Journal of Experimental & Theoretical Artificial Intelligence,
Vol. 17, No. 4, December 2005, 317–342

From apples and oranges to symbolic dynamics: a framework for conciliating notions of cognitive representation

RICK DALE* and MICHAEL J. SPIVEY

Department of Psychology, Cornell University, Ithaca, NY 14853, USA

(*Received 3 May 2005; accepted 26 June 2005*)

We introduce symbolic dynamics to cognitive scientists with the aim of furthering constructive debate on representation. Symbolic dynamics is a mathematical framework in which both continuous and discrete states of a system can be considered jointly. We discuss a number of theoretical implications this framework has for cognitive science, and offer some consideration of the way in which it might be employed for comparing or conciliating discrete and continuous representational theories. Symbolic dynamics may thus serve as a common, level playing field for debate in theories of cognitive representation.

Keywords: Dynamical systems; Symbolic dynamics; Cognition; Representation

1. Introduction

Since its inception, cognitive science has offered up a wide array of hypothetical constructs, intervening somewhere between our sensors and our effectors, to explain our observable behaviour. Many of these constructs can be filed under the umbrella term 'representation'. Representations might 'stand for' things in the world (Bloom and Markson 1998), be asymmetrically dependent with worldly objects (Fodor 1987), they might get stored or processed or recalled (Atkinson and Shiffrin 1968), and they surely change somehow during development and learning (Danovitch and Keil 2004). This generic construct has appeared and reappeared in a variety of forms, labelled variously with the terms 'traces' (e.g. Rosen 1975), 'schemata' (e.g. Bartlett 1932, Neisser 1976), 'categories' (e.g. Rosch 1975), 'concepts' (e.g. Medin 1989), 'object files' (e.g. Feigenson and Carey 2003), and so on—perhaps describable as different *forms* of representation.

There is no single agreed upon theory or definition of representation among cognitive scientists (Dietrich and Markman 2003). The details about any particular

*Corresponding author. Email: rad28@cornell.edu

Journal of Experimental & Theoretical Artificial Intelligence
ISSN 0952–813X print/ISSN 1362–3079 online © 2005 Taylor & Francis
http://www.tandf.co.uk/journals
DOI: 10.1080/09528130500283766

brand of representation are mostly specific to the theory in which it plays a role—but each brand can be characterized in terms of some basic features. Nonetheless, even these most fundamental properties of representation are the subject of continuing debate in cognitive science. One such property concerns the temporal and spatial extent of representational states. There are two basic sides to the traditional version of this debate. One family of theories may be described as 'discrete-symbolic', because they claim that internal representational states involve discrete computational information structures that are manipulated in logical algorithmic processes. Here, 'computational' can be understood intuitively as structures that take the form of something a digital computer would process—content that is discrete in space and time. A competing family of theories may be described as 'continuous-distributed', because they instead invoke representational states that are spread out in space, and extended in time. These states are graded, statistical and probabilistic—they cannot be individuated discretely in time, or uniquely in their informational content. Continuous-distributed representations contain probabilistic informational patterns that might blend into other such representations, whereas discrete-symbolic representations are by definition independent uniquely identifiable states that are each separate from, yet used in conjunction with, other discrete representations (Dietrich and Markman 2003).

Debate continues about which composition is the most appropriate foundation for cognitive explanation. Recently, Dietrich and Markman (2003, Markman and Dietrich 2000) have offered persuasive arguments about the crucial role of symbolic representation in higher-order cognition, such as conceptual organization, problem solving, and language (see also Pinker 1997, Marcus 2001). Meanwhile, Spivey and Dale (2004) argue that a continuous composition is extensively evidenced throughout even complex cognition, offering examples from real-time language processing and visual cognition (see also Port and Van Gelder 1995, Elman *et al*. 1996). In this article, we provide a review and discussion of a mathematical terrain in which these two representational formats can be directly compared and evaluated. We suggest that a kind of 'mathematization' of the problem space, in terms of nonlinear dynamical systems and symbolic dynamics, can aid in a variety of ways. The descriptive power of dynamical systems, and the computational power of symbolic dynamics based on them, can reveal an epistemological synthesis of this debate, and offer an illuminating framework for exploring such conceptual conciliation. For reasons that we describe later, we do not expect the framework of symbolic dynamics to make moot the debate between discrete and continuous descriptions of mental activity, but rather may pose as the level playing field on which the debate may actually achieve a consensual resolution.

A case may be made for mathematization of scientific domains as a course toward resolving theoretical disputes, clarifying conceptual confusions, and making potential decisions concerning the greater validity of one verbalized scientific description over another (e.g. for a discussion of this in psychology, see Meehl 1998). What early calculus did for Newtonian mechanics, tensor calculus for general relativity, symbolic logic for computer science, among other possible examples, is to provide a formal framework for exploring the relationships among observables, thereby making explicit predictions that can be tested empirically. A common mathematical framework, within which different theories compete, permits more rigorous evaluation of hypotheses that otherwise would be couched merely in a verbalized form.

For example, despite the growing popularity of quantitative models, it is not difficult to find theories in cognitive psychology whose existence enjoys only verbal description. This certainly does not invalidate the potential contribution of these theories – observers should always be reminded of the youth of cognitive science. The bigger problem is that multiple competing theories in verbal form may be conducive to debate with little chance of resolution. For one thing, without a common formal framework, it may be difficult to tell if two competing theoretical schemes are in fact mutually exclusive, or perhaps even extensionally equivalent. In other words, without explicit formulation of the relationship among theoretical entities, in more or less formal terms, it may be difficult to determine whether two competing entities are two distinct incommensurable accounts, two different aspects of one process, or merely two different descriptions of the same process. Secondly, theorists of differing persuasions may be talking past one another, preventing a Hegelian 'thesis–antithesis–synthesis' resolution that may be revealed by a common framework permitting conciliation of competing notions.

There are pursuits in cognitive science that benefit from aspects of formalization at present. For example, connectionist models have been used as a common information-processing framework for evaluating competing theoretical accounts of cognitive processes involved in language. McRae *et al.* (1998) used a localist attractor network to compare directly the immediate information-integration predictions from the constraint-based theory of sentence processing (MacDonald *et al.* 1994) to the architecturally delayed integration predictions from the Garden-Path theory (Frazier 1995). Drawing from work by Elman (1990) and Schütze (1994), Spivey-Knowlton and Saffran (1995) used a connectionist-like framework to directly compare the advantages of incremental prediction and explicit negative evidence in learning a simple language structure. Also, particularly relevant to our concerns here, homogeneous versus hybrid simulations using connectionist principles have recently been developed to compare dual- vs. single-route models of reading (Harm and Seidenberg 1999, Coltheart *et al.* 2001). Other formal frameworks, such as Bayesian modelling (Tenenbaum and Griffiths 2001), genetic algorithms (Chater *et al.* 2004), and statistical models of sentence processing (Chater and Redington 1996) have been manipulated in ways that allow comparison of competing theories.

The overarching theoretical concern, however, is that many of these models involve too many degrees of freedom to make them a sufficiently agreed-upon common ground for comparing theoretical constructs whose properties are highly disparate. Comparing individual models of particular processes is surely valuable and inevitable, yet fundamental theoretical differences in cognitive science cannot be contacted through manipulating already-existing architectures.† The debate over representation is particularly illustrative in this respect. Those who propose

†This is not to say that theoretical debate cannot proceed by comparing distinct architectures and their ability to capture the data – because this is already an area of productive debate in cognitive science (e.g. Roberts and Pashler 2000, Pitt *et al.* 2002). We are instead recommending the use of a single formal framework that permits comparison of different theoretical constructs that could exist within that framework. The relative contribution of these different constructs for fitting experimental data, within the same set of agreed mathematical or formal principles, would then adjudicate between the competing theories.

symbolic rules and representations have often urged core qualitative differences between these kinds of states and the probabilistic distributed states that are the hallmark of statistical models, such as connectionist simulations. For this reason, choosing a formalization that has a pre-existing affiliation with a particular theoretical framework, such as a production system or a connectionist model, biases the enterprise toward the theory from which the model originated. A common ground should instead derive from a formalization that can already adequately incorporate and implement both sides of the theoretical debate. A mathematization or formalization of the debate over representation needs a common framework for directly comparing symbolic states and dynamic processes within the same explanatory arena.

In this article, we propose that a branch of dynamical systems theory may serve as this common ground. Symbolic dynamics has both complemented understanding of the continuous-time nature of systems and provided groundbreaking insight into the computational power inherent in dynamical systems (e.g. Crutchfield 1994). A few proposals concerning symbolic dynamics have already been offered from contributors outside of cognitive science. Below we introduce symbolic dynamics for the cognitive scientist, and review some of these proposals. Before describing this framework, we first offer discussion supporting the position that some portion of our perceptual/cognitive processes is already awash in continuity: that the best physical description of the mind/brain must invariably invoke, at some level, continuous (or discretely-approximated continuous) bases for understanding the substrate of cognition. A theory of cognition is superimposed on this continuity in two broad ways historically: discrete-symbolic or continuous-distributed representational states and processes as the theoretical basis for cognitive explanation. We then introduce symbolic dynamics as a framework that can incorporate (and, in certain special cases, show the equivalence of) both kinds of explanation. We finally offer our own consideration of symbolic dynamics, with its potential contribution to and limitations in cognitive science.

2. Continuity

A key point Dietrich and Markman (2003) use to support discrete representations is our cognitive system's ability to form categories for objects in our world: 'If a system categorizes environmental inputs then it has discrete representations' (2003: 101). Moreover, they argue, continuous accounts of categorization would miss the mark, since categorization by definition involves consistent responses to completely distinct elements in our environment—it makes no room for continuity (historically this may be arguable; cf. Rips *et al*. 1973, Rosch 1975). Although Dietrich and Markman offer extensive discussion to forestall possible replies, there remains a problem with this perspective. What the authors dub 'enduring classes of sameness' (2003: 101), in an environment that our system must categorize, involve discrete internal representations whose primary evidence comes from what might be called *time-course irrelevant* responses during a cognitive task. An outcome-based response measure, such as a forced-choice categorization task, is time-course irrelevant because the temporal dynamics of representational activations leading up to the forced choice go undetected by the response measure. Such response measures may artificially exaggerate the degree to which the enduring classes exhibit their sameness. For example, even the cognitive literature's darling of discrete mental events, categorical speech perception (e.g. Harnad 1987, Liberman *et al.* 1957), exhibits some graded sensitivity to continuous phonetic feature information when its temporal dynamics is measured with reaction times and eye movements (Pisoni and Tash 1974, McMurray *et al.* 2003; see section 2.2 below). In this article, we consider two realms, visual cognition and language comprehension, in which an unmistakable continuity is observed even in seemingly discrete, categorical tasks (see Spivey and Dale 2004 for further review).

