions that are definitely not romantic symphonies, there are many other compositions whose membership is not so clear (such as Beethoven's Ninth Symphony).

combinatorial syntax and semantics for mental representations, and (2) structure sensitivity of process. In other words, in classical but not in connectionist theories knowledge is represented by formulas with a complex constituent structure (as opposed to simple images) and is processed by rules sensitive to the structure of these formulas (as opposed to simple association). They further argue that since these characteristics are necessary for any adequate theory of cognition, connectionism cannot be adequate.[13] In this section we address these issues and show that connectionist knowledge representation provides a fundamentally different view of structured information and of its rule-like processing.

4.1 Constituent Structure

Defenders of traditional approaches to knowledge representation and inference have criticized connectionism for its inability to represent constituent structure (e.g., Fodor and Pylyshyn, 1988; Pinker and Prince, 1988), and as a result there has been considerable effort devoted to representing and processing trees in connectionist systems (e.g., Pollack, 1988, 1989; Dolan and Dyer, 1989; Dolan and Smolensky, 1989; Smolensky, 1987; MacLennan, in press). Although the results of these investigation have some use — especially in answer to the critics — we believe that connectionism has more to offer. In fact, we will go further, and assert that *a post-Chomskian linguistic theory will not be developed until we learn to think of language in some terms other than trees.* Unfortunately this is very difficult. Nevertheless, we believe that by forgetting trees we will reveal the possibility of non-graph-structured representations for linguistic structures in particular, and for cognitive structures in general. Unfortunately, we can only hint at what might be possible.

4.1.1 Decomposition of Images

To understand the nature of the connectionist representations of constituent structure, it will be helpful to review some characteristics of traditional discrete representations. First observe that discrete languages are defined *generatively*, that is, by specifying a processes, expressed in rules, for the construction of formulas from atomic constituents. Conversely in unambiguous languages (the usual case), formulas can be uniquely analyzed into their constituents. This means that the constructor functions are injective (one to one) and that their ranges are disjoint. In this way the construction of a formula is always uniquely determined.

In contrast, the definition of continuous languages (image spaces) is problematic. But if we consider some examples, such as defining a space of visual or auditory images, then it seems that circumscribing a background is one part of the definition. Thus we might specify a particular two-dimensional region for visual images or a particular range of frequencies for auditory images. Another

[13] We are of course grossly oversimplifying their 70 page argument.

part of the definition identifies the significant kinds and degrees of variation on this background. For example, we might specify a range of colors and intensities for visual images and of amplitudes for auditory images. In mathematical terms an image space is often simply a function space (e.g. $\mathcal{L}_2(\Omega)$), and specifying the background and range of variation amounts to specifying the domain and range of the image functions. In general the problem of defining continuous languages is open, but specific cases are easy.

The process we have outlined may seem informal, but it is no more so than the definition of calculi. Observe that calculi are defined by example, by exhibiting instances of the atomic types and by showing the outputs of the formation and transformation rules when applied to formulas containing special tokens (which we call *variables*). Such a process of ostensive definition must take place upon a background of assumptions about what is relevant and irrelevant to the definition. These assumptions are so familiar that they are no longer visible to us; the background has receded and is hidden from criticism. The ostensive definition of simulacra is also against a background of assumptions, but in this case they are less familiar, so they are visible and hence exposed to criticism.

Our discussion of the top-down definition of image spaces might seem to support the critics' contention that connectionist representations have no structure, but in fact it only implies that their structure is not *given*. Unlike (unambiguous) discrete formulas, continuous images need not have a unique structure. Analysis is usually problematic for images, since there are usually many ways in which they can be decomposed, and for some decompositions there may be no "bottom" to the analysis. Consider some of the possible decompositions of a visual image: pixels, edges, structural units, polygons, cylinders, etc. Consider also the many suggested "primitives" for representing images in visual cortex (oriented edges, differences of Gaussians, wavelets, Gabor elementary functions, etc.). These examples show that many image spaces — far from having no constituent structure — have an extremely rich constituent structure. We consider briefly some decompositions of image spaces.

