

Continuous Dynamics in Real-Time Cognition

Michael J. Spivey¹ and Rick Dale²

¹Cornell University and ²University of Memphis

ABSTRACT—*Real-time cognition is best described not as a sequence of logical operations performed on discrete symbols but as a continuously changing pattern of neuronal activity. The continuity in these dynamics indicates that, in between describable states of mind, mental activity does not lend itself to the linguistic labels relied on by much of psychology. We discuss eye-tracking and mouse-tracking evidence for this temporal continuity and provide geometric visualizations of mental activity, depicting it as a continuous trajectory through a state space (a multidimensional space in which locations correspond to mental states). When the state of the system travels toward a frequently visited region of that space, the destination may constitute recognition of a particular word or a particular object; but on the way there, the majority of the mental trajectory is in intermediate regions of that space, revealing graded mixtures of mental states.*

KEYWORDS—*cognition; continuity; dynamical systems; eye movements*

When introspecting on one's thought processes, one often feels as though a discrete concept is considered all by itself, followed by another discrete concept, followed by another. Such a sequence of distinct, nonoverlapping mental constituents is consistent with the string of logical symbols on the tape of a Turing machine. In the 1950s, computing theory devised the theoretical construct of a universal Turing machine, on which the computation of any algorithm could be implemented by moving the tape forward or backward so that a programmed tape head could read the symbols. Such a computing system can, by modifying the discrete symbols on the tape, exhibit a variety of intelligent behaviors that resemble our own cognitive skills, such as performing arithmetic or playing chess. The theoretical construct of a universal Turing machine was quickly imported into psy-

chology as the mathematical foundation underlying the information-processing approach to cognitive psychology (Pylyshyn, 1984).

We draw an analogy from this series of mental symbols to a series of stitches on the hem of a curtain. On the surface, you see a half-inch line of thread, followed by a half-inch gap, followed by another half-inch line of thread, and so on. This sequence of seemingly nonoverlapping stitches could cause an observer to conclude that each thread starts at one end of a stitch and stops at the other end. However, a deeper inspection of the fabric will, of course, reveal that there is actually one continuous thread, portions of which are above the surface of the fabric and portions of which are below the surface. Likewise, with the mind, one's thoughts often appear as though they are composed of "individuated elements"—perhaps because people tend to use the discrete results of their (otherwise continuous) actions, such as carefully chosen words or a deliberated chess move, as informal evidence for those thoughts. However, when you look more closely, especially with continuous online experimental measures such as eye-tracking and computer-mouse-tracking (instead of outcome-based measures such as reaction time and accuracy), it is possible to see that mental activity is also being conducted in between those seemingly discrete thoughts. Thus, we argue that cognition is best analyzed as a continuously dynamic biological process, not as a staccato series of abstract computer-like symbols.

In this brief article, we hope to make progress toward dispelling the illusion of sequential discrete thoughts, or symbol strings, and toward revealing the continuous thread of thought that weaves its way through the fabric of the mind. We employ a description of mental contents that treats cognition as living in a high-dimensional state space, in which locations correspond to mental states (and proximity among those locations is equivalent to similarity among those mental states). Different regions of that space exert an attractive force, with the strength of that attraction gradually falling off as a function of distance (a bit like gravitational forces among planets in a solar system). These regions are called attractor basins, and they correspond to patterns of neuronal activity that are elicited when someone is recognizing

Address correspondence to Michael Spivey, Department of Psychology, Cornell University, Ithaca, NY 14853; e-mail: spivey@cornell.edu.

	CDIR	437	B	Dispatch: 17.10.06	Journal: CDIR	CE: Blackwell
	Journal Name	Manuscript No.		Author Received:	No. of pages: 5	PE: Sarvanan/ananth

a spoken word or a visual object. Mental activity is seen as a continuous trajectory through this state space, traveling to word-recognition attractor basins, object-recognition attractor basins, semantic-category attractor basins, and everywhere in between.

CONTINUITY IN COGNITION

Although many researchers examine the continuous temporal dynamics of cognitive development, tracking the time course of weeks to months (e.g., Elman et al., 1996; Thelen & Smith, 1994), there is comparatively little research on the continuous temporal dynamics of real-time cognition, tracking the time course of dozens to hundreds of milliseconds. Just as dynamical accounts have been discovering that the transitions between putative stages of cognitive development are marked by continuous nonlinear change, not discrete jumps, so does a dynamical analysis of real-time cognition reveal that transitions between “mental states” consist of continuous nonlinear change, not discrete jumps.