2.1. *Vision*

Vision research is replete with examples of continuity in real-time perception. The gradual settling of a population code of neurons, over the course of hundreds of milliseconds, is a typical way to think about how the visual system recognizes objects and faces. Compelling visualizations of the continuous manner in which sensory input gradually produces a percept can easily be found in visual neuroscience. We briefly consider three cases: object and face recognition, visual search, and perceptual decisions.

Rolls and Tovee (1995) recorded from neurons in macaque inferotemporal cortex, and found that it takes a few hundred milliseconds for a responsive population of cells to achieve their appropriate firing rates indicating full identification of a fixated object or face. The cumulative information (in bits) provided by an inferotemporal neuron in the service of recognizing a face or object accrues continuously (though nonlinearly) over the course of about 350 milliseconds until asymptote. Perrett *et al.* (1998) demonstrate similar patterns of gradual accumulation of neuronal evidence during face recognition. When an object or face is partly rotated away from the frontal view, recognition or matching will generally take longer as a function of how far it is rotated (e.g. Shepard and Metzler 1971, Cooper and Shepard 1973, Jolicoeur 1985; see also Georgopoulos *et al.* 1989). Perrett *et al.* (1998) describe recordings from cells in the monkey temporal cortex during viewing of frontal, three-quarter profile, profile and quarter profile schematic faces. When the accumulated action potentials are plotted over time, these curves gradually rise to asymptote over the course of several hundred milliseconds, but at different rates as a function of how canonical the face orientation is.

The same kind of gradual accumulation of perceptual evidence can be observed when multiple objects are competing for attention during visual search. Although a serial-processing account has argued that the observer allocates attentional resources wholly and discretely to individual objects, one at a time (e.g. Treisman and Gelade 1980, Treisman 1988; see also Wolfe 1992), a parallel-processing account is being developed in which attention is best characterized as involving partially active representations of objects simultaneously competing for probabilistic mappings onto motor output (e.g. Desimone and Duncan 1995, Reynolds and Desimone 2001). In fact, a wide range of studies have been suggesting that the traditional distinction between putatively 'serial' and 'parallel' search functions is best revised as a continuum of search efficiency rather than two separate mechanisms of visual search (Duncan and Humphreys 1989, Nakayama and Joseph 1998, Wolfe 1998, Olds *et al.* 2000; see also Spivey *et al.* 2001).

Finally, Gold and Shadlen (2000) examined decision processes in macaque visual perception. A common task in visual psychophysics involves presenting a display of quasi-randomly moving dots. As the experimenter increases the proportion of dots that move in a roughly consistent direction, the perception of a coherent direction of flow amidst the dots becomes more apparent (Britten *et al.* 1992). Monkeys were trained to indicate the perceived direction of dot flow, upon offset of the stimulus, by making an eye movement to one peripheral location or an opposite one. After identifying a relevant frontal-eye field (FEF) region, electrical microstimulation evoked an involuntary saccade that was perpendicular to the two voluntary response saccades. On some of the direction-of-flow judgment trials, this region was microstimulated immediately after the moving dot display disappeared, i.e. exactly when the monkey was supposed to produce a voluntary eye movement that would indicate his response regarding the perceived direction of flow of the dots. By incrementally increasing viewing time of the stimulus before this microstimulation, the experimenters were able to observe the gradual increase in 'strength' or 'confidence' of the perceptual decision over time, as indicated by the degree to which that voluntary saccade 'leaked into' the execution of FEF-microstimulated evoked saccade. Thus, the population of FEF cells that produced the evoked saccade was already somewhere in the process of settling toward a pattern of activation that would produce the voluntary response saccade. If the microstimulation took place early on in this decision process, rather little effect of the voluntary response would be apparent in the direction of the evoked saccade, but if the miscrostimulation took place later on in the decision process, a significant amount of the voluntary response would be apparent in the direction of the evoked saccade. This finding suggests that decision processes themselves may be coextensive with the gradual settling of partially active and competing neural representations in oculomotor areas of cortex (Schall 2000, Gold and Shadlen 2001; see also Georgopoulos 1995).

2.2. *Language*

Much like visual perception, language comprehension also exhibits a form of temporal dynamics that reveals underlying continuous-distributed formats of representation. There is considerable evidence that the multiple levels of linguistic complexity—comprehension of speech sounds, words and sentences—are driven by graded, partially active representations. At the level of speech sounds, the phenomenon of categorical perception was long adduced as evidence for discrete representational states (Liberman 1982). Lately it has been subject to extensive empirical investigation, and made consistent with more temporally dynamic approaches to categorization (Damper and Harnad 2000; see also Anderson *et al.* 1977). For example, McMurray and Spivey (1999) tracked participants' eye movements while they performed the standard categorical identification task. This task involves categorizing sounds that lie on a voice-onset-time continuum between 'bah' and 'pah', by clicking a relevant icon on one or the other side of the computer screen. Thus, in addition to recording the participants' explicit choice, there was also a semi-continuous record of how the eyes tended toward one or the other response icon during categorization. With 'pah' or 'bah' sounds near the categorical boundary, eye movements exhibited conspicuous vacillation between categories before the overt mouse-click response was made. Despite the apparent categorical nature of the eventual choice, eye movements revealed a more continuous decision process that is sensitive to some of the graded acoustic–phonetic variation in the stimulus. These temporary phonemic ambiguities exhibit their effects not just in phoneme categorization tasks but also in spoken word recognition tasks (McMurray *et al.* 2002, 2003).

At the level of word recognition, Spivey-Knowlton *et al.* (1994) demonstrated cohort effects in eye-movement patterns by having subjects follow instructions to manipulate real objects on a table. Participants sat in front of a table containing a central fixation cross and various objects around it (e.g. a fork, a mug, a candle). In some trials, objects whose names had similar initial phonemes were present on the table, available for manipulation (e.g. a bag of candy and a candle). Even before the spoken word was completed, eye-movements to both objects were often observed, such as briefly fixating the candle when instructed to 'Pick up the candy'. This phonologically similar object conspicuously attracting eye movements is indicative of the competing lexical representation being partially active during, and perhaps shortly after, delivery of the spoken word. Headband-mounted eye-tracking studies like this have demonstrated this real-time lexical competition using computer-displayed objects (Allopenna *et al.* 1998), using artificial lexicons (Magnuson *et al.* 2003), with young children (Fernald *et al.* 2001), and even across two languages in bilingual participants (Spivey and Marian 1999, Marian and Spivey 2003).

Finally, in sentence processing, eye movements can again reveal the continuous intake and use of information during comprehension of a spoken utterance. For example, when presented with a real 3D display containing an apple on a towel, another towel and an empty box, and then instructed to 'Put the apple on the towel in the box', participants often look briefly at the irrelevant lone towel near the end of the spoken instruction before returning their gaze to the apple, grasping it and then placing it inside the box (Tanenhaus *et al.* 1995, Spivey *et al.* 2002). In this case, the syntax is ambiguous as to whether the prepositional phrase 'on the towel' is attached to the verb 'put' (as a movement destination) or to the noun 'apple' (as a modifier). Given the actions afforded by the display, the latter syntactic structure is the correct one. However, the brief fixation of the irrelevant lone towel indicates a temporary partially activated incorrect parse of the sentence. To demonstrate the influence of visual context on this syntactic ambiguity resolution process, the display was slightly altered to include a second apple (resting on a napkin). In this case, the visual co-presence (in Herb Clark's (1992) terms) of the two potential referents for the phrase 'the apple' should encourage the listener to interpret the ambiguous prepositional phrase 'on the towel' as a modifier (in order to determine which apple is being referred to) rather than as a movement destination (cf. Crain and Steedman 1985, Altmann and Steedman 1988, Spivey and Tanenhaus 1998). Indeed, with this two-referent display, participants rarely fixated the irrelevant lone towel, indicating that visual context had exerted an immediate influence on the incremental syntactic parsing of the spoken sentence (Tanenhaus *et al.* 1995, Spivey *et al.* 2002; see also Knoeferle *et al.* 2003).

The current state of affairs in the field of sentence processing is at a consensus with regard to the continuity of information flow, and has been gradually approaching consensus with regard to the rapid integration of syntax, semantics and pragmatic context (Trueswell and Tanenhaus 2004). Just as the processing of speech sounds, at the scale of tens of milliseconds, appears to be characterized by multiple partially active phonemic representations competing over time (McMurray *et al.* 2002, 2003), and the comprehension of spoken words, at the scale of hundreds of milliseconds, appears to be characterized by multiple partially active lexical representations competing over time (McClelland and Elman 1986, Marslen-Wilson 1987, Allopenna *et al.* 1998), so does the resolution of syntactic ambiguity, at the scale of seconds, appear to be characterized by multiple partially active syntactic representations competing over time (MacDonald *et al.* 1994, Stevenson 1994, Spivey and Tanenhaus 1998, Tabor and Tanenhaus 1999).

2.3. *Summary*

From perception, such as visual processing, to cognition, such as the various levels of linguistic processing, there seems to be extensive evidence for continuous-distributed representation (see Spivey and Dale 2004, for further discussion and examples). There nevertheless remains considerable debate about the nature of representation in other areas of cognition. In particular, in 'high-level' cognitive processes such as reasoning and problem solving, there seems to be markedly slower success with continuous-distributed frameworks. This situation is exacerbated further by the comparatively rapid rise, and longer history, of discrete-symbolic accounts of reasoning and problem solving (e.g. Weizenbaum 1966, Winograd 1970, Newell and Simon 1976).

If it can be granted that perception is largely driven by continuous change in processing states, then for the discrete-symbolic perspective to be right about cognition, there must be a 'discretization' that happens somewhere in between perception and motor output. The debate can then be placed in the following terms: *how early in the system do our theories need to postulate this discretization, thus invoking a language of discrete symbols generated through causal influences of continuous processes?* A purely continuous-distributed account of cognition might place this discretization at the extreme end, only in between the motor action itself and its effects on the problem-solving environment. For example, although you may be trying to decide between moving your rook four squares up or three squares up in a game of chess, and this vacillation may even be visible in the continuous motor movement, in the end, only one of those alternatives actually happens. In contrast, the discrete-symbolic account of cognition urges an earlier discretization, recommending theories to work with symbolic states and algorithmic state-transition rules not long after perceptual processing. In such a case, the decision to move one's rook three squares or four squares would be discretely made in an internal cognitive stage, and any vacillation observed in the motor movement would be best interpreted as a vestigial or epiphenomenal echo of the earlier temporary cognitive uncertainty.