4.1.2 Nonrecursive Decomposition

A decomposition of an image space is an analysis of that space into component spaces that are in some way more basic. The simplest case is when a higher-dimensional space is expressed as a product of two or more spaces of lower dimension. For example, the space of tongue positions in vowel articulation can be decomposed into a product of two linear continua (high/low, front/back). Recall however that an image space is not simply a set of images, it is also characterized by a metric; therefore it is not sufficient to express the set of given images as a Cartesian product of sets of simpler images; it is also necessary to show how the given metric can be expressed in terms of metrics on the component spaces. Thus, decomposing the image space $(\mathcal{X}, d)$ into spaces $(\mathcal{X}_1, d_1)$ and $(\mathcal{X}_2, d_2)$ entails showing the product decomposition of the set of images, $\mathcal{X} = \mathcal{X}_1 \times \mathcal{X}_2$, but also entails establishing a relation between the metric d and the component metrics d_1 and d_2. This cannot be assumed to be the obvious Euclidean (or ℓ_2) relation, $d^2[(x_1, x_2), (y_1, y_2)] = d_1^2(x_1, y_1) + d_2^2(x_2, y_2)$.

For real perceptual and motor spaces the isometric decomposition into lower dimensional spaces may be quite difficult to accomplish and often requires statistical techniques such as principal components analysis and multidimensional scaling (Shepard, 1980). This difficulty illuminates an important distinction between traditional and connectionist knowledge representation: *In calculi the structure of the formulas is manifest and unique; in simulacra it is hidden and manifold.* Although this can be interpreted as a disadvantage of connectionist knowledge representation, it is actually a source of useful representational richness. Instead of having a single distinguished decomposition, images have many decompositions, from which an appropriate analysis can be selected for the information-processing purpose at hand.

We have seen that image spaces can often be expressed as a product of two or more simpler image spaces. We turn our attention now to a special case, the decomposition of an image space into an infinity of simple (real) continua. One way to accomplish this follows from the Riesz-Fischer Theorem, which says that any separable Hilbert space, such as the space $\mathcal{L}_2(\Omega)$ of square-integrable (or "finite energy") functions on a domain Ω, is isomorphic and isometric to the Hilbert space ℓ_2 of square-summable sequences of real numbers. That is, infinite-dimensional images, such as visual and auditory images, can be analyzed into infinite sequences of reals (one-dimensional "images"). Further, the $\mathcal{L}_2$ metric on images is equivalent to the Euclidean metric on the sequences.[14]

The simplest example of such a decomposition is the familiar Fourier series in which an image is represented as the coefficients of a linear superposition of elementary images (sinusoids):

$$\phi = \sum_k c_k e_k,$$

where ϕ is the given image, the e_k are the elementary images, and the c_k are the Fourier coefficients of ϕ. More generally, if $(e_0, e_1, e_2, \ldots)$ is any basis (complete orthonormal system of images) for the image space, then an image may be represented by its *generalized Fourier coefficients* with respect to this basis. Moreover, the decomposition need not even be in terms of orthogonal images; for example, it is quite possible that the elementary images of the visual system are nonorthogonal Gabor wavelets (Daugman, 1988; MacLennan, 1991).

4.1.3 Recursive Decomposition

Supporters of traditional, discrete symbolic knowledge representation have criticized connectionist knowledge representation for an inability to represent recursive structures, such as trees (Fodor and Pylyshyn, 1988). Such structures

[14]The $\mathcal{L}_2$ metric on functions over a domain Ω is given by $d^2(x, y) = \int_\Omega [x(s) - y(s)]^2 \, \mathrm{d}s$. Of course, there is no requirement that the $\mathcal{L}_2$ metric be the metric on the image space, but if it is, then the decomposition is straightforward.

have been a mainstay of linguistic theory and are considered by some to be essential to a theory of human cognition. Although, some cognitive scientists have disagreed with this claim (Rumelhart, McClelland, et al., 1986, pp. 119–120) and, as we have argued above, the pervasive appeal to recursive structures may be more a limitation than a strength of traditional theories, it is nevertheless worthwhile to investigate connectionist approaches to recursive structure.

