

From *Encyclopedia of Cognitive Science* (2002)
Lynn Nadel (editor), Nature Publishing Group, Macmillan Company,
London, Volume 1, pp. 1027-1032.

The Dynamical Systems Hypothesis in Cognitive Science

Robert F. Port, Indiana University

A. Overview

The dynamical hypothesis in cognition identifies various research paradigms applying the mathematics of dynamical systems to understanding cognitive function. The approach is allied with and partly inspired by research in neural science over the past fifty years for which dynamical equations have been found to provide excellent models for the behavior of single neurons (Hodgkins and Huxley, 1952). It also derives inspiration from work on gross motor activity by the limbs (e.g., Bernstein, 1967, Fel'dman, 1966). In the early 1950s, Ashby made the startling proposal that all of cognition might be accounted for with dynamical system models (1952), but little work directly followed from his speculation due to a lack of appropriate mathematical methods and computational tools to implement practical models. More recently, the connectionist movement (Rumelhart and McClelland, 1986) provided insights and mathematical implementations of perception and learning, for example, that have helped restore interest in dynamical modeling.

The dynamical approach to cognition is also closely related to ideas about the *embodiment* of mind and the environmental *situatedness* of human cognition, since it emphasizes commonalities between behavior in neural and cognitive processes on one hand with physiological and environmental events on the other. The most important commonality is the dimension of *time* shared by all of these domains. This permits real-time coupling between domains, where the dynamic of one system influences the timing of another. Humans often couple many systems together, such as when dancing to music -- where the subject's auditory perception system is coupled with environmental sound, and the gross motor system is coupled to both audition and musical sounds. Because of this commonality between the world, the body and cognition, the method of differential equations is applicable to events at all levels of analysis over a wide range of time scales. This approach directs explicit attention to change over time of relevant system variables.

B. Mathematical Context

The mathematical models employed by dynamical systems research derive from many sources in biology and physics. Of the two schemas to be pointed out here, the first is the neural network idea, partially inspired by the remarkable equations of Hodgkins and Huxley (1952) which account for many known phenomena of neurons in terms of the dynamics of cell membrane. They proposed a set of differential equations for the flow of sodium and potassium ions through the axonal membrane during the passage of an action potential down the axon. These equations, which turn out to apply with slight modification to all neurons, inspired extensions to account for whole cells (rather than just a patch of membrane) in terms of its likelihood to fire given various excitatory and inhibitory inputs. Interesting circuits of neuron-like units were also constructed and

simulated on computer. The Hodgkin-Huxley equations also inspired many psychological models, like those of Grossberg (1982;1986), the connectionist network models (Rumelhart and McClelland, 1986; Hinton, 1986) and models of neural oscillations (Kopell, 1995).

In this general framework, each cell or cell group in a network is hypothesized to follow an equation like:

$$\text{Equation 1: } dA/dt = -\gamma A(t) + \delta[aE(t) - bI(t) + cS(t)] + \textit{bias}$$

indicating that the change in activation (i.e., the likelihood of firing) at time t , dA/dt , depends on the decay, $-\gamma$, of the current value of A plus a term representing inputs from other cells that are either excitatory, $E(t)$, (tending to increase the likelihood of firing), or inhibitory, $-I(t)$ (tending to decrease the likelihood of firing). For some units there may be an external physical stimulus, $S(t)$. A nonlinear function, $\delta(x)$, encourages all-or-none firing behavior and the bias term adjusts the value of the firing threshold. An equation of this general form can describe any neuron. Over the past 50 years, networks of units like these have demonstrated a wide variety of behaviors, including many specific patterns of activity that animal nervous systems exhibit.