If this formulation of the question is agreeable to both sides of the debate, then there exists a 'common format of explication' that future research in high-level cognition might fruitfully use in order to consensually adjudicate between theories that propose an *internal* discretization of the brain's continuous dynamics and theories that propose only an *external* discretization of them. The mathematical arena of symbolic dynamics (e.g. Crutchfield 1994, Devaney 2003, Goertzel 1998, Shalizi and Albers submitted; see also Cleeremans *et al.* 1989, Tabor 2002, for related discussions) has exactly the ingredients for building systems that implement continuous temporal dynamics in a high-dimensional state space (of perception and of action) and can convert that continuous trajectory into an emitted string of formal logical symbols for describing external action-effects in a problem-solving environment, and also for describing internal cognitive states.† We next offer a very simple introduction to symbolic dynamics, and then discuss a number of issues relevant to its application in cognitive science.

3. Symbolic dynamics

A continuous-distributed perspective on representation in perceptual and cognitive processes is often couched in model systems that change in time (be it continuous-time or discrete-time): dynamical systems. A dynamical systems framework provides a rich set of conceptual tools for cognitive science. The geometric entities in the study of dynamical systems can serve as an intuitive, and potentially mathematically rigorous, vocabulary for visualizing state changes within and between perception and cognition. As already mentioned, this strategy is widely used in many areas of cognitive science, and is often considered its own framework for the study of cognition (Thelen and Smith 1994, Kelso 1995, Port and Van Gelder, 1995, Ward 2002). In order to lay out this descriptive vocabulary, we briefly consider a simple iterative dynamical system, surely familiar to many readers, that illustrates a number of these geometric metaphors. Consider a function $F(x)$ that maps real numbers onto real numbers by iteration: $F^2(x)$ is given by $F(F(x))$, and $F^3(x)$ by $F(F(F(x)))$, etc. The logistic map is given by the equation

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

The time dimension is here represented by progressive iteration of the real value x into the function F, scaled by μ. The iterative process in this simple equation illustrates stability, meta-stability, and transition into chaotic behaviour. For example, when μ is between 0 and about 3.5, iteration of $F(x)$ from any starting point of x will settle into stable attractor states—namely, the value of $F^n(x)$, as n becomes very large, stabilizes on one or more particular precisely-repeated values. These values are termed attractors in the logistic map's dynamics. As μ approaches about 3.6 or so, the logistic map exhibits chaotic behaviour, where there is no stable attractor state, and its series of values can superficially appear random. One way of representing the transitions in state space of this system is through a phase plot, shown in figure 1. By tracking the value of x at each iteration, we can visualize the trajectory of the system from some random initial x into its attractor states (figure 2).

The logistic map is used extensively in textbooks on dynamical systems. Its curiosity lies in the rich complexity that emerges from iterating such a simple equation. In fact, the same issues considered for discerning the nature of the logistic map are a concern for innumerable systems of practical and theoretical interest. A wide variety of pure and applied mathematical techniques can be used to study

†In fact, while defending discrete representations in cognition, Dietrich and Markman (2003) essentially describe the basic concept of symbolic dynamics, without referring to it by name, in their fourth argument (their discussion of figure 3).

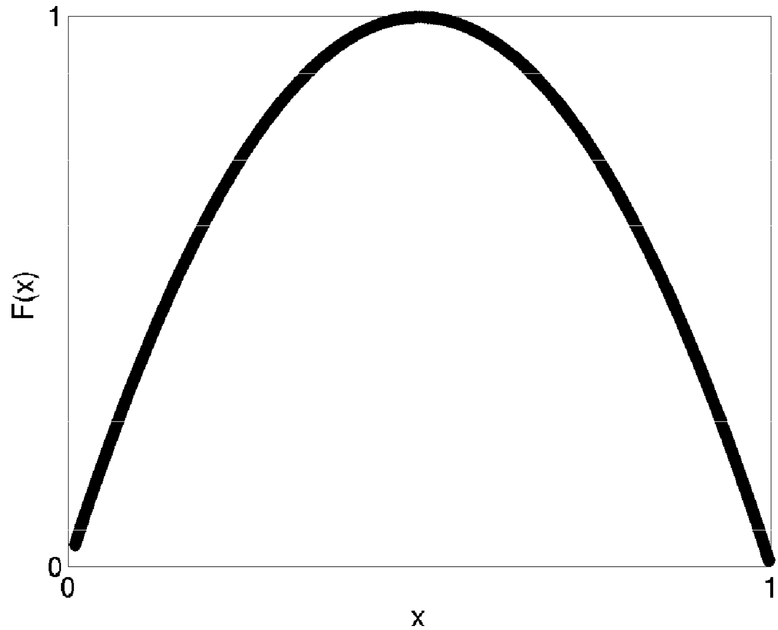

Figure 1. A phase-plot for the logistic map. Provided μ is above 0 and below 4, the system lies within the interval of [0, 1]. In other words, given its current value of x at time t, its subsequent iteration, time $t + 1$, will be on this line (here with $\mu = 3.9$).

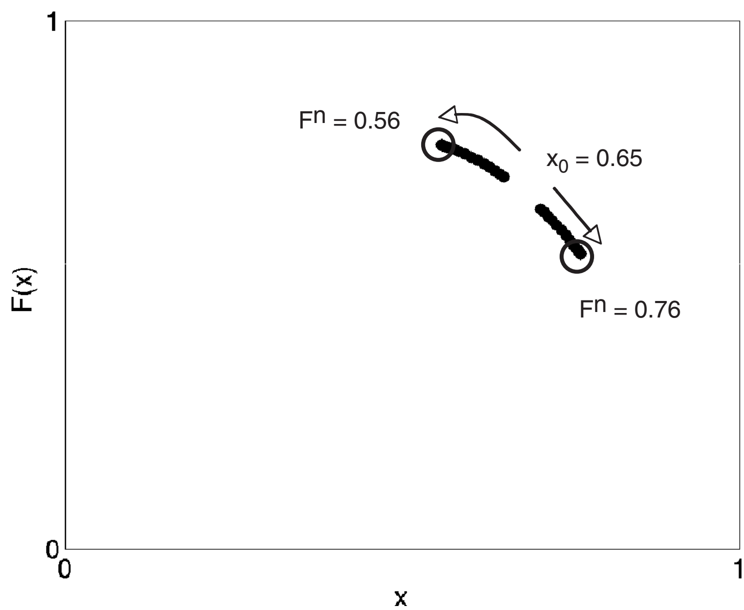

Figure 2. In this 'phase flow' diagram, the logistic map moves into two stable attractors with $\mu = 3.1$. The system starts at $x_0 = 0.65$. As the x is iterated through F, it settles into two attractor states, approximately 0.56 and 0.76, between which it will alternate indefinitely.

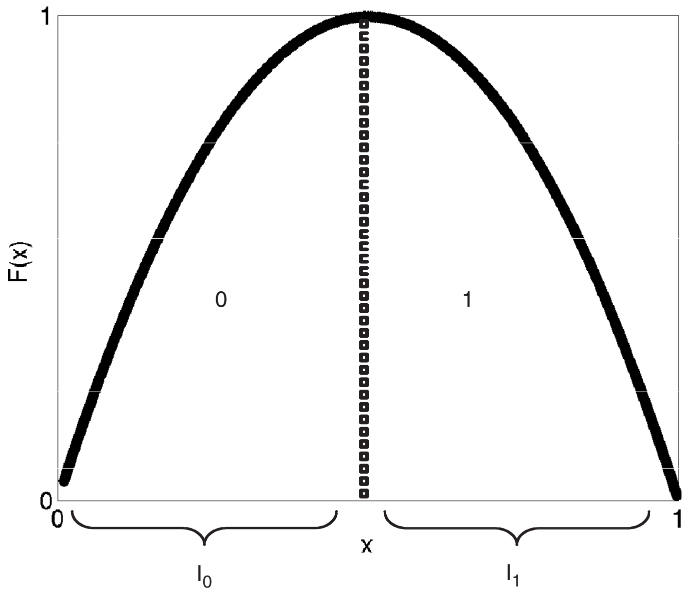

Figure 3. The phase space of the logistic map can be carved into two intervals. Each time the system enters the interval, it outputs the symbol corresponding to that region ('0' or '1'). Here $\mu = 3.9$ again.

these systems. Nevertheless, it is not uncommon that these methods can be outstripped by a system's complexity. A specific technique available to overcome such limitations, around for decades and gaining much attention of late, is termed symbolic dynamics, and offers a means of simplifying analysis (e.g. early on, Morse and Hedlund 1938; see Devaney 2003 and Williams 2002 for review). A system's dynamics can be rendered symbolic by carving partitions or regions into its phase space, and assigning a unique numeric state or label to that partition. As the dynamical system's state changes in time, this trajectory is transformed into a sequence of emitted symbols corresponding to partitions in the space. Take, for example, the logistic map. We can represent its phase space as in figure 3, and divide the plot into two intervals, $I_0 = [0, \frac{1}{2}]$ and $I_1 = (\frac{1}{2}, 1]$. When iterations of the system enter the first interval, the symbol '0' is output, and likewise '1' for the second. The dynamics of the logistic map may therefore be represented by a sequence of 1s and 0s, indicating its approximate position in these partitions at each iteration.

At $\mu = 2.9$, the system quickly descends onto a particular stable attractor state at approximately $x = 0.65$. The symbol sequence generated by this system is extremely simple—'1111...'. The system in fact never leaves this interval, never passing the threshold into the other, and therefore emits the symbol '1' for any subsequent iteration once $F^n(x)$ reaches its attractor. However, when $\mu = 3.55$, for example, the map fluctuates for a bit and then reaches eight distinct and perfectly repeated attractors. Once it reaches this meta-stable state, while tracing the transitions across intervals I_1 and I_0, this trajectory generates the sequence '011101110...'. This may be simplified using the notation $(01110111)^n$ or even $(0111)^n$. Contained in this simple sequence rule is the original dynamics: transitions between eight separate attractor states.