The basic problem can be explained as follows. The generative rules of most calculi are recursive, which means that the space of well-formed formulas is closed under the construction operations. In contrast, the state spaces of neural networks, which are usually taken to be finite-dimensional vector spaces, are not closed under construction operations. In intuitive terms, if we combine into a single vector two n-dimensional vectors we will get a $2n$-dimensional vector, not an n-dimensional vector. The impossibility of representing recursive structures in finite-dimensional spaces is implied by an important theorem of Brouwer's, which states that Euclidean spaces of different dimension are not homeomorphic. In other words, for finite n, we cannot have a continuous constructor operation $f : E^n \times E^n \rightarrow E^n$ with a continuous inverse $f^{-1} : E^n \rightarrow E^n \times E^n$. Thus we cannot continuously combine two finite-dimensional images into another of the same dimension.

On the other hand, recursive construction is possible for infinite-dimensional spaces (such as the Hilbert spaces $\mathcal{L}_2 = E^\infty$). We give one example. Suppose that the images $(e_0, e_1, e_2, \ldots)$ are a basis for an infinite-dimensional image space; then the images ϕ, ψ can be represented by their generalized Fourier series:

$$\phi = \sum_{k=0}^{\infty} c_k e_k, \quad \psi = \sum_{k=0}^{\infty} d_k e_k.$$

We can then construct the pair $\chi = (\phi, \psi)$ by interleaving their Fourier coefficients:

$$\chi = \sum_{k=0}^{\infty} c_k e_{2k} + d_k e_{2k+1}.$$

It is easy to show that this construction operation is a homeomorphism, since both it and its inverse are one-to-one and continuous (in fact, isometries).

If we consider the corresponding deconstruction operations (i.e., extracting the odd or even Fourier coefficients of an image) then we see that there is no "bottom" to the deconstruction of infinite dimensional images, a property we saw to be characteristic of many continuous representations. In contrast, there is always a bottom to the analysis of discrete formulas, namely, the atomic tokens. We take this to be another argument in favor of infinite-dimensional representations.

The representation of recursive structures is intimately related to the *competence/performance issue* in cognitive science (Rumelhart, McClelland, et al., 1986, pp. 119–120; Fodor and Pylyshyn, 1988) because it is apparent that human beings cannot handle arbitrarily deeply embedded structures. Thus cognitive science has traditionally distinguished an idealized theoretical *competence*, under which embedding is unlimited, from the observed *performance*, which is limited — and uninteresting to the theoreticians. Connectionism, in contrast, has tried to include performance issues in its theories from the beginning (Rumelhart, McClelland, et al., 1986, pp. 12, 130–131).

The competence/performance distinction appears very straight-forwardly in the current model. In an ideal continuous medium, which supports images with arbitrarily high-frequency components, trees of any depth can be constructed and deconstructed. On the other hand, real, physical media have limitations that interfere with the representation of higher frequency components, so we find that in practice the more deeply nested components of trees are the more degraded by noise and other mechanisms, and therefore are progressively less accessible. Such graceful degradation is biologically more robust than a hard limit imposed by, for example, a fixed stack depth.[15]

Representations such as our Fourier construction and those developed by other connectionist (e.g., Pollack, 1988, 1989; Dolan and Smolensky, 1989) show that it is possible to do LISP-style information processing (decomposing and recomposing symbolic structures) in neural networks, but this is just putting a fresh coat of paint on old, rotting theories. It would be better to avoid decomposition of formulas altogether and to understand how images are processed *as wholes*. A key problem is to show how such processing looks — to a first approximation — like processing that operates in terms of constituents.[16]