In fact, one of the oft-touted examples of sudden cognitive transitions, the “Aha!” effect in insight problem solving, actually has a gradual quality to it. Insight problems are labeled as such because they tend to induce a period of frustration, in which the solution has not been found, new ideas are not forthcoming, and there is no hint of being close to the solution. This is known as the “impasse.” Then, seemingly out of nowhere, the correct solution arrives in an instantaneous inexplicable epiphany. This description fits nicely with the discrete state transitions associated with the computer metaphor of the mind (Pylyshyn, 1984). However, Bowden and Beeman (1998) demonstrated that partial activation of an insight solution was detectable in the form of priming. They had participants try to solve insight-based problems known as compound remote associates (e.g., “What one word makes a common word pair with each of the words *opera*, *hand*, and *dish*?”), and followed each problem with a word-naming task. Even on trials in which participants failed to find a solution, these participants were faster to read out loud the correct solution word (e.g., “soap”) compared to a neutral control word.¹ Moreover, electroencephalography data reveal that, more than a second before the gamma-frequency burst associated with the conscious “Aha!” experience, alpha-frequency activity emerges in the right posterior parietal cortex (Jung-Beeman et al., 2004). Thus, well before the insight solution is discovered, and even when it is not discovered, the mind is nonetheless preparing for and “getting close to” the solution.

Another popular example of putatively discrete processing is categorical speech perception. When two speech sounds, such as “bah” and “pah,” are synthetically blended into a continuum of sounds, the sounds on one half of the continuum are perceived as indistinguishable “bah” sounds and the sounds on the other

half as indistinguishable “pah” sounds. On the surface, this phenomenon seems consistent with the information-processing framework’s assumption that perceptual categories are discrete logical forms involving no graded variation. However, when participants’ eye movements were recorded during the speech-identification task (with a headband-mounted eye tracker that recorded where participants looked while they mouse-clicked either the “ba” or “pa” response boxes), speech sounds near the category boundary elicited eye movements that initially vacillated between the two response options (McMurray et al., 2003). In fact, the proportion of time looking at a response option (e.g., “pah”) corresponded with the typicality of the speech sound. Thus, the gradations in the speech stimuli were not instantaneously discarded during perception; instead they influenced a dynamic competition between response options.

This competition process was a crucial property of Anderson and colleagues’ (Anderson, Silverstein, Ritz, & Jones, 1977) attractor-network simulation of categorical speech perception. When the system is exposed to a potentially ambiguous stimulus, the internal pattern of neuronal activity will be partially consistent with multiple different population codes (groups of neurons that cooperate with one another to form a stable percept or concept). This pattern of neuronal activity can be mathematically described as a location in a state space, of which the dimensions are the activation levels of neurons (see Rolls & Tovee, 1995). Deciding on the identity of a stimulus is thus a pattern-completion process whereby the internal state moves toward some attractor basins (regions of “gravitational pull”), away from others, and finally settles on a unique point attractor (stable location).

CONTINUITY IN SPOKEN-WORD RECOGNITION

We (Spivey and Dale, 2004) review a variety of cases in which perception and cognition involve a temporally continuous competition process rather than stage-by-stage computation. These include examples from visual cognition, such as the neurophysiology of object recognition and of perceptual decisions, and from attractor-network simulations of visual-search processes. A particularly compelling case for this continuity comes from studying eye movements in response to spoken-word recognition.

Tanenhaus, Spivey-Knowlton, Eberhard, and Sedivy (1995) report evidence that partway through hearing a spoken word, visual objects in the environment attract eye movements when their names match the first few phonemes of the speech signal. For example, when presented with a display of real objects and instructed to “pick up the candy,” participants occasionally looked first at a candle and then at the candy. When the proportion of eye movements is plotted over hundreds of milliseconds, looks at the target object (e.g., the candy) and the competitor object (e.g., the candle) initially rise in tandem. A few hundred milliseconds later, looks at the competitor object

¹Interestingly, this only occurred with words presented in the left visual field and thus initially processed by the right hemisphere.