The above example is deliberately simplified for the sake of introducing rudimentary dynamical systems and symbolic dynamics. The logistic map affords this simplification. The strategy of employing symbolic dynamics, however, is somewhat more complex in most contexts. Symbolic dynamics rapidly served to help explore chaotic dynamical systems in more theoretical contexts (see Williams 2002, for a review). In a further simplified example, we can straightforwardly introduce what this application entails. Consider an alphabet of N symbols that we might use in our partition of a system's phase space, $A = \{0, \ldots, N-1\}$, and the space of all possible sequences constructed from this alphabet:

$$\sum_2 = \{S | S = s_0 s_1 s_2 \ldots, \text{ and } s_i \in A\}$$

Here, s_0 is the first symbol emitted by the dynamical system, and the sequence continues *ad infinitum*. The set $\sum_2$ is the space of all such possible sequences. A particular system's dynamics can be captured by shifting its infinite sequence, $S \in \sum_2$, to the left, so that $s_0 s_1 s_2 \ldots \to s_1 s_2 s_3 \ldots$, and the new sequence begins at the next emitted symbol, s_1. This shift operation captures the progression in time of emitted symbols, and is often represented by σ, so that $S' = \sigma S$, where $s'_i = s_{i+1}$. This shift operation can act as a mapping on a continuous space, σ: $\sum_2 \to \sum_2$, by specifying a distance measure or metric between sequences, $d(S, S')$. In other words, the trajectory of a system can be represented in terms of an ordered set of infinite symbol sequences, formed by progressive shifting.

From here, a means of exploring dynamical systems involves demonstrating that the space of symbols $\sum_2$ and its shift map σ have a certain geometrical equivalence to a dynamical system's own continuous mapping and the set of states which it visits. The famous Smale horseshoe can be studied through partitions of its phase space—and through its symbolic dynamics, it can be shown to have particular dynamical features (e.g. chaos; Medio and Lines 2001). The logistic map has also been explored through its symbolic dynamics. Consider the case where the control parameter μ in F is larger than 4. It is easy to see that most initial states will have F^n approaching $-\infty$ as n gets larger. Specifically, since $x = 0.5$ grants the product $x(1 - x)$ its largest value (0.25), any value for μ that is greater than 4 will take F outside the interval [0, 1], and thus at the next iteration, on a path towards infinity. However, inspection of the phase plot for the logistic map in figure 4 reveals the simple observation that not all values of x take F out of [0, 1]. The set C of all values that avoid this escape, along with the function F, can be shown to have this kind of equivalence with $\sum_2$ and σ, and allow certain conclusions about the properties of this set: once again, the map F on C is chaotic (Devaney 2003).

These textbook examples of the theoretical and mathematical benefits of symbolic dynamics merely scratch the surface of its recent role in the study of dynamical systems. Recent excitement has instead been concerned with the extent to which symbolic dynamics is informative about more complex systems through statistical analysis of its output. Symbolic dynamics is thus intriguing because it offers structures of sequences that can be subjected to a wide variety of 'tricks for predicting discrete stochastic processes' (Shalizi 2004a). Such statistical analysis has offered insight into complex dynamic processes in a wide variety of fields, including astronomy, biology, chemistry and computational linguistics (see Daw *et al.* 2003 for a review). The past two decades have also seen symbolic dynamics

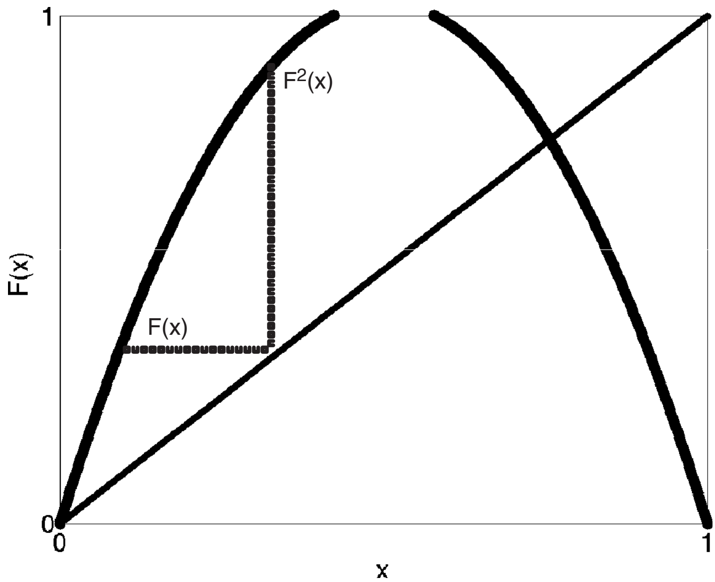

Figure 4. The logistic map phase plot with $\mu = 4.1$. A portion of the phase space is outside the interval [0, 1], and points leaving will tend to $-\infty$ through F iteration. However, a set of points does not leave this interval, illustrated with one iteration of some value x (dotted lines). The initial value x becomes $F(x)$, and remains in the [0, 1] interval. Symbolic dynamics allows investigation into the nature of these iterations that do not escape (see text for more detail).

make explicit connections between the study of digital computation and that of continuous dynamical systems (Crutchfield 1994).

Symbolic dynamics also has the interesting property of sometimes exhibiting equivalence with the continuous system from which it originates. As long as a partition is adequately selected, analysis of the symbol sequences can actually be used to reconstruct the continuous dynamics of the original system. A specific kind of partition, termed *generating partition*, can in fact yield 'approximately complete and precise descriptions of the system' (beim Graben 2004: 47). Perfect definition of a generating partition requires knowledge of the original dynamical system, but there exist techniques for approximating these demarcations (e.g. Davidchack *et al.* 1999, Kennel and Buhl 2003). Such generating partitions allow the symbolic dynamics to be topologically equivalent to the original continuous dynamics (Kitchens 1998, beim Graben 2004, Shalizi and Albers in press). However, finding generating partitions is very difficult in systems consisting of more than two dimensions (Kennel and Buhl, 2003), and they only work for deterministic dynamical systems (Crutchfield and Packard 1982). Therefore, much of the practical applicability of symbolic dynamics may lie in iteratively refined approximations of generating partitions, rather than true generating partitions. For example, non-generating partitions in symbolic dynamics have been used for describing the phase-space of bimanual rhythmic co-ordination (Engbert *et al.* 1998) and of heart rate variability (Kurths *et al.* 1995). However, with even slightly misplaced partitions, the threshold-crossing method for emitting symbol strings from continuous trajectories can very easily introduce severe compounded misrepresentations of the original continuous dynamics, i.e., grammatical errors in the symbol sequences (Bollt *et al.* 2000, 2001). Symbolic sequences from more complex systems have therefore often been subject to more experimental or empirical styles of analysis (Daw *et al.* 2002).

To summarize, there have been two broad areas in which symbolic dynamics have made a clear impact. In the first, and ultimately its origin, it is explored extensively in pure theoretical contexts in mathematics to study tractable systems. In another, it has played a role in simplified descriptions and statistical analyses of more complex iterated mathematical processes, and even in application to dynamics of complex physical systems. It may, we argue, offer something to cognitive science theory as well. In the debate on representational format, symbolic dynamics could make headway toward formalizing theoretical debate. Current discussion on symbolic dynamics raises a number of important questions in this respect. We next consider these, and then introduce some reflections on future directions for symbolic dynamics in cognitive science.

4. Symbolic dynamics and cognitive science

Dietrich and Markman (2003) actually describe something very close to symbolic dynamics in a short segment of their paper supporting discrete representations. They offer a number of properties that cognition must have, which only discrete representations endow. One of these properties is compositionality: representations best explaining many cognitive processes must have component parts that are combined (see also Fodor and Pylyshyn 1988, Marcus 2001). They argue that any representational subsystem, if continuous, can only have parts if there is some other system that discretely interprets its regions, and takes in discrete representations as input. This is in fact a description of symbolic dynamics, though there are details to be worked out. For one, the resultant symbols, if not time-course irrelevant, might encode the original dynamics of the system, as mentioned in the way of generating partitions above. This would mean an equivalence relation holds between the two systems, at least in the sense that the symbols carry some of the continuous information in the original dynamic subsystem. Secondly, it has been demonstrated recently that dynamical systems actually do have considerably surprising computational powers. In fact, a number of these properties often considered hallmarks of discrete-symbolic algorithmic processing can be approached with symbolic dynamics.

For example, one such property, *discrimination*, is easy to achieve through translation into symbol sequences. Multi-stable one-dimensional dynamical systems can emit symbols pertaining to any stable point (and a given interval around it) in its phase space—as a matter of discrete, symbolic output from that system. This scenario may indeed be superior to verbalized discrete theories since symbolic output from an iterated map retains some information about time. For example, a metastable system that drifts slowly will produce symbol sequences with long strings of identical symbols, indicating its inhabiting of some categorical state. The output is therefore discretely representational, but also reveals patterns of change in time. Perceptual state-space is of course not a matter of collapsing over a single dimension—the situation becomes very complex when we consider the number of categories (symbols) that need to be represented, and the fact that our visual system, for example, is translation invariant, so whatever partition can define the symbolic output from visual to cognitive processes must exist in a very large number of dimensions. A second problem concerns delineating the stages at which such collapsing from continuous mappings to discrete symbol strings occurs. As mentioned, to account for continuous perceptual states moving into something akin to sets of sameness, there must be one or more stages of 'analogue-to-digital' (A-to-D) conversion.

Probably the most studied and established property that dynamical systems exhibit through their symbolic dynamics is their digital computation – that a description of dynamical systems can take the form of explicating its information-processing capacities via symbol sequences. This feature approaches the well-known issue of *systematicity*, a property that many have argued cognitive systems must have (especially human ones; e.g. Fodor and Pylyshyn 1988, Hadley 1994, Marcus 2001). This discussion fits into the scope of symbolic dynamics in three ways. First, it is possible to characterize the dynamics of a system through computational descriptive schemes. For example, the classic paper by Crutchfield and Young (1989) introduced an approach to nonlinear dynamical systems that quantifies their computational qualities. Subsequent research has pursued the extraction of such intrinsic computation from nonlinear dynamical systems, among other systems (see Crutchfield 1994 and Andrews 2002 for reviews). Sought after qualities of systematic computation urged by Hadley (1994) and Fodor and Pylyshyn (1988) may very well be encoded in the edge-of-chaos dynamics of even simple systems (Crutchfield and Young 1990).