4.2 Rule-based Processing

Rules define structure-sensitive processes, that is, a rule applies to an image or formula that is in a given configuration or pattern, and rearranges the constituents of that pattern into a new configuration. We take a typical rule to have the form

$$A(x_1, \ldots, x_n) \Longrightarrow S(x_1, \ldots, x_n)$$

which means that an image constructed according to the pattern $A(x_1, \ldots, x_n)$ should be rearranged into the pattern $S(x_1, \ldots, x_n)$. The rule is *general*, because it applies for any (legal) constituents $(x_1, \ldots, x_n)$, and it is *structure-sensitive* because its applicability is conditional on the input matching the pattern represented by the constructor A. We will consider the processes defined by individual rules and finite rule-sets operating on image spaces.[17]

[15] What would be the behavioral effect of trying to comprehend a center-embedded sentence that exceeds our hypothesized stack depth? Would the too-deeply nested clauses be discarded completely, so that we would have no awareness of their content? Or would the outermost clauses be pushed out of the bottom of the stack, and so become lost? The actual effect is not so catastrophic; instead we find our comprehension to be *progressively* impacted.

[16] For example, rule-based operation might be a discrete approximation to continuous operation in terms of one of the classification propositions (e.g. Prop. 7).

[17] A more characteristically connectionist view would be that "rules" themselves belong to (infinite) continua and that they process information continuously, but we restrict our attention here to the finite/discrete case to simplify the comparison with traditional approaches.

We noted above that rules in calculi are specified in terms of general pattern constructors, which are defined ostensively by exhibiting "templates," instances of formulas containing tokens of a special type that are interpreted to be "variables." The convention for recognizing variables is part of the background of assumptions for the process of ostensive definition. Continuous pattern constructors can be defined ostensively in an analogous way. For a simple example, suppose the constructor $A(\phi)$ represents placing the image ϕ in a frame in a black and white image For the sake of the example, we assume that the only part of ϕ that will appear in the result is that part in the same place as the frame. The constructor A can be defined ostensively by adopting a convention for indicating the variable part of the image; for example it could be marked by an otherwise unused color or texture.

So far the ostensive definition of constructors in simulacra has seemed similar to that in calculi, but now we find some differences. Because images are continuous, the variable portion of the template must blend continuously with the fixed portion. Or, to put it another way, in constructing a pattern, we must classify regions of the template as "variable" or "fixed," but we have seen (Proposition 4) that exact classification is impossible. Therefore, let μ represent a possible classification function (e.g., allowed by Propositions 5, 6 or 7) which is 0 on the definitely fixed part of the template and 1 on the definitely variable part, but between 0 and 1 on the indefinite parts. Then the constructor $A : \mathcal{E} \rightarrow \mathcal{I}$ can be defined by a formula such as this (there are of course many others):

$$A(\phi) = (1 - \mu)\alpha + \mu\phi.$$

That is, if $A(\phi) = \psi$ then

$$\psi(p) = [1 - \mu(p)]\alpha(p) + \mu(p)\phi(p).$$

Thus, although part of the result is fixed and comes from the template α, and part is variable and comes from ϕ, there is also a unavoidable boundary region that blends α into ϕ. It is easy to show (with the $\mathcal{L}_2$ metric) that constructors of this kind are continuous (and, indeed, satisfy a Lipschitz condition) and one-to-one. Furthermore, if we assume that the allowable component images ϕ are restricted so that they are nonzero only where that classification function μ is greater than some fixed threshold θ, then we find that the inverse constructor A^{-1} is also continuous, and so A is a homeomorphism between $\mathcal{E}$ and $A[\mathcal{E}] = \{A(\phi) | \phi \in \mathcal{E}\}$.[18]

We have seen that a constructor for a complex image is a continuous injective (one to one) mapping from one image space into another (possibly the same as the first). For simplicity we restrict our attention to one-place constructors (i.e., $n = 1$); the extension to more places is routine. So suppose $A : \mathcal{E} \rightarrow$

[18]The requirement for the threshold θ can be understood as follows. If ϕ is allowed to be nonzero wherever μ is nonzero, then as $\mu(p) \rightarrow 0$ the product $\mu(p)\phi(p)$ in the definition of A allows arbitrarily large changes in ϕ to be canceled by small values of μ. Conversely, then, $\phi = A^{-1}(\psi)$ can magnify small changes of ψ into arbitrarily large changes in ϕ.