A second concrete schema inspiring the dynamical approach to cognition is the classical equation for a simple oscillator like a pendulum. Indeed, it is obvious that arms and feet have many of the properties of pendula. Students of motor control have discovered that pendular motion is a reasonable prototype for many limb motions. A nearly identical system (lacking the complication of arc-shaped motion) is the equation for a system of a mass and spring. In this form:

$$\text{Equation 2: } m(d^2x)dt^2 + d(dx/dt) + k(x-x_0) = 0$$

it specifies simple harmonic motion in terms of the mass, m , times the acceleration, $(d^2x)dt^2$, the damping, d , scaling the velocity, (dx/dt) , and the spring's stiffness, k , proportional to the distance from the neutral position, x_0 , of the mass. Fel'dman (1966) used heavily damped harmonic motion to model a simple reach with the arm. If the neutral position, x_0 (the attractor position when damped) can be externally set to the intended target position (e.g., by adjusting the stiffness in springs representing flexor and extensor muscles), then an infinity of movements from different distances and directions toward the target will result – simply by allowing the neuromuscular system for the arm to settle to its fixed point, x_0 . A number of experimental observations – reaching maximum velocity in the middle of the gesture, higher maximum velocity for longer movements, automatic correction for an external perturbation plus the naturalness and ease of oscillatory motions at various rates – can be accounted for with a model using a mass and a spring with controllable stiffness, rest length and damping.

In the most general terms, a *dynamical system* may be defined as a set of quantitative variables (e.g., distances, activations, rates of change, etc.) that change simultaneously in real time due to influences on each other. These mutual influences can be described by differential or difference equations (van Gelder and Port, 1995). Defined this way, Newton's equations of motion for physical bodies were the earliest dynamical models. Mathematical developments over the past 30 years have revolutionized the field. Whereas up until the 1950s, the analysis of dynamical models was restricted to linear systems (as in Equations 1 and 2) and only when they contain no more than 2 or 3

variables, now, through the use of simulations by discrete computer programs and using computer graphics, practical methods for studying nonlinear systems with many variables are now possible (Strogatz, 1994).

C. Perceptual Models

Dynamical models seem particularly appropriate to account for motor control and for perceptual recognition since research on temporal aspects of perception has been conducted for many years.

One well-known example of a dynamical model for general perception is the *adaptive resonance theory* (ART) model of Grossberg (1995). This neural network is defined by a series of differential equations, similar to the network Equation 1 above, describing how the activation of any given node is increased or decreased by stimulus inputs, excitation and inhibition from other nodes and intrinsic decay. This depends on weights (represented as matrices for a , b and c in Equation 1) which are modified by previous successful perceptual events (simulating learning from experience). The model can discover the low-level features that are most useful for differentiating frequent patterns in its stimulus environment (using *unsupervised learning*) and identify specific high-level patterns even from noisy or incomplete inputs. It can also reassign resources whenever a new significant pattern appears in its environment without forgetting earlier patterns. Notions like "successful perception" and "significant pattern" are provided with mathematical specifications that drive the system toward greater "understanding" of its environment.

To recognize an object such as a letter from visual input, the signal from a spatial retina-like system excites low-level feature nodes. The pattern of activated features here feeds excitation through weighted connections to a higher set of identification nodes. These nodes compete through mutual inhibition to identify the pattern. The best matching unit quickly wins by suppressing all its competitors. When the match is good enough, a "resonance loop" is established between some sensory feature units and a particular classification unit. Only at this point is successful (and, according to Grossberg, conscious) identification achieved. This perceptual model is dynamic because it depends on differential equations that increase or decrease the activation of nodes in the network at various rates. Grossberg's group has shown that variants of this model can account in a fairly natural way for many phenomena of visual perception, including those involving backward masking, reaction time and so on.

D. High-level Models

Dynamical models have also been applied to higher-level cognitive phenomena. First, Grossberg and colleagues have elaborated the ART model with mechanisms like *masking fields* that extend the model to tasks like word recognition from temporally arriving auditory input. Several time-sensitive phenomena of speech perception can be successfully modeled this way (Grossberg, 1986). In a second example, models of human decision making have for many years applied *expected utility theory* where evaluation of the relative advantages and disadvantages of each choice is made at a single point in time. But Townsend and Busemeyer (1995) have been developing their *decision field theory* that not only accounts for the likelihood of each eventual choice, but also accounts for many time-varying aspects of decision making, such as "approach-

avoidance' or vacillatory effects, and the fact that some decisions need more time than others.