The second way systematicity can fit into symbolic dynamics is through exploring the ability of dynamic systems to acquire formal languages. The system that learns the language may again be characterized in terms of symbolic dynamics. For example, the well-known simulations in Pollack (1991) demonstrated that a neural network can learn context-free languages and classify novel sentences from such grammars via a decision process akin to symbolic dynamics (see also Cleeremans *et al*. 1989 for related examples). More recently, Rodriguez *et al*. (1999) investigated a very simple recurrent network in its ability to learn deterministic context-free grammars. Networks that learned successfully performed a form of 'counting' in their phase space. This allows successful learning of the context-free grammar without explicitly implementing a pushdown automaton. Also, Tabor (2001) recently used a neural network model trained to predict sequences of symbols from four languages of differing levels of complexity (see also Tabor 2000). Networks trained on context-free languages, as opposed to the regular languages, exhibited edge-of-chaos effects (or intermittency), revealing the kind of intrinsic computational qualities outlined in other nonlinear dynamical systems by Crutchfield and colleagues (Crutchfield and Young 1989).

Finally, one may simply take the symbols for granted at some level. In this case, though the solution appears simplistic, the difficulty is in delimiting and discovering the nature of the interface between continuous perceptual states and the resultant discrete cognitive informational states that undergo algorithmic manipulation. Resorting to this idealization, and thereby taking for granted straightforward algorithmic parlance about representations, requires explication of A-to-D conversions between high-dimensional continuous perceptual states and their entryway into cognitive processing. Although the above review suggests the continuous dynamics of simple systems can already exhibit surprising computational qualities, it may be in the domain of this symbolic demarcation that debate between formats is best mitigated. Following our discussion here of theoretical issues in the use of symbolic dynamics, we offer some sketches of the ways in which this mitigation might take place in computational models.

There are numerous issues with both symbolic dynamic theory and application that are relevant to theoretical frameworks in cognitive science. We introduce three of these issues here, and elaborate further in the subsequent sections of this article. The first issue concerns the consequences of generating partitions. Some have argued that the equivalence between symbolic dynamics (from a generating partition) and the continuous mapping from which it originated, renders moot the debate between continuous and discrete states in the mind (Crutchfield 1998, Shalizi 2004b). Though compelling at first pass, the argument is based on simple, low-dimensional systems—ones whose consequences cannot be handily generalized to noisy, high-dimensional (and likely highly non-stationary) dynamics in neural systems at the level of cognitive processing. As a second issue, we consider the implications of recent discussion concerning the epistemic problems of finding appropriate partitions for meaningful symbolic dynamics (beim Graben 2004). This has direct relevance to conceptualization of 'error' in competence and performance, and the nature of language comprehension and production, among other cognitive processes. The final issue concerns how continuous and symbolic dynamical systems function in tandem during perception and cognition. Presumably, if discrete-symbolic descriptions are most suitable for 'higher' cognition, then there must be some stage at which continuous dynamics of perceptual or 'lower' cognition gets transmogrified into interpretable symbolic states.

4.1. *Continuous-symbolic equivalence*

As discussed above, there are many reasons to study continuous dynamic maps via symbol sequences. An interesting fact for many such maps is that there exist generating partitions that emit symbol sequences exactly reflecting the original dynamics of the system. This has led some to argue that the debate concerning discrete-symbolic and continuous-distributed representations is ill posed (e.g. Shalizi 2004b). Since a dynamical system can be seen as identical with some symbolic dynamics, it might be inappropriate to suppose that two formats of representation are at odds when they are mathematically equivalent. As already mentioned, Crutchfield and collaborators (Crutchfield 1994) conceive of dynamics as inherently computational, and offer numerous techniques for generating computational machinery from symbolic sampling of continuous states (e.g. ε-machines; Crutchfield 1994). Elsewhere, Crutchfield has argued that supposing dynamics can replace discrete computation (e.g. Van Gelder 1998) neglects the intrinsic computational nature of dynamical systems themselves (Crutchfield 1998).

Early in Crutchfield's seminal paper (1994), he distinguishes between two concepts of computation. The first, 'useful' computation, refers to specific instantiations of input-output mappings in some computational architecture. The second, 'intrinsic' computation, concerns the basic capacities and limitations of a computational system, dynamical or otherwise. This involves exploration or specification of information-processing capacities of a system, without reference to any specific 'useful' input–output accomplishment. This perspective has led to extremely fruitful research on discovering the underlying computational aspects of nonlinear dynamical systems. For example, early work by Crutchfield and Young (1989) sought to specify and measure the complexity of minimal stochastic automata whose state transitions (emitting symbols) embody the logistic map's dynamics at differing values of μ (see also Crutchfield and Young 1990).

There are, however, reasons for remaining cautious about the direct implications in our understanding of complex cognitive states. An in-principle statement concerning the equivalence of continuous and symbolic dynamics in systems is not sufficient on its own to alleviate debate. There are at least two related reasons for this. First, relying on such equivalence neglects the very crucial and substantial details that debate on representational format carries. A rather straightforward one is the kind of characterization that symbolic and distributed formats receive. For example, Andy Clark (2001) characterizes much symbolic cognitive theory as resting on representations whose contents are *semantically transparent*. A classical computational theory of language deals in representations of words, their meanings, and the structures that they compose. These representational formats are highly 'scrutable', their significance in a system's computation immediately accessible. However, systems relying on probabilistic and distributed representations or states, such as connectionist systems, often rely on formats that are *semantically opaque*. For example, establishing the function of a hidden-unit manifold often involves detailed statistical analysis of the hidden-unit activations under varying circumstances. The resultant function may be very nonlinear and complex, and not easily describable through commonsensical or folk-psychological labels.

For this reason, simply saying that the two kinds of descriptive machinery, continuous and symbolic, both serve the same functions actually skirts some substantive issues. The debate concerns explanation in terms of specific kinds of computational mechanics—concretely identifiable words in our 'language of thought' (Fodor 1983), or some other more or less semantically transparent discrete states. These are pitted against models accounting for behaviour in terms of distributed representations whose interpretations are less obvious, or perhaps 'subsymbolic' (Smolensky 1988). In fact, cognitive science has already had a number of battles concerning whether these two systems are equivalent, or the second is just a special case of the first, and so on (e.g. Fodor and Pylyshyn 1988, Lachter and Bever 1988).

The second reason to be cautious about the lesson from symbolic-dynamic equivalence is that 'useful' computation has been considerably less explored than 'intrinsic' dynamics in the study of computational mechanics. Although the current accomplishments can only be described as some of the most exciting and relevant to cognitive science, they have yet to delve into systems whose complexity can match a level of description needed for understanding *cognitive* processes. Van Gelder (1998) replies in this manner, remarking that when 'it comes time to model the complexities of real cognition—to publish in Psychological Review rather than Physica D—they may find that the dynamics drops out of the picture, and the relevant story is cast entirely at the level of the emergent computation. Alternatively, they may find (as have many dynamicists) that the computational aspects play second fiddle to the dynamics' (1998: 13). This seems to misunderstand what is accomplished in symbolic analyses of dynamical systems: The descriptions are two sides of the same computational coin. A more direct concern at present is whether meaningful partitions can be established. As we approach a level of complexity that matches what is accomplished in a neural substrate, or proposed cognitive processes of multiple dimensions, the likelihood of finding generating partitions drops radically (see next section).

So, whether or not we embrace the equivalence of symbolic and continuous dynamics through generating partitions, we are still left with some confusion. Are the discrete symbolic states of our cognitive system available for scientific scrutiny, and the dynamics more complex (yet equivalent)? Or are discrete symbolic states of our cognitive system inadequate explanatory constructs, and we should reach for continuous dynamic descriptions of our mind/brain? Churchland (1992) offers discussion relevant to these more substantive issues in the domain of neural networks, and considers partitions of their hidden-unit state space that can reflect conceptual structure in human cognition. Adopting a set of partitions, Churchland argues, 'may suffice for the accurate short-term prediction of its behaviour, but that knowledge is inadequate to predict or explain the evolution of those partitions over the course of time' (1992: 178). We argue that, in the domain of higher levels of cognitive processing, this position has considerable merit, but is very much without consensus in the field. It is thus through these substantive issues that the two formats of representation stake their respective claims.

4.2. *Epistemic issues*

Similar to Andy Clark's (2001) distinction between transparent symbols and opaque distributed representations, Atmanspacher (2000) makes a distinction between epistemic and ontic types of description of chaotic systems. An ontic description is exhaustive concerning the dynamical system—it comprehensively encapsulates the composition of the dynamics. An epistemic description is framed in terms of knowledge or ignorance of an observer evaluating these ontic 'states'. Epistemic descriptions are achieved by evaluation of an observed or measured dynamical system, through statistical quantification or characterization of it. This terminology of *ontic* and *epistemic* can be used to frame the previous section's discussion of equivalence through a generating partition. The pure equivalence between a symbolic dynamics and its origin map can involve only ontic descriptions for any system of sufficient complexity. That is, urging the equivalence of some symbolic dynamics and the original system implies a kind of ontic state that is inaccessible to us as observers. Instead, we are confined to epistemic descriptions for complex cognitive systems. In order for a chosen symbolic dynamics to adequately explain or represent the cognitive process under study, it must be chosen appropriately. This is not a trivial matter. We consider two problems of these epistemic issues relating to finding a good partition for a continuous map.

First, beim Graben (2004, beim Graben and Altmanspacher 2004) argues that incompatible, but equally accurate, symbolic epistemic descriptions are possible with multiple non-generating partitions. This means that two sets of different symbolic dynamics may be equally adequate as formal descriptions of the original dynamics, yet mutually incompatible with one another (cf. Quinean indeterminacy: Quine 1960). beim Graben (2004) provides an example of a Hopfield network as a multi-stable dynamical system living in a space of many dimensions (Balkenius and Gärdenfors 1991). There is no generating partition for this space, and indeed multiple descriptions via symbolic dynamics can be mutually incompatible, while remaining equally good (or bad) partitions of the underlying dynamical system. Consequently, non-generating partitions can provide 'conceptual' descriptions of the continuous system, and there may indeed be many such descriptions. While all of them may serve as formal descriptions at a symbolic-discrete level, they can be mutually incompatible with each other. Quine (1960) early on made very similar points concerning the study of linguistic meaning by indicating that many equally good (or bad) rule sets can exist for translation from one language to another; yet these equivalent translation strategies may be mutually incompatible when compared directly (see also Moore 1956 for a related classic theorem; Gauker 2003). It is crucial to point out that this is not merely a fact of further exploring the ontic description of the given system in order to select the better of these incompatible symbolic accounts. In fact, Quine, and beim Graben and colleagues in symbolic dynamics, reveal that even given the full set of ontic states themselves, there are still mutually incompatible, yet thoroughly equally good, symbolic descriptions. In short, they cannot be reliably adjudicated among.