$\mathcal{I}$ is a constructor defined on image space $\mathcal{E}$; thus $A(\phi)$ is the image in $\mathcal{I}$ constructed from the image $\phi \in \mathcal{E}$. As suggested above we also require A^{-1} to be continuous.[19]

Now a rule $A(x) \Longrightarrow S(x)$, where $S : \mathcal{E} \rightarrow \mathcal{J}$, defines a function $R_o : A[\mathcal{E}] \rightarrow \mathcal{J}$ defined by $R_o = S \circ A^{-1}$. That is, if ψ is an image in the range of A, $\psi = A(\phi)$, then $R_o(\psi) = S(\phi)$. Since S and A^{-1} are both continuous and one-to-one, so is the rule R_o.[20] Notice, however that since R_o is defined only on the patterns constructed by A, it is in general only *partial* on $\mathcal{I}$. Thus we must extend R_o to a (continuous) function $R : \mathcal{I} \rightarrow \mathcal{J}$ defined on the entire image space.

There are many ways this extension may be accomplished — we give an example below — but they all result in some ability to generalize. For example if $\psi \in \mathcal{I}$ is an image not constructed by A, and thus not strictly in the domain of applicability of the rule, but if ψ is close to $A(\phi)$ then the continuity of R requires that $R(\psi)$ be close to $S(\phi)$. In other words, the apparently exact rule has an inevitable penumbra of applicability to images that do not exactly match its pattern.

Next consider the process defined by a rule-set comprising two rules:[21]

$$A(x) \Longrightarrow S(x), \qquad B(x) \Longrightarrow T(x)$$

If this process is to be deterministic, then there must be no images to which both rules are applicable; that is, there is no ψ such that $A(\phi) = \psi = B(\phi')$, since otherwise we could have inconsistent results $S(\phi)$ and $T(\phi')$. Therefore we assume that $A[\mathcal{E}]$ and $B[\mathcal{E}]$ are disjoint. The rules individually define mappings $R_1 = S \circ A^{-1} : A[\mathcal{E}] \rightarrow \mathcal{J}$ and $R_2 = T \circ B^{-1} : B[\mathcal{E}] \rightarrow \mathcal{J}$. Therefore the rule set defines a mapping $R_o = (S \circ A^{-1} \cup T \circ B^{-1}) : (A[\mathcal{E}] \cup B[\mathcal{E}]) \rightarrow \mathcal{J}$, which must be extended continuously to $R : \mathcal{I} \rightarrow \mathcal{J}$.[22] The effect of this continuous extension is that as an image varies from one matching pattern A, say $A(\phi)$, to one matching pattern B, say $B(\phi')$, the output of the process will vary from $S(\phi)$ to $T(\phi')$. Thus the postulates of simulacra imply that any rule set provides smooth transitions from the applicability of each rule to the applicability of the others.

We consider briefly one specific kind of continuous extension. Suppose the patterns A and B are continuously transformable into one another. For example, $A(\phi)$ might represent putting a circle around ϕ and $B(\phi)$ might represent putting a square around it. Or $A(\phi)$ might represent putting a minus sign in front of ϕ and $B(\phi)$ might represent putting a plus sign behind it. In both cases it is clear that in general $A(\phi)$ is continuously transformable into $B(\phi)$, for any

[19]Since A is a constructor, A^{-1} is a *deconstructor* or *decomposition*. It performs an *analysis* complementary to the *synthesis* performed by S. Thus, both analysis and synthesis must be continuous.