Fig. 1. Path of computer-mouse movement (green circles) when an individual was instructed to “click the candy” on a computer screen. The upper plane depicts the x- and y-axes on a computer screen, with pictures of a candle and a piece of candy at the top. The curvature of the trajectory from the starting point at the bottom reflects a graded spatial attraction toward the competitor object (the candle) due to its phonological similarity. The lower plane depicts a simplified version of an attractor manifold in a low-dimensional state space (representing the possible patterns of neuronal activation), with two attractor basins for the lexical alternatives and a mental trajectory (circle shadows) that initially gravitates toward the midpoint of the two attractors and then curves into the correct one. (Although hand movements are slower than eye movements, they are nonetheless tightly yoked to the dynamic activity of population codes in the premotor cortex; cf. Cisek & Kalaska, 2005.)

decline, and the proportion of looks at the target object approaches 1.0. This pattern of simultaneous partial activation of lexical alternatives, ensuing competition, and eventual resolution is consistent with dynamical models of spoken-word recognition.

Instead of recording ballistic eye movements, which either look at the competitor object or not on a given trial, recent work has made use of continuous manual measures to more richly flesh out continuous processing in time. Spivey, Grosjean, and Knoblich (2005) recorded the streaming x-y coordinates of continuous computer-mouse movements, which can reveal a graded spatial attraction toward the competitor object even on an individual trial. When participants saw a piece of candy and a candle on the computer screen and were instructed to “click the candy,” the trajectory of their mouse movements tended to exhibit spatial attraction to both objects, thus initially gravitating toward their midpoint before eventually curving toward the target object.² The green circles in Figure 1 plot the actual data from an individual trial. Underneath the mouse trajectory, Figure 1 depicts an idealized attractor landscape, where one might imagine the mental state as a marble that smoothly rolls downhill (as shown by the circle shadows) and eventually settles into one of the basins. The constraints on the mouse-movement task es-

entially force those neuronal state-space dynamics to be expressed in the two-dimensional action space of the computer screen, allowing us to witness a low-dimensional record of the high-dimensional mental trajectory. The key observation is that the majority of the trajectory’s time is spent in intermediate regions of state space (i.e., ambiguous patterns of neural activity) that are partially consistent with multiple lexical representations. That is, on the way toward distinctly recognizing a unique word, the mind is entertaining a continuously evolving *mélange* of words.

CONTINUITY IN SEMANTIC CATEGORIZATION

The continuity of mind is further revealed at the level of semantic processing in categorization. McRae (2004) describes an attractor-network model that encodes vectors of features corresponding to the semantic properties of word forms. For example, the word form *whale* would have a large semantic vector with ones for present features (e.g., “has fins”) and zeros for absent features (e.g., “flies”). The network simulates accessing these semantic properties when a word is presented to it, and one can track its processing over time as it settles into a semantic interpretation—reaching a stable point attractor in the state space of semantic features. McRae’s model captures a wide variety of results in semantic memory and categorization, including a

²In the control condition, where the name of the alternative object was not similar to the spoken word, the trajectory was significantly straighter.

detailed simulation of the temporal dynamics of semantic processing.

McRae's (2004) attractor-network account of semantic memory suggests that richly interconnected featural representations may underlie semantic categorization. As in spoken-word recognition, where phonemic overlap draws the eyes and hands toward competitors, featural overlap in semantic categorization does so as well. A recent study by Dale, Kehoe, and Spivey (in press) demonstrates this continuous nonlinear processing in semantic categorization. Participants categorized an animal word (e.g., *cat*), presented at the bottom of the screen, by mouse-clicking one of two taxonomic classes (e.g., the words *mammal* and *fish*) in the upper corners of the screen. Critical trials involved atypical animals (e.g., *whale*) and included an incorrect competitor category that had considerable overlap in terms of semantic and visual features (e.g., *fish*). Although participants reliably clicked the appropriate category (e.g., *mammal*), these mouse-movement trajectories exhibited substantially more spatial attraction toward the competitor category than in the typical-animal trials (where, for example, *cat* was categorized as a mammal).³ The incorrect response box in the atypical-animal trials acted as a competitive attractor basin, because of featural overlap, generating some gravitational pull in the dynamics of motor output. Thus, the categorization process guiding motor movement had not fully "discretized" its decision, and the partial activation of multiple interpretations was continuously cascading into the motor-execution phase of the task.