It's also important to keep in mind that some phenomena that at first glance seem to depend on high-level reasoning skills may turn out to reflect more low-level properties of cognition. One startling result of this kind is in the "*A-not-B problem*". Infants (9-12 mo.) will sometimes reach to grab a hidden object yet when the object is moved to a new location, infants often reach to the first site – to A not to B. This puzzle was interpreted by Piaget (1954) as demonstrating a lack of the concept of '*object permanence*', that is, that children have an inadequate understanding of objects, thinking that they somehow intrinsically belong to the place they are first observed. Recently Thelen, Schöner, Schrier and Smith (2001) demonstrated a dynamical model for control of reaching that predicted sensitivity to a variety of temporal variables that are supported by experimental tests. The lesson is that sometimes what seems at first to be a property of abstract, high-level, static representations may turn out to result from less abstract time-sensitive processes naturally modeled using dynamical equations.

E. Relation to Situated Cognition and Connectionism

From the perspective of *situated cognition*, the world, the body and the cognitive functions of the brain can all be analyzed using the same conceptual tools. This is important because it greatly simplifies our understanding of the mapping between these domains, and is readily interpreted as an illustration of the biological adaptation of the body and brain to the environment.

Connectionist models are discrete dynamical systems and so are the learning algorithms used with them. But the touchstone of a thoroughly dynamical approach to cognition is the study of phenomena occurring in continuous time – something not all connectionist models do. Of course, neural networks are frequently used to study time-varying phenomena, but other dynamical methods are also available not employing connectionist networks. The development of connectionist modeling since the 1980s has certainly helped to move the field in the direction of dynamical thinking, but connectionist models are not always good illustrations of the dynamical hypothesis of cognition.

F. Contrasting the dynamical systems framework with traditional approaches.

Of course, the most widespread conceptualization of the mechanism of human cognition proposes that cognition resembles computational processes, like deductive reasoning or long division, by using symbolic representations of objects and events in the world that are manipulated by cognitive operations modeling time only as serial order -- not in real time. These operations reorder or replace symbols, and draw deductions from them. The *computational approach* has its best-known articulation in the *physical symbol system hypothesis* (Newell and Simon, 1972). The theoretical framework of modern linguistics (Chomsky, 1963, 1965, 1967) also falls squarely within this tradition since it views sentence generation and interpretation as a serially ordered process of manipulating word-like symbols (e.g., *table* and *go*), abstract syntactic symbols (like *NounPhrase* or *Sentence*) and letter-like symbols representing minimal speech sounds (such as */t/*, */a/* or features like [*Voiceless*] or [*Labial*]) in discrete time. In application to skills like the perceptual recognition of letters and sounds or recognizing a person's distinctive gait, or the motor control that produces actions like reaching, walking or

pronouncing a word, the traditional approach hypothesizes that essentially all processes of cognition are computational operations that manipulate digital representations in discrete time. The mathematics of such systems is based on the algebra of strings and graphs of symbol tokens. Chomsky's work on the foundation of such abstract algebras (Chomsky, 1963) served as the theoretical foundation for computer science as well as modern linguistic theory.

It should be noted that the dynamical systems hypothesis for cognition is in no way incompatible with serially ordered operations on discrete symbols. However proponents of the dynamical systems approach deny that most cognition can be satisfactorily understood in computational terms. They propose that any explanation of human symbolic processing must sooner or later include an account of its implementation in continuous time. The dynamical approach points out the inadequacy of simply assuming that a 'symbol processing mechanism' is somehow available to human cognition, the way a computer happens to be available to a programmer. A fundamental contrast between these frameworks is that the discrete time of computational models is replaced with continuous time for which first and second time derivatives are meaningful at each instant and where critical time points are specified by the environment and/or the body rather than by the needs of a discrete-time device jumping from one clock tick to the next.