[20]In general a rule need not be one-to-one, since it may discard variables, for example, $A(x, y) \Longrightarrow S(x)$. If a rule doesn't discard variables, then it is a homeomorphism between $A[\mathcal{E}]$ and $S[\mathcal{E}]$.

[21]The extension to more then two rules is routine.

[22]For convenience of exposition we interpret the functions as sets of ordered pairs (i.e., we identify the functions and their *graphs*).

ϕ. Therefore we introduce a real parameter θ that represents the distance along any fixed transformation so that $\theta = 0$ represents A and $\theta = 1$ represent B. Specifically, we define a higher order pattern $C : [0,1] \times \mathcal{E} \rightarrow \mathcal{I}$ so that $C(\theta, \phi)$ varies continuously from $C(0, \phi) = A(\phi)$ at $\theta = 0$ to $C(1, \phi) = B(\phi)$ at $\theta = 1$. For simplicity we assume C is surjective (onto), so every $\psi \in \mathcal{I}$ can be written $\psi = C(\theta, \phi)$. Now the meaning of the rule set can be defined as a composition $R = D \circ C$ where D varies continuously between the patterns S and T, for example,

$$D(\theta, \phi) = \theta S(\phi) + (1 - \theta) T(\phi).$$

(This presumes that the image space $\mathcal{J}$ is a linear space, for example a Hilbert space.) Thus an input of the form $A(\phi)$ will lead to an output of the form $S(\phi)$, and an input $B(\phi)$ will lead to $T(\phi)$, but other images in $\mathcal{I}$ will lead to some mixture of the S and T patterns. In particular, if $\psi = C(\theta, \phi) \in \mathcal{I}$, then $R(\psi) = \theta S(\phi) + (1 - \theta) T(\phi)$.

More generally, we draw the parameter θ from a continuum $\mathcal{H}$ (with metric d) so that $C : \mathcal{H} \times \mathcal{I} \rightarrow \mathcal{J}$ is surjective. If, without loss of generality, we assume that $A(\phi) = C(\alpha, \phi)$, $B(\phi) = C(\beta, \phi)$, $d(\alpha, \beta) = 1$, then an appropriate D is:

$$D(\theta, \phi) = d(\theta, \beta) S(\phi) + d(\theta, \alpha) T(\phi).$$

The result is that the rule set is applicable to all images in $\mathcal{I}$, and the result of applying it to an image not strictly in the domain of any rule will be a kind of interpolation of rule outputs for nearby images. The quality of the generalization depends, of course, on the continuous extension that is chosen.

The connectionist rules we have described are defined analogously to discrete rules. In this sense they are not consistently connectionist. More typically in a connectionist system individual rules are not glued together to form a rule set; instead the "rule set" emerges from a uniform background transformation by some adaptive process (e.g. learning). Thus in some cases approximately discrete rules may emerge from this process, but connectionism also permits the inferential processes to take the form of a continuum that is not expressible as rules. This is a significant strength of connectionism both for cognitive modeling and for implementing flexible information processing.

5 Decidability

Since simulacra provide a fundamentally different framework for computation from calculi, the question naturally arises of whether they are subject to the corresponding undecidability and noncomputability results (e.g., Gödel's and Turing's). There is reason to believe that they are not, since simulacra permit inference rules that are effectively infinite, and some systems with inference rules of this kind are both ω-consistent and complete (Heijenoort, 1967, p. 355). This question deserves further investigation.