A second issue concerns the kinds of errors that result from inappropriate partitions. Bollt *et al.* (2001) analyse the tent map, whose generating partition is known, and measure the topological entropy resulting from shifts of that partition to varying degrees. The effects of shifted partitions are quite drastic, with topological entropy being affected immediately, and in an irregular (non-monotonic) fashion. The upshot, according to the authors, is that arbitrary partitions (e.g. Kurths *et al.* 1995) can result in 'severe under-representation of a dynamical system' (2001: 281). It is important to point out that these results were based on a well-known deterministic and simple dynamical system. The effects of noise (Crutchfield and Packard 1982) and increasing complexity (Kennel and Buhl 2003) in degrading the fidelity of the emitted symbolics are also well documented.

The foregoing remarks on epistemic limitations on symbolic dynamics have two implications for current discussion. First, they serve to underscore the points made in the previous section on equivalence. The possibility of the existence of a generating partition is not sufficient to dissolve debate on symbolic versus continuous representation. Instead, the epistemological limitations on more complex, noisy dynamical systems suggests that there is considerable room for debate concerning the adequacy of either continuous or symbolic accounts for some cognitive process. Indeed, in most cases (if not all), we do not have sufficient knowledge of the ontic conditions of some cognitive process. The upshot for cognitive science is that continuous or symbolic accounts are (1) highly unlikely to be resolved by mere recognition of equivalence, and (2) are likely to offer differing amounts of coverage of the human data regarding these ontic states—about which the field has much to discover.

4.3. *A-to-D conversions*

Despite these limitations, we argue that the promise of symbolic dynamics lies in articulating the transition from dynamical, continuous descriptions of perception, into the theoretical language of discrete, algorithmic processes for high-level cognition. Whatever the *ontic* states underlying cognition, our *epistemic* descriptions and theories ought to be couched in structures or processes that bear causal relationships to others, and ultimately, to our observable behaviour. If it is the case that some are discrete and symbolic, there must occur a transition into them from a continuous state-space of perceptual or early-cognitive processing. These A-to-D conversions consist of collapsing the continuous-distributed representations onto discrete-symbolic ones that cause inherent information loss about the perceptual states feeding into them. However, this loss may be merely 'lossy', as in image compression algorithms, where the resulting compact representations still carry information appropriate for cognitive processing.

The question is then not merely *when* this transmogrification occurs, but also *what kind* of information from continuous states these discrete states need in order to account fully for observable behaviour. For example, in information-processing accounts of cognition, Miller's (1982) exploration of the concept of information 'grain' provided an early challenge to discover what kind of discrete representations there are: what level of 'granularity' do discrete representations need to have to account for cognition. For example, along with Miller's work on the 'response preparation effect', Abrams and Balota (1991) and Spivey *et al.* (2005) demonstrate that dynamic response measures (e.g. force and velocity measures from a response bar or continuously recorded computer-mouse movements) exhibit graded properties depending on the continuous strength or reliability of the information that produced the motor output. In addition, a wide variety of eye-movement research (reviewed above) suggests that metabolically cheap movements such as saccades reveal a decision process that does not appear perfectly discrete. These studies demonstrate that 'echoes' of continuous information states can be observed in the *dynamic properties of resultant responses*. The discrete states that may have mediated the transitions from sensors to effectors must carry at least some relevant information from early graded states. In other words, while reaction time may reveal information about the decision process during discrete, algorithmic processing, the concomitantly graded manual output from the system indicates that even when these discrete decision processes collapse onto the effectors, there remains some fine granularity.

We can frame the situation simply by defining an idealized 'problem space' for some cognitive process. The space may be maximally simple, from an idealized or simulated continuous perceptual space into one or two symbolic processes. Here, A-to-D conversion performs a *useful* computation in the Crutchfield sense, described above. Whatever the intrinsic computational properties of the initial continuous perceptual space, the system of continuous representation feeding into discrete symbolic processing has an informational function in a problem space of, say, evolutionary relevance (for some such thing as reproduction, or running away from something that might eat you). A simple example is perceptual categorization. An idealized continuous space of perceptual information can be manipulated so as to output symbols feeding into some discrete process. This idealized scheme may be suited for existence proofs of granularity of the A-to-D conversion to adequately account for such graded effects outlined above (see Churchland 1992 for some early possible examples).

Symbolic dynamics offers a playground in which this conceptual problem can be formally explored. Given a dynamical system living in a state space of m dimensions, a set of stable or meta-stable attractors can be explored via simulation.

Like perceptual processes, this dynamical system can feed into a separate system, described in a variety of ways (e.g. a Turing machine, or finite-state machine), that receives symbolic input via threshold crossing in a partition carving that state space. This collapse involves maximal loss: from m dimensions into 1 dimension of Q possible states defining the partition. Of course, for simple systems such as the logistic map (where $m=1$), this collapse can still carry the original dynamics, and entail some interesting computational properties. As for higher-dimensional state spaces, there seem to be two ways that models of this kind might begin to approach existence proofs for symbolization of a perceptual space. On the simplest side, one might explore resultant algorithmic processing on symbolic output of just one dimension with Q regions of one partitioning. These states may be numerous, or *refined*, enough to carry some echoes of the original space. Another possibility is to consider collapsing the state space of m dimensions into more than one partition. The m-dimensional system may be collapsed onto an n-tuple of symbols, each element of which is the output from some separate partition that uniquely carves the state space. Exploration of this system would involve subsequent A-to-D/D-to-D conversions, permitting sequenced levels of granularity in the various stages of perceptual-cognitive processing (that is, without considering feedback projections).

There exist a number of 'useful' computational models performing symbolization of this sort. For example, the decision processes of Pollack's (1990) dynamical recognizer and the Hopfield network of Balkenius and Gärdenfors (1991) that implements non-monotonic logics are two relatively early models. Recently, Tabor has specifically addressed the learning and processing of formal languages by such systems (Tabor 2000, 2001), along with beim Graben and colleagues' sophisticated analyses (beim Graben *et al.* 2004). These are just a few of the enticing invitations to employ symbolic dynamics in a way we recommend here: devising existence proofs relating the stages of continuous-discrete transitions in a simplified problem space akin to cognitive processes and their output. Though only promissory at this point, symbolic dynamics may make it possible to reconcile both the dynamic and discrete descriptions of the states and processes underlying cognition.

5. Conclusion

We do not have full privileges in our access to the ontic states of our mind/brain. An inevitable fact about higher-order cognitive theories is that they are descriptive at a very coarse level—it is currently an intractable problem to specify, even partially, the dynamics underlying neural computation in a cognitive task of any nontrivial complexity (e.g. Uttal 2001). Nevertheless, for the majority of cognitive scientists, this daunting state of affairs does not invalidate proposals for structures and processes of our cognitive system. In this article, we have limned the surface of a framework within which competing theoretical accounts of representational structures and processes may have equal opportunity to contribute to our understanding of cognition. Symbolic dynamic investigations of idealized problem spaces may provide a common arena for exploring the interplay between continuous-distributed and discrete-symbolic representational accounts. Moreover, as a framework for further discussion, it may help both representational formats overcome the limitations of time-course irrelevant descriptions of cognition. Given a set of input information and informational goals, symbolic dynamics offers both informational and temporal insight into the transition from continuous perceptual trajectories into more or less fine-grained discretized states for higher cognitive processes. There will be, of course, some conceptual and technical obstacles in the way ahead, and we have considered a number of these above.

Given the deep epistemic problems experienced by all theories of cognition, and the complexity of the brain onto which they are imposed, it seems we might forever be confined to epistemic descriptions of the ontic states of our cognitive system. Regarding the state of the art in cognitive science, the dispute between two such families of description, discrete-symbolic and distributed-continuous, thus seems just as likely to be evenly conciliated than for one or the other to win permanent prominence. This article offers some further considerations of symbolic dynamics to contribute to ongoing debate (see also, for example, Goertzel 1998, beim Graben 2004). Mathematization of simplified problem spaces, such as perceptual categorization or computation in 'chaotic itinerancy' (Tsuda 2001), may be the route to a formal terrain permitting cohabitation of both kinds of theoretical constructs—or, at least, a mutually supportive arena in which they can have a fair fight.

It is perhaps a striking illusion, at the physical level, that there exist discrete states of the mind/brain. This at least seems to be the case if you grant that the substrate is in constant motion, like Heraclitus's river. The illusion is nevertheless difficult to overcome, because our phenomenology seems to be in an inescapable embrace with experiences that have strict boundaries. At the epistemic level, rather than the phenomenological level, it may be inevitable that boundaries need to be placed around neurophysiological complexity to construct sufficiently explanatory, and tractable, theories of cognitive processes. This becomes evermore troublesome when one considers functional redundancy and feedback loops within the substrate, as well as between it and its environment. So, while theoretical debate may continue concerning whether the mind is a system that imposes boundaries on a continuous information flow, symbolic dynamics may offer a mathematical terrain in which these boundaries can be rigorously explored.

References

R. Abrams and D. Balota, "Mental chronometry: Beyond reaction time", *Psychological Science*, 2, pp. 153–157, 1991.

P.D. Allopenna, J.S. Magnuson and M.K. Tanenhaus, "Tracking the time course of spoken word recognition using eye movements: Evidence for continuous mapping models", *Journal of Memory and Language*, 38, pp. 419–439, 1998.

G. Altmann and M. Steedman, "Interaction with context during human sentence processing", *Cognition*, 30, pp. 191–238, 1988.

J.A. Anderson, J.W. Silverstein, S.A. Ritz and R.S. Jones, "Distinctive features, categorical perception, and probability learning: Some applications for a neural model", *Psychological Review*, 84, pp. 413–451, 1977.

M. Andrews, Unpublished PhD dissertation, Cornell University, 2003.

R.C. Atkinson and R.M. Shiffrin, "Human memory: A proposed system and its control processes", in *The Psychology of Learning and Motivation: Advances in Research and Theory*, vol. 2, K.W. Spence and J. T. Spence, Eds, New York: Academic Press, 1968, pp. 89–195.

H. Atmanspacher, "Ontic and epistemic descriptions of chaotic systems", in *Computing Anticipatory Systems: CASYS 99*, D.M. Dubois, Ed., Berlin: Springer, 2000, pp. 465–478.

C. Balkenius and P. Gärdenfors, Non-monotonic inferences in neural networks, in *Principles of Knowledge Representation and Reasoning: Proceedings of the Second International Conference*, J.A. Allen, R. Fikes and E. Sandewall, Eds, San Mateo, CA: Morgan Kaufman, 1991, pp. 32–39.

