

A Complex-Systems Perspective on the “Computation vs. Dynamics” Debate in Cognitive Science

Melanie Mitchell
Santa Fe Institute
1399 Hyde Park Road
Santa Fe, NM 87501 USA
E-mail: mm@santafe.edu

In M. A. Gernsbacher and S. J. Derry (eds.), *Proceedings of the 20th Annual Conference of the Cognitive Science Society—Cogsci98*, 710-715.

Abstract

I review the purported opposition between computational and dynamical approaches in cognitive science. I argue that both computational and dynamical notions will be necessary for a full explanatory account of cognition, and give a perspective on how recent research in complex systems can lead to a much needed rapprochement between computational and dynamical styles of explanation.

The “Computation vs. Dynamics” Debate

Cognition and computation have been deeply linked for at least fifty years, particularly in the symbolic AI tradition. The origin of the electronic digital computer lies in Turing’s attempt to formalize the kinds of symbolic logical manipulations that human mathematicians can perform, and computation was later viewed by Newell, Simon, and others as the correct conceptual framework for understanding thought in general (Newell & Simon, 1976).

Another tradition for understanding thought is rooted in dynamical systems theory. Dynamical approaches to cognition go back at least to the cybernetics era in the 1940s in which information theory, dynamics, and computation were brought together in studying the brain (Ashby, 1952). However, with the dominance of symbolic AI and “information-processing psychology” in the 1960s and 1970s, dynamical-systems-based approaches were not extensively pursued.

More recently, the idea that dynamics is a relevant framework for understanding cognition has become popular again. For example, Thelen and Smith (1994) describe the development of kicking and reaching in infants in terms of dynamical notions such as the stability of attractors in a phase space defined by body and environmental parameters. Movements to new stages in development are explained in terms of bifurcations to new attractors as a

result of change in order parameters—infant weight, body length, etc.—as the infant grows. Thelen and Smith believe that “higher cognition” is ultimately rooted in these types of spatial skills learned in infancy, and thus that higher cognition will itself be best understood dynamically. They contrast their account with traditional “information processing” theories of development, in which new developmental stages are caused by brain maturation and the increasing ability of maturing infants to reason logically.

Many proponents of dynamical approaches in cognitive science, like Thelen and Smith, take the adversarial position that computation and information processing are misleading and incorrect notions to use in understanding cognition. Some go even further and dismiss more basic notions such as representation and “symbols” as being harmful to the cognitive science enterprise (e.g., Freeman & Skarda, 1990).

In a recent overview article, van Gelder and Port (1995) seek to show that the “computational approach”—i.e., that “Cognitive operations are transformations from one static symbol structure to the next” (p. 1)—is false and that the “dynamical hypothesis”—that thought is best understood in the language of dynamical systems theory—is true. In this paper, I briefly review the main points of this debate, argue that both computational and dynamical notions will be necessary for a full explanatory account of cognition, and give my view on how complex systems research can lead to an important rapprochement between computational and dynamical notions in cognitive science.

Although the various dynamical approaches described by van Gelder and Port do not yet yield a concise single formulation, the general idea is that cognition should be characterized as a continual coupling among brain, body, and environment that unfolds in *real time*, as opposed to the discrete time steps of digital computation. The emphasis of the dynamical approach is on how the brain/body/environment system as a whole changes in real time, and dynamics is proposed as the best framework for capturing that change. This is said to contrast with computation’s focus on “internal structure”—i.e., its concern with the static organization of information processing and representational structure in a cognitive system.

This opposition—between dynamics as focused on *change* and computation as focused on internal *structure*—brings to mind a similar debate that has gone on for years in the evolutionary biology community, and whose resolution will, I believe, be instructive for the dynamics/computation debate in cognitive science. The predominant explanatory framework in evolution has been neo-Darwinism, a theory of change *par excellence* (inherited random change from one generation to the next leads to adaptation via natural selection). But some evolutionary theorists have questioned the adequacy of classical neo-Darwinism as either an explanatory or a predictive theory, and argue instead for the primacy of historical contingency (Gould, 1989a) or the self-organization of biological structure not due to natural selection (Fontana & Buss, 1996; Goodwin, 1990; Kauffman, 1993). These “historicists” and “structuralists” are the connectionists of the evolutionary biology community—the people questioning