G. Strengths and Weaknesses of Dynamical Models

Dynamical modeling offers many important strengths relative to traditional computational cognition. First, the biological plausibility of digital, discrete-time models remains a problem. How and where might there be in the brain, a device that would behave like a computer chip, clicking along performing infallible operations on digital units? The answer often put forward in the past was 'We don't really know how the brain works, anyway, so this hypothesis is as plausible as any other' (Chomsky, 1965 and even 2000). Such an argument does not seem as reasonable today as it did 30 or 40 years ago. Certainly neurophysiological function exhibits many forms of discreteness. But that does not justify simply assuming whatever kind of units and operations would be useful for a digital model of cognition.

Second, temporal data can finally, by this means, be incorporated directly into cognitive models. Phenomena like (a) processing time (e.g., reaction time, recognition time, response time, etc.), (b) temporal structure in motor behaviors (like reaching, speech production, locomotion, dance), and (c) temporal structure in stimulation (e.g., for speech and music perception, interpersonal coordination while watching a tennis match, etc.) can now be linked together if critical, domain-spanning events can be predicted in time.

The language of dynamical systems provides a conceptual vocabulary that permits unification of cognitive processes in the brain with physiological processes in our bodily periphery and with environmental events external to the organism. Unification of processes across these fuzzy and partly artificial boundaries makes possible a truly embodied and situated understanding of human behavior of all types. The discrete-time modeling of traditional approaches is always forced to draw a boundary somewhere to separate the discrete-time, digital aspects of cognition from continuous-time physiology (as articulated in Chomsky's, 1965, distinction of *Competence vs. Performance*).

Third, cognitive development and runtime processing can now be integrated, since learning and perceptuo-motor behavior are governed by similar processes even if on different time scales. Symbolic or computational models were forced to treat learning and development as totally different processes unrelated to motor and perceptual activity.

Finally, trumping the reasons given above, is the fact that *dynamical models include discrete-time, digital models as a special case* whereas the other way around is not possible. (The sampling of continuous events permits discrete simulation of continuous functions, but the simulation itself remains discrete and only models a continuous function up to one half its sampling rate. See Port, Cummins and McAuley, 1995). Thus, any actual digital computer is, in fact, also a dynamical system with real voltage values in continuous time that are discretized by an independent clock. Of course, computer scientists prefer not to look at them as continuous valued dynamical systems (because it is much simpler to exploit their digital properties) but computer engineers have no choice. Hardware engineers have learned to constrain computer dynamics to be governed reliably by powerful attractors for each binary cell that assure that each bit settles into either state *zero* or state *one* before the next clock tick comes round.

These strengths of dynamical modelling are of great importance to our understanding of human and animal cognition. As for weaknesses of dynamical modelling, there are several. First, the mathematics of dynamical models is more inscrutable and less developed than the mathematics of digital systems. It is clearly much more difficult, for the time being, to construct actual models except for carefully constrained simple cases.

Second, during some cognitive phenomena (e.g., a student performing long division, or designing an algorithm, and possibly some processes in the use of language) humans appear to rely on ordered operations on discrete symbols. Although dynamical models are capable of exhibiting digital behavior, how a neurally plausible model could do these tasks remains beyond our grasp for the time being. It seems that computational models are, at the very least, simpler and more direct, even if they remain inherently insufficient.

H. Discrete vs. Continuous Representations

One of the major intuitive strengths of the classical computational approach to cognition has been the seeming clarity of the traditional notion of a cognitive representation. Since cognition is conceived as functioning somewhat like a program in *Lisp*, the representations are constructed from parts that resemble Lisp *atoms* and *s-expressions*. A *representation* is a distinct data structure that happens to have semantic content (with respect to the world outside or inside the cognitive system). They can be moved around or transformed as needed. Of course, such tokens have an undeniable resemblance to words and phrases in natural language (cf. Fodor, 1975). Thus, if one considers making a sandwich from bread and ham in the refrigerator, one can imagine employing cognitive tokens standing for bread, the refrigerator, etc. Thinking about sandwich assembly might be cognitively modeled using representations of sandwich components. Similarly, constructing the past tense of *walk* can be modeled by concatenating the representation of *walk* with the representation of *-ed*. However, this traditional view runs into more difficulties when we try to imagine thinking about actually slicing the bread or spreading the mayonnaise. How could discrete, wordlike representations be deployed to yield successful slicing of bread? But if this is instead to be handled by a nonrepresentational

system (such as a dynamical one), then how could we combine these two distinct and seemingly incompatible types of systems?

The development of *connectionist models* in the 1980s, employing networks of interconnected nodes, provided the first alternative to the view of representations as context-invariant, manipulable tokens. In connectionist models, the result of a process of identification (of, say, an alphabetic character or a human face) is only a temporary pattern of activations across a particular set of nodes, not something resembling a context-free object. The possibility of representation in this more flexible form led to the notion of *distributed representations*, where no apparent 'object' can be found to do the work of representing, but only a particular pattern distributed over the same set of nodes as are used for many other patterns. Connectionists emphasized that such a representation would not seem to be a good candidate for a *symbol* as conventionally conceived in the formalist or computational tradition, yet can still function as a representation for many of the same purposes.

The development of dynamical models of perception and motor tasks has led to further extension of the notion of the representational function to include time-varying trajectories, limit cycles, coupled limit cycles and attractors toward which the system state may tend. From the dynamical viewpoint, static, computational representations will play a far more limited role in cognition. Indeed, a few researchers in this tradition deny that static representations are ever needed for modeling any cognitive behavior (Brooks, 1997).

References

- Ashby, R. (1952) *Design for a Brain*. Chapman-Hall, London.
- Brooks, Rodney (1997) Intelligence without representation. In J. Haugeland (ed.) *Mind Design II*. MITP, Cambridge, Mass. pp. 395-420.
- Chomsky, Noam (1963) Formal properties of grammars. In R. D. Luce, R. R. Bush, & E. Galanter (eds.) *Handbook of Mathematical Psychology, Vol. 2*, Wiley, New York. pp. 323-418.
- Chomsky, Noam (1965) *Aspects of the Theory of Syntax*. MITP, Cambridge, Mass.
- Chomsky, Noam (2000) Linguistics and brain science. In A. Marantz, Y. Miyashita and W. O'Neil (eds.) *Image, Language and Brain*. MITP, Cambridge, Mass, pp. 13-28.
- Fel'dman, A. G. (1966) Functional tuning of the nervous system with control of movement or maintenance of a steady posture-III. Mechanographic analysis of the execution by man of the simplest motor tasks. *Biophysics 11*, 766-775.
- Fodor, Jerry (1975) *The Language of Thought*. Harvard University Press, Cambridge, Mass.
- Grossberg, Stephen (1982) *Studies of mind and brain: Neural principles of learning, perception, development, cognition, and motor control*. Kluwer Academic, Norwell, Mass.

Grossberg, Stephen (1986) The adaptive self-organization of serial order in behavior: Speech, language and motor control. In N. E. Schwab and H. Nusbaum (eds.) *Pattern Recognition by Humans and Machines: Speech Perception*. Academic Press, Orlando, Florida. pp.187-294.

Grossberg, Stephen (1995) Neural dynamics of motion perception, recognition, learning and spatial cognition. In Port and van Gelder (eds) *Mind as Motion: Explorations in the Dynamics of Cognition*. MITP, Cambridge, Mass, pp. 449-490.

Haken, H., Kelso, J. A. S., & Bunz, H. (1985). A theoretical model of phase transitions in human hand movements. *Biological Cybernetics*, 51, 347-356.

Haugeland, John. (1985). *Artificial Intelligence: The Very Idea*. MITP, Cambridge, Mass.

Kopell, Nancy (1995) Chains of coupled oscillators. In M. Arbib (ed.) *Handbook of Brain Theory and Neural Networks*. MITP, Cambridge Mass, pp. 178-183.

Newell, Allen, & Herbert Simon. (1975) Computer science and empirical inquiry. *Communications of the ACM*, pp. 113-126.

Piaget, Jean (1954) *The Construction of Reality in the Child*. Basic Books, New York.

Port, Robert F., Fred Cummins, and J. Devin McAuley. (1995) Naive time, temporal patterns and human audition. In Robert F. Port and Timothy van Gelder, editors, *Mind as Motion: Explorations in the Dynamics of Cognition*. MITP, Cambridge, Mass. pp. 339-371.

Port, Robert & Timothy van Gelder, editors. (1995) *Mind as Motion: Explorations in the Dynamics of Cognition*. MITP, Cambridge Mass.

Strogatz, S. H. (1994). *Nonlinear Dynamics and Chaos with Applications to Physics, Biology, Chemistry, and Engineering*. Reading, Mass., Addison Wesley.

Thelen, Esther, G. Schöner, C. Schrier and L. B. Smith (2001) The dynamics of embodiment: A field theory of infant perseverative reaching. *Behavioral and Brain Sciences*. Cambridge University Press. (in press).

Townsend, James and Jerome Busemeyer (1995) Dynamic representation of decision making. In Robert F. Port and Timothy van Gelder, editors, *Mind as Motion: Explorations in the Dynamics of Cognition*. MITP, Cambridge, Mass. pp. 101-120.

van Gelder, Timothy, and Robert Port. (1995) It's about time: Overview of the dynamical approach to cognition. In Robert Port and Timothy van Gelder, editors, *Mind as motion: Explorations in the dynamics of cognition*. MITP, Cambridge, Mass. pp. 1-43.

FURTHER READING

Books

Abraham, Ralph H. and Christopher D. Shaw (1982) *Dynamics: The Geometry of Behavior*. Vol 1. Ariel Press; Santa Cruz, CA.

Clark, Andy (1997) *Being There: Putting the Brain, Body and World Together Again*. MITP, Cambridge, Mass.

Haugeland, John. (1985). *Artificial Intelligence: The Very Idea*. MIT Press, Cambridge, Mass.

Kelso, J. A. Scott (1995) *Dynamic Patterns: The Self-organization of Brain and Behavior*. MIT Press; Cambridge, Mass.

Port, Robert and Tim van Gelder (1995) *Mind as Motion: Explorations in the Dynamics of Cognition*. MIT Press, Cambridge, Mass.

Thelen, Esther and Linda Smith (1994) *A Dynamical Systems Approach to the Development of Cognition and Action*. MITP; Cambridge Mass.

Articles

Thelen, E., G. Schöner, C. Schrier and L. B. Smith (In press) The dynamics of embodiment: A field theory of infant perseverative reaching. *Behavioral and Brain Sciences* Cambridge Univ Press.

Port, Robert F., Fred Cummins, and J. Devin McAuley. (1995) Naive time, temporal patterns and human audition. In Robert F. Port and Timothy van Gelder, editors, *Mind as Motion: Explorations in the Dynamics of Cognition*. MIT Press, Cambridge, Mass. pp. 339-371.

van Gelder, Tim, & Robert Port. (1995) It's about time: Overview of the dynamical approach to cognition. In Robert Port and Timothy van Gelder (eds.) *Mind as Motion: Explorations in the dynamics of cognition*. MITP, Cambridge, Mass, pp. 1-43.

Glossary

dynamical system. A set of real variables changing over time (real time or discrete) due to mutual influence, described using differential or difference equations.

state space. The set of all possible values of all variables in a dynamical system. For a small number of variables, this is usually displayed graphically as a line, circle, torus or other geometric figure. The dynamical equations for the system constrain possible instantaneous system states and determine possible trajectories through the state space.

attractor. Any location or closed trajectory in the state space of a dynamical system toward which the system tends over time.

settling. The process in which the state of a dynamical system approaches an attractor state -- whether a static attractor or a periodic trajectory.

vector field. A graphic representation of the behavior of a dynamical system specifying for each point in the state space which direction the state will move and at what rate.

coupling. A situation of unidirectional or mutual influence between two or more oscillating dynamical systems, as when a musician plays in time with a metronome or in time with another musician.