F.C. Bartlett, *Remembering*, Cambridge: Cambridge University Press, 1932.

P. beim Graben, "Incompatible implementations of physical symbol systems", *Mind and Matter*, 2, pp. 29–51, 2004.

P. beim Graben and H. Atmanspacher, "Complementarity in classical dynamical systems", *Chaotic Dynamics*, arXiv: nlin.CD/0407046, 2004.

P. beim Graben, B. Jurish, D. Saddy and S. Frisch, "Language processing by dynamical systems", *International Journal of Bifurcation and Chaos*, 14, pp. 599–621, 2004.

P. Bloom and L. Markson, "Intentionality and analogy in children's naming of pictorial representations", *Psychological Science*, 9, pp. 200–204, 1998.

E. Bollt, T. Stanford, Y. Lai and K. Zyczkowski, "Validity of threshold crossing analysis of symbolic dynamics from chaotic time series", *Physical Review Letters*, 85, pp. 3524–3527, 2000.

E. Bollt, T. Stanford, Y. Lai and K. Zyczkowski, "What symbolic dynamics do we get with A misplaced partition? On the validity of threshold crossings analysis of chaotic time-series", *Physica D*, 154, pp. 259–286, 2001.

K. Britten, M. Shadlen, W. Newsome and J. Movshon, "The analysis of visual motion: a comparison of neuronal and psychophysical performance", *Journal of Neuroscience*, 12, pp. 4745–4767, 1992.

N. Chater, M. H. Christiansen and F. Reali, "Is coevolution of language and language genes possible?", paper presented at the *Fifth International Conference on Language Evolution*. Leipzig, Germany, 2004.

P. Churchland, *A Neurocomputational Perspective: the Nature of Mind and the Structure of Science*, Cambridge, MA: MIT Press, 1992.

A. Clark, *Mindware: An Introduction to the Philosophy of Cognitive Science*, Oxford: Oxford University Press, 2001.

H. Clark, *Arenas of Language Use*, Chicago, IL: University of Chicago Press, 1993.

A. Cleeremans, D. Servan-Schreiber and J. McClelland, "Finite state automata and simple recurrent networks", *Neural Computation*, 1, pp. 372–381, 1989.

M. Coltheart, K. Rastle, C. Perry, R. Langdon and J. Ziegler, "DRC: A Dual Route Cascaded model of visual word recognition and reading aloud", *Psychological Review*, 108, pp. 204–256, 2001.

L. Cooper and R. Shepard, "The time required to prepare for a rotated stimulus", *Memory and Cognition*, 1, pp. 246–250, 1973.

S. Crain and M. Steedman, "On not being led up the garden path", in D. Dowty, L. Karttunen and A. Zwicky, Eds, *Natural Language Parsing*, Cambridge, MA: Cambridge, 1985.

J.P. Crutchfield, "The calculi of emergence: Computation, dynamics, and induction", *Physica D*, *75*, pp. 11–54, 1994.

J.P. Crutchfield, "Dynamical embodiments of computation in cognitive processes", *Behavioral and Brain Sciences*, 21, pp. 635–637, 1998.

J.P. Crutchfield and N. Packard, "Symbolic dynamics of one-dimensional maps: Entropies, finite precision, and noise", *International Journal of Theoretical Physics*, 21, pp. 433–466, 1982.

J.P. Crutchfield and K. Young, "Inferring statistical complexity", *Physical Review Letters*, 63, pp. 105–108, 1989.

J.P. Crutchfield and K. Young, "Computation at the onset of chaos", in *Entropy, Complexity, and Physics of Information*, W. Zurek, Ed., Reading, MA: Addison-Wesley, 1990, vol. 8, pp. 223–269.

R.I. Damper and S.R. Harnad, "Neural network models of categorical perception", *Perception & Psychophysics*, 62, pp. 843–867, 2000.

J.H. Danovitch and F.C. Keil, "Should you ask a fisherman or a biologist?: Developmental shifts in ways of clustering knowledge", *Child Development*, 5, pp. 918–931, 2004.

R. Davidchack, Y. Lai, A. Klebanoff and E. Bollt, "Towards complete detection of unstable periodic orbits in chaotic systems", *Physics Letters A*, 287, pp. 99–104, 2001.

C. Daw, C. Finney and E. Tracy, "A review of symbolic analysis of experimental data", *Review of Scientific Instruments*, 74, pp. 916–930, 2003.

R. Desimone and J. Duncan, "Neural mechanisms of selective visual attention", *Annual Review of Neuroscience*, 18, pp. 193–222, 1995.

D. Devaney, *An Introduction to Chaotic Dynamical Systems*, 2nd edn, Boulder, CO: Westview Press, 2003.

E. Dietrich and A.B. Markman, "Discrete thoughts: Why cognition must use discrete representations", *Mind and Language*, 18, pp. 95–119, 2003.

J. Duncan and G. Humphreys, "Visual search and stimulus similarity", *Psychological Review*, 96, pp. 433–458, 1989.

J. Elman, E. Bates, M. Johnson, A. Karmiloff-Smith, D. Parisi and K. Plunkett, *Rethinking Innateness*, MIT Press: Cambridge, MA, 1996.

J.L. Elman, "Finding structure in time", *Cognitive Science*, 14, pp. 179–211, 1990.

R. Engbert, C. Scheffczyk, R.T. Krampe, J. Kurths and R. Kliegl, "Symbolic dynamics of bimanual production of polyrhythms", in *Nonlinear analysis of physiological data*, H. Kantz, J. Kurths and G. Mayer-Kress, Eds, Berlin: Springer, 1998.

L. Feigenson and S. Carey, "Tracking individuals via object files: Evidence from infants' manual search", *Developmental Science*, 6, pp. 568–584, 2003.

A. Fernald, D. Swingley and J. Pinto, "When half a word is enough: Infants can recognize spoken words using partial phonetic information", *Child Development*, 72, pp. 1003–1015, 2001.

J.A. Fodor, *Modularity of Mind*, Cambridge, MA: Bradford, 1983.

J.A. Fodor, *Psychosemantics*, Cambridge, MA: MIT Press, 1987.

J.A. Fodor and Z.W. Pylyshyn, "Connectionism and cognitive architecture: a critical analysis", *Cognition*, 28, pp. 3–71, 1988.

L. Frazier, "Constraint satisfaction as a theory of sentence processing", *Journal of Psycholinguistic Research*, 24, pp. 437–468, 1995.

C. Gauker, *Words without Meaning*, Cambridge, MA: MIT Press, 2003.

A. Georgopoulos, "Motor cortex and cognitive processing", in *The Cognitive Neurosciences*, M. Gazzaniga, Ed., Cambridge, MA: The MIT Press, 1995, pp. 507–517.

A. Georgopoulos, J. Lurito, M. Petrides, A. Schwartz and J. Massey, "Mental rotation of the neuronal population vector, *Science*, 243, pp. 234–236, 1989.

B. Goertzel, "Learning the language of mind: Symbolic dynamics for modeling adaptive behaviour", in *Models of Action: Mechanisms for Adaptive Behaviour*, C. Wynne and J. Staddon, Eds, Mahwah, NJ: Lawrence Erlbaum Associates, 1998, pp. 1–27.

J. Gold and M. Shadlen, "Representation of a perceptual decision in developing oculomotor commands", *Nature*, 404, pp. 390–394, 2000.

J. Gold and M. Shadlen, "Neural computations that underlie decisions about sensory stimuli", *Trends in Cognitive Sciences*, 5, pp. 10–16, 2001.

R. Hadley, "Systematicity revisited", *Mind and Language*, 9, pp. 431–444, 1994.

M.W. Harm and M.S. Seidenberg, "Phonology, reading acquisition, and dyslexia: Insights from connectionist models", *Psychological Review*, 106, pp. 491–528, 1999.

S. Harnad, Ed., *Categorical Perception: The Groundwork of Cognition*, New York: Cambridge University Press, 1987.

P. Jolicoeur, "The time to name disoriented natural objects", *Memory and Cognition*, 13, pp. 289–303, 1985.

J.A.S. Kelso, *Dynamic Patterns: The Self-organization of Brain and Behaviour*, Cambridge, MA: MIT Press, 1995.

M.B. Kennel and M. Buhl, "Estimating good discrete partitions from observed data: symbolic false nearest neighbors", *Physical Review Letters*, 91, 084102, 2003.

P. Knoeferle, M.W. Crocker, C. Scheepers and M.J. Pickering, "The influence of the immediate visual context on incremental thematic role-assignment: Evidence from eye-movements in depicted events", *Cognition*, 95, pp. 95–127, 2004.

J. Kurths, A. Voss, P. Saparin, A. Witt, H. Kleiner and N. Wessel, "Quantitative analysis of heart rate variability, *Chaos*, 5, pp. 88–94, 1995.

J. Lachter and T. Bever, "The relation between linguistic structure and associative theories of language learning: A constructive critique of some connectionist learning models", *Cognition*, 28, pp. 195–247, 1988.

A. Liberman, "On finding that speech is special", *American Psychologist*, 37, pp. 148–167, 1982.

A. Liberman, K. Harris, H. Hoffman and B. Griffith, "The discrimination of speech sounds within and across phoneme boundaries", *Journal of Experimental Psychology*, 53, pp. 358–368, 1957.

M. MacDonald, N. Pearlmutter and M. Seidenberg, "The lexical nature of syntactic ambiguity resolution", *Psychological Review*, 101, pp. 676–703, 1994.

J. Magnuson, M. Tanenhaus, R. Aslin and D. Dahan, "The time course of spoken word learning and recognition: Studies with artificial lexicons", *Journal of Experimental Psychology: General*, 132, pp. 202–227, 2003.

G.F. Marcus, *The Algebraic Mind: Integrating Connectionism and Cognitive Science*, Cambridge, MA: MIT Press, 2001.

V. Marian and M. Spivey, "Competing activation in bilingual language processing: Within-and between-language competition", *Bilingualism: Language and Cognition*, 6, pp. 97–115, 2003.

A.B. Markman and E. Dietrich, "In defense of representation", *Cognitive Psychology*, 40, pp. 138–171, 2000.

W. Marslen-Wilson, "Functional parallelism in spoken word recognition", *Cognition*, 25, pp. 71–102, 1987.

J. McClelland and J. Elman, "The TRACE model of speech perception", *Cognitive Psychology*, 18, pp. 1–86, 1986.

B. McMurray, M. Tanenhaus and R. Aslin, "Gradient effects of within-category phonetic variation on lexical access", *Cognition*, 86, pp. B33–B42, 2002.