Although decidability questions of this kind have been addressed by others who have investigated various forms of continuous computation (e.g., Blum, 1989; Blum, Shub and Smale, 1989; Stannett, 1990), we believe that their approach has not taken a fully consistent continuous perspective, for they have approached the problems of continuous computation from the perspective of discrete symbol systems. That is, they have asked what a *calculus* can decide or compute about a continuous computational system. To ask questions corresponding to Gödel's and Turing's, we must ask what a *simulacrum* can decide or compute about *simulacra.*

When we take a fully consistent continuous perspective, then we find that we cannot even ask the same questions. In particular, since a simulacrum cannot do exact classification, it *a priori* cannot answer a yes-or-no decidability or computability question. Thus the questions must be reformulated before they can be addressed. For example, we could ask about the existence of a simulacrum that classifies a certain subset of the images as definitely decidable and another subset as definitely undecidable, so long as these subsets do not exhaust the space (i.e., answer in accord with Proposition 6). Further discussion can be found in MacLennan (in press).

6 Conclusions

We don't claim this is the definitive and final statement of the properties of simulacra or of connectionist knowledge representation systems. First, it is unlikely that there is one set of postulates that applies to all systems, just as there is not a single characterization of calculi (some are deterministic, others nondeterministic; some have a finite number of atomic types, others an uncountable infinity of them; etc.). Second, we have much less experience with simulacra — as a class — than with calculi, so it is to be expected that our first enumeration of postulates will require revision.

Nevertheless we think that it is clear that connectionism provides a fundamentally new view of knowledge and of information representation and processing. This will extend the range of our theories of both natural and artificial intelligence, and open new possibilities in classification, linguistics, structured knowledge representation, rule-like processing and decidability. Our intent is not to claim that the continuous perspective is better or more fundamental than the discrete, or that the discrete view should be abandoned. Rather it is to argue that the continuous be given equal time. Over the past 2500 years we have identified knowledge and cognition with calculi. Connectionism offers an essentially different perspective, the view that knowledge and cognition are simulacra.

References

[Blum, 1989] Blum, Lenore. "Lectures on a Theory of Computation and Complexity over the Reals (or an Arbitrary Ring)." Report TR-89-065, International Computer Science Institute, December 22, 1989.

[Blum, Shub and Smale, 1989] Blum, Lenore; Shub, Mike; and Smale, Steve. "On a Theory of Computation and Complexity over the Real Numbers: NP Completeness, Recursive Functions and Universal Machines." *The Bulletin of the American Mathematical Society*, Vol. 21, No. 1 (July 1989), pp. 1–46.

[Daugman, 1988] Daugman, J. G. "An Information-theoretic View of Analog Representation in Striate Cortex." In: *Computational Neuroscience*, E. L. Schwartz (ed.), Cambridge, MA: MIT Press, 1988, pp. 403–423.

[Dolan and Dyer, 1989] Dolan, Charles P.; and Dyer, Michael G. "Parallel Retrieval and Application of Conceptual Knowledge." *Proceedings of the 1988 Connectionist Models Summer School*, San Mateo, CA: Morgan Kaufmann, 1989, pp. 273–280.

[Dolan and Smolensky, 1989] Dolan, Charles P.; and Smolensky, Paul. "Implementing a Connectionist Production System Using Tensor Products." *Proceedings of the 1988 Connectionist Models Summer School*, San Mateo, CA: Morgan Kaufmann, 1989, pp. 265–272.

[Fodor and Pylyshyn, 1988] Fodor, J. A.; and Pylyshyn, Z. W. "Connectionism and Cognitive Architecture: A Critical Analysis." In: *Connections and Symbols*, Steven Pinker and Jacques Mehler (eds.), Cambridge, MA: MIT Press, 1988, pp. 3–71; reprinted from *Cognition*, Vol. 28 (1988).

[Hausdorff, 1957] Hausdorff, Felix. *Set Theory*, tr. John R. Aumann et al. New York, NY: Chelsea, 1957.

[Heijenoort, 1967] Heijenoort, J. van. "Gödel's Theorem." In: *The Encyclopedia of Philosophy*, Paul Edwards (ed.), New York, NY: Macmillan, 1967, Vol. 3, pp. 348–357.

[Hirsch and Smale, 1974] Hirsch, Morris W.; and Smale, Stephen. *Differential Equations, Dynamical Systems, and Linear Algebra*. San Diego, CA: Academic Press, 1974.