THE EFFABLE AND THE INEFFABLE

It is important to note that, in real life, the environment generally does not deliver an isolated stimulus and then patiently wait for one to emit an isolated response. Therefore, conceiving of cognition as a trajectory that starts in a neutral location in state space, begins moving toward some attractors, and then settles motionlessly on one of them, is only a small beginning of the story. Flowing arrays of stimulation, such as hearing a sentence comprising several words or watching people move about in a scene, elicit a series of near visits to various different attractors. The resulting dynamics constitute a smooth trajectory through state space, as the neuronal population codes gradually transition—not discretely "teleport"—from one recognizable pattern to another. In fact, as each next word in speech and each next object in free viewing tend to arrive at the senses every third of a second (Spivey & Dale, 2004) and the time needed for a neuronal population code to achieve full activation tends to be around half a second (Rolls & Tovee, 1995), the vast majority of

³In fact, even the very first time step of mouse movement revealed a significant difference in angle for typical and atypical animals. Typical animals elicited initial movements that were more directly aimed at the correct response box, suggesting that these mouse-tracking data are not the result of a routinized upward-movement command followed by an early or late subsequent command to turn left or right.

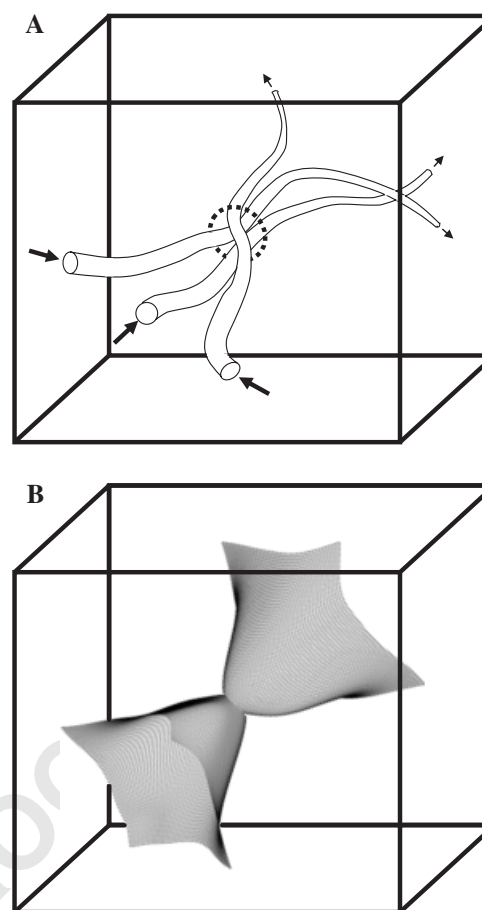


Fig. 2. Three-dimensional representations of mental state space. A given location (i.e., a specific pattern of neuronal activation) may be visited via a wide variety of trajectories (panel A). The dashed circle could indicate the attractor for a word like *eat*; Various instances of hearing this word in different contexts would involve the mental trajectory traveling through this region, with subtle nuances in its use being reflected in subtle spatial variation of the entry point (cf. Elman, 2004). Since different kinds of things perform eating events, the trajectory can come from a variety of previous regions, and since different kinds of things are eaten, the trajectory can then continue to a variety of subsequent regions. Interpolating over the many previous directions and the many subsequent directions that trajectories have taken relative to this location in space, we can envision graded semantic cones (panel B) for the past contexts and for the future contexts that better depict the temporal continuity, and the state-space contiguity, of what a word or concept can mean in its various uses.

the mental trajectory's time is spent between attractors rather than in them.

To provide a visualization of this temporally continuous change in mental activity, Figure 2A offers a cartoon rendition of state-space trajectories that might, on different occasions, visit a common attractor. If the many trajectories that visited this attractor were overlaid on one another (Figure 2B), the envelope of their convergence and divergence might look something like past and future light cones in physics: a past thought cone and a future thought cone, if you will. This kind of spatiotemporal hourglass in the state space of the mind can be treated as an abstracted way to envision semi-bounded "individuated elem-

ents” in cognition, while nonetheless acknowledging that when they actually happen, these mental events are partially overlapping with one another in time and in state space.

More real-time experimental evidence for the continuity of mind is accumulating in psycholinguistics, visual perception, categorization, and even reasoning and problem-solving. The long-term consequences for contemporary theories in psychology are that the putative classifications between different cognitive processes (e.g., word recognition, object recognition, decision making, action) become less important than the gradations that intertwine such classifications and that map the partial overlap among them all. For future work in this area, the continuous state space will eventually need to include among its dimensions not only neuronal parameters but also biomechanical and ecological parameters, such as degrees of freedom for particular limb movements and action possibilities that the surroundings make available to a particular body (e.g., affordances)—thereby treating the larger animal–environment system (not just the brain) as the arena in which cognition emerges (cf. Turvey & Shaw, 1999).

Recommended Reading

- Kelso, J.A.S. (1995). *Dynamic patterns*. Cambridge, MA: MIT Press.
- Port, R., & van Gelder, T. (Eds.) (1995). *Mind as motion*. Cambridge, MA: MIT Press.
- Rogers, T., & McClelland, J. (2004). *Semantic cognition*. Cambridge, MA: MIT Press.
- Spivey, M. (in press). *The continuity of mind*. New York: Oxford University Press.
- Tabor, W., & Tanenhaus, M. (1999). Dynamical modeling of sentence processing. *Cognitive Science*, 23, 491–515.
-

Acknowledgments—We are grateful to Harry Reis and two anonymous reviewers for helpful comments in revising earlier versions of this article. This work was supported by NIMH Grant R01-63961 to MJS and by a Paller-Dallenbach Fellowship to RD.

REFERENCES

- Anderson, J.A., Siverstein, J., Ritz, S., & Jones, R. (1977). Distinctive features, categorical perception, and probability learning: Some

applications of a neural model. *Psychological Review*, 84, 413–451.

- Bowden, E., & Beeman, M. (1998). Getting the right idea: Semantic activation in the right hemisphere may help solve insight problems. *Psychological Science*, 6, 435–440.
- Cisek, P., & Kalaska, J. (2005). Neural correlates of reaching decisions in dorsal premotor cortex. *Neuron*, 45, 801–814.
- Dale, R., Kehoe, C., & Spivey, M. (in press). Graded motor responses in the time course of categorizing atypical exemplars. *Memory and Cognition*.
- Elman, J. (2004). An alternative view of the mental lexicon. *Trends in Cognitive Sciences*, 8, 301–306.
- Elman, J., Bates, E., Johnson, M., Karmiloff-Smith, A., Parisi, D., & Plunkett, K. (1996). *Rethinking innateness*. Cambridge, MA: MIT Press.
- Jung-Beeman, M., Bowden, E., Haberman, J., Frymiare, J., Arambel-Liu, S., Greenblatt, R., Reber, P., & Kounios, J. (2004). Neural activity observed in people solving verbal problems with insight. *PLOS Biology*, 2, 500–510.
- McMurray, B., Tanenhaus, M., Aslin, R., & Spivey, M. (2003). Probabilistic constraint satisfaction at the lexical/phonetic interface: Evidence for gradient effects of within-category VOT on lexical access. *Journal of Psycholinguistic Research*, 32, 77–97.
- McRae, K. (2004). Semantic memory: Some insights from feature-based connectionist attractor networks. In B. Ross (Ed.), *The psychology of learning and motivation* (Vol. 45, pp. 41–86). San Diego, CA: Elsevier.
- Pylyshyn, Z. (1984). *Computation and cognition*. Cambridge, MA: MIT Press.
- Rolls, E., & Tovee, M. (1995). Sparseness of the neuronal representation of stimuli in the primate temporal visual cortex. *Journal of Neurophysiology*, 73, 713–726.
- Spivey, M., & Dale, R. (2004). On the continuity of mind: Toward a dynamical account of cognition. In B. Ross (Ed.), *The psychology of learning and motivation* (Vol. 45, pp. 87–142). San Diego, CA: Elsevier.
- Spivey, M., Grosjean, M., & Knoblich, G. (2005). Continuous attraction toward phonological competitors. *Proceedings of the National Academy of Sciences, U.S.A.*, 102, 10393–10398.
- Tanenhaus, M., Spivey-Knowlton, M., Eberhard, K., & Sedivy, J. (1995). Integration of visual and linguistic information during spoken language comprehension. *Science*, 268, 1632–1634.
- Thelen, E., & Smith, L. (1994). *A dynamic systems approach to the development of cognition and action*. Cambridge, MA: MIT Press.
- Turvey, M., & Shaw, R. (1999). Ecological foundations of cognition: I. Symmetry and specificity of animal–environment systems. *Journal of Consciousness Studies*, 6, 111–123.