B. McMurray, M. Tanenhaus, R. Aslin and M. Spivey, "Probabilistic constraint satisfaction at the lexical/phonetic interface: Evidence for gradient effects of within-category VOT on lexical access", *Journal of Psycholinguistic Research*, 32, pp. 77–97, 2003.

R. McMurray and M. Spivey, "The categorical perception of consonants: The interaction of learning and processing", in *Proceedings of the Chicago Linguistic Society*, 1999.

K. McRae, M.J. Spivey-Knowlton and M.K. Tanenhaus, "Modeling the influence of thematic fit (and other constraints) in on-line sentence comprehension", *Journal of Memory and Language*, 38, pp. 283–312, 1998.

D.L. Medin, "Concepts and concept structure", *American Psychologist*, 44, pp. 1469–1481, 1989.

A. Medio and M. Lines, *Nonlinear Dynamics: A Primer*, Cambridge: Cambridge University Press, 2001.

P.E. Meehl, "The power of quantitative thinking", address at APS, May 1998. Available online: http://www.tc.umn.edu/~pemeehl

J. Miller, "Discrete versus continuous stage models of human information processing: In search of partial output", *Journal of Experimental Psychology: Human Perception and Performance*, 8, pp. 273–296, 1982.

E.F. Moore, "Gedanken-experiments on sequential machines", in *Automata Studies*, C.E. Shannon and J. McCarthy, Eds, Princeton, NJ: Princeton University Press, 1956, pp. 129–153.

M. Morse and G.A. Hedlund, "Symbolic dynamics", *American Journal of Mathematics*, 60, pp. 815–866, 1938.

I. Myung and S. Zhang, "Toward a method of selecting among computational models of cognition", *Psychological Review*, 109, pp. 472–491, 2002.

K. Nakayama and J. Joseph, "Attention, pattern recognition, and pop-out visual search", in *The Attentive Brain*, R. Parasuraman, Ed., Cambridge, MA: MIT Press, 1998, pp. 279–298.

U. Neisser, *Cognition and Reality*, San Francisco: Freeman, 1976.

A. Newell and H. Simon, "Computer science as empirical inquiry: Symbols and search", *Communications of the ACM*, 19, pp. 113–126, 1976.

E. Olds, W. Cowan and P. Jolicoeur, "The time-course of pop-out search", *Vision Research*, 40, pp. 891–912, 2000.

D. Perrett, M. Oram and E. Ashbridge, "Evidence accumulation in cell populations responsive to faces: An account of generalization of recognition without mental transformations", *Cognition*, 67, pp. 111–145, 1998.

S. Pinker, *How the Mind Works*, New York: Norton, 1997.

D. Pisoni and J. Tash, "Reaction times to comparisons within and across phonetic categories", *Perception and Psychophysics*, 15, pp. 285–290, 1974.

J.B. Pollack, "The induction of dynamical recognizers", *Machine Learning*, 7, pp. 227–252, 1991.

R. Port and T. van Gelder, *Mind as Motion: Explorations in the Dynamics of Cognition*, Cambridge, MA: MIT Press, 1995.

W.V.O. Quine, *Word and object*, Cambridge, MA: MIT Press, 1960.

M. Redington and N. Chater, "Transfer in artificial grammar learning: A re-evaluation", *Journal of Experimental Psychology: General*, 125, pp. 123–138, 1996.

J. Reynolds and R. Desimone, "Neural mechanisms of attentional selection", in *Visual Attention and Cortical Circuits*, J. Braun, C. Koch and J. Davis, Eds, Cambridge, MA: The MIT Press, 2001, pp. 121–135.

L. Rips, E. Shoben and E. Smith, "Semantic distance and the verification of semantic relations", *Journal of Verbal Learning and Verbal Behavior*, 12, pp. 1–20, 1973.

S. Roberts and H. Pashler, "How persuasive is a good fit? A comment on theory testing", *Psychological Review*, 107, pp. 358–367, 2000.

P. Rodriguez, J. Wiles and J. Elman, "A recurrent neural network that learns to count", *Connection Science: Journal of Neural Computing, Artificial Intelligence and Cognitive Research*, 11, pp. 5–40, 1999.

E. Rolls and M. Tovee, "Sparseness of the neuronal representation of stimuli in the primate temporal visual cortex", *Journal of Neurophysiology*, 73, pp. 713–726, 1995.

E. Rosch, "Cognitive representations of semantic categories", *Journal of Experimental Psychology*, 104, pp. 192–233, 1975.

D. Rosen, "An argument for the logical notion of a memory trace", *Philosophy of Science*, 42, pp. 1–10, 1975.

J. Schall, "Decision making: From sensory evidence to a motor command", *Current Biology*, 10, pp. R404–R406, 2000.

H. Schütze, "A connectionist model of verb subcategorization", *Proceedings of the Sixteenth Annual Conference of the Cognitive Science Society*, Hillsdale, NJ: Erlbaum, 1994, pp. 784–788.

C. Shalizi, "Notebooks: Dynamics in cognitive science", 2004a. Available online: http://cscs.umich.edu/~crshalizi/notebooks/dynamics-cognition.html

C. Shalizi, Notebooks: Symbolic dynamics, 2004b. Available online: http://cscs.umich.edu/~crshalizi/notebooks/symbolic-dynamics.html

C. Shalizi and D.J. Albers, Symbolic dynamics for discrete adaptive games, submitted.

R. Shepard and J. Metzler, "Mental rotation of three dimensional objects", *Science*, 171, 701–703, 1971.

P. Smolensky, ''On the proper treatment of connectionism'', *Behavioral and Brain Sciences*, 11, pp. 1–74, 1988.
M. Spivey and V. Marian, ''Crosstalk between native and second languages: Partial activation of an irrelevant lexicon'', *Psychological Science*, 10, pp. 281–284, 1999.
M. Spivey and M. Tanenhaus, ''Syntactic ambiguity resolution in discourse: Modeling the effects of referential context and lexical frequency'', *Journal of Experimental Psychology: Learning, Memory and Cognition*, 24, pp. 1521–1543, 1998.
M. Spivey, M. Grosjean and G. Knoblich, ''Continuous attraction toward phonological competitors'', *Proceedings of the National Academy of Sciences*, 102, pp. 10393–10398, 2005.
M. Spivey, M. Tanenhaus, K. Eberhard and J. Sedivy, ''Eye movements and spoken language comprehension: Effects of visual context on syntactic ambiguity resolution'', *Cognitive Psychology*, 45, pp. 447–481, 2002.
M. Spivey, M. Tyler, K. Eberhard and M. Tanenhaus, ''Linguistically mediated visual search'', *Psychological Science*, 12, pp. 282–286, 2001.
M.J. Spivey and R. Dale, ''The continuity of mind: Toward a dynamical account of cognition'', *Psychology of Learning and Motivation*, 45, pp. 85–142, 2005.
M. Spivey-Knowlton and J. Saffran, ''Inducing a grammar without an explicit teacher: Incremental distributed prediction feedback'', in *Proceedings of the 17th Annual Conference of the Cognitive Science Society*. Hillsdale, NJ: Erlbaum, 1995.
M. Spivey-Knowlton, J. Sedivy, K. Eberhard and M. Tanenhaus, ''Psycholinguistic study of the interaction between language and vision'', in *Proceedings of the 12th National Conference on Artificial Intelligence: Workshop on the Integration of Natural Language and Vision Processing*, Menlo Park, CA: AAAI Press, 1994, pp. 189–192.
S. Stevenson, ''Competition and recency in a hybrid network model of syntactic disambiguation'', *Journal of Psycholinguistic Research*, 23, pp. 295–322, 1994.
W. Tabor, ''Fractal encoding of context-free grammars in connectionist networks'', *Expert Systems: The International Journal of Knowledge Engineering and Neural Networks*, 17, pp. 41–56, 2000.
W. Tabor, ''Dynamical assessment of symbolic processes with backprop nets'', in *Proceedings of the INNS-IEEE International Joint Conference on Artificial Neural Networks*, 2001.
W. Tabor, The value of symbolic computation, *Ecological Psychology*, 14, 21–51, 2002.
W. Tabor and M. Tanenhaus, ''Dynamical models of sentence processing'', *Cognitive Science*, 23, pp. 491–515, 1999.
M. Tanenhaus, M. Spivey-Knowlton, K. Eberhard and J. Sedivy, 1995, ''Integration of visual and linguistic information during spoken language comprehension'', *Science*, 268, pp. 1632–1634, 2001.
J.B. Tenenbaum and T.L. Griffiths, ''Generalization, similarity, and Bayesian inference'', *Behavioral and Brain Sciences*, 24, pp. 629–664, 2001.
E. Thelen and L. Smith, *A Dynamic Systems Approach to the Development of Cognition and Action*, Cambridge, MA: MIT Press, 1994.
A. Treisman, ''Features and objects: The Fourteenth Bartlett Memorial Lecture'', *Quarterly Journal of Experimental Psychology*, 40A, pp. 201–237, 1988.
A. Treisman and G. Gelade, ''A feature integration theory of attention'', *Cognitive Psychology*, 12, pp. 97–136, 1980.
I. Tsuda, 2001, Toward an interpretation of dynamic neural activity in terms of chaotic dynamical systems, *Behavioral and Brain Sciences*, 25, pp. 793–848, 1980.
W. Uttal, *The New Phrenology: The Limits of Localizing Cognitive Processes in the Brain*, Cambridge, MA: The MIT Press, 2001.
T. Van Gelder, The dynamical hypothesis in cognitive science, *Behavioral and Brain Sciences*, 21, pp. 615–665, 1998.
L. Ward, *Dynamical Cognitive Science*, Cambridge, MA: MIT Press, 2002.
J. Weizenbaum, ''ELIZA A computer program for the study of natural language communication between man and machine'', *Communications of the ACM*, 9, pp. 36–45, 1966.
S.G. Williams, ''Introduction to symbolic dynamics'', in *Proceedings of Symposia in Applied Mathematics: Symbolic Dynamics and its Applications*, S.G. Williams, Ed., Providence, RI: American Mathematical Society, 2002, vol. 60, pp. 1–13.
T. Winograd, *Understanding Natural Language*, Oxford: Academic Press, 1972.
J. Wolfe, ''What can 1 million trials tell us about visual search?'', *Psychological Science*, 9, pp. 33–39, 1998.