[Ihde, 1977] Ihde, Don. *Experimental Phenomenology: An Introduction*. New York, NY: Capricorn Books, G. P. Putnam's Sons, 1977.

[MacLennan, 1987] MacLennan, B.J. "Technology Independent Design of Neurocomputers: The Universal Field Computer." *Proceedings, IEEE First International Conference on Neural Networks* , 1987, Vol III, pp. 39–49.

[MacLennan, 1988] MacLennan, Bruce. "Logic for the New AI." In: *Aspects of Artificial Intelligence*, James H. Fetzer (ed.). Dordrecht, Holland: Kluwer Academic Publishers, 1988, pp. 163–192.

[MacLennan, 1990] MacLennan, B. J. "Field Computation: A Theoretical Framework for Massively Parallel Analog Computation, Parts I–IV." Technical report CS-90-100, University of Tennessee, Knoxville, Computer Science Department, February 1990.

[MacLennan, 1991] MacLennan, Bruce. "Gabor Representations of Spatiotemporal Visual Images." Technical report CS-91-144, University of Tennessee, Knoxville, Computer Science Department, September 1991.

[MacLennan, in press] MacLennan, Bruce. "Continuous Symbol Systems: The Logic of Connectionism." *Neural Networks for Knowledge Representation and Inference*, Daniel S. Levine and Manuel Aparicio IV (eds.). Hillsdale, NJ: Lawrence Erlbaum, in press.

[Nemytskii and Stepanov, 1989] Nemytskii, V. V.; and Stepanov, V. V. *Qualitative Theory of Differential Equations*. New York, NY: Dover, 1989; previously published, Princeton, NJ: Princeton University Press, 1960.

[Pepper, 1942] Pepper, Stephen C. *World Hypotheses: A Study in Evidence*. Berkeley, CA: University of California Press, 1942.

[Pinker and Prince, 1988] Pinker, S.; and Prince, A. "On Language and Connectionism: Analysis of a Parallel Distributed Processing Model of Language Acquisition." *Cognition*, Vol. 28 (1988), pp. 73–193.

[Pollack, 1988] Pollack, J. B. "Recursive Auto-Associative Memory: Devising Compositional Distributed Representations." In: *Proceedings of the Tenth Annual Conference of the Cognitive Science Society*, 1988, pp. 33–39.

[Pollack, 1989] Pollack, J. B. "Implications of Recursive Distributed Representations." In: *Advances in Neural Information Processing Systems*, David S. Touretzky (ed.), pp. 527–536. San Mateo, CA: Morgan Kaufman, 1989.

[Pribram, 1991] Pribram, Karl H. *Brain and Perception: Holonomy and Structure in Figural Processing*. Hillsdale, NJ: Lawrence Erlbaum, 1991.

[Rumelhart, McClelland, et al., 1986] Rumelhart, David E.; McClelland, James L.; and the PDP Research Group. *Parallel Distributed Processing: Explorations in the Microstructure of Cognition, Volume 1: Foundations*. Cambridge, MA: MIT Press, 1986.

[Shepard, 1980] Shepard, Roger N. "Multidimensional Scaling, Tree-Fitting, and Clustering." *Science*, Vol. 210, pp. 390–398 (24 October 1980).

[Smolensky, 1987] Smolensky, Paul. "On Variable Binding and the Representation of Symbolic Structures in Connectionist Systems." Technical report CU-CS-355-87, Department of Computer Science, University of Colorado at Boulder, 1987.

[Smolensky, 1988] Smolensky, Paul. "On the Proper Treatment of Connectionism." *Behavioral and Brain Sciences*, Vol. 11 (1988), pp. 1–74.

[Stannett, 1990] Stannett, Mike. "X-Machines and the Halting Problem: Building a Super-Turing Machine." *Formal Models of Computing*, Vol. 2 (1990), pp. 331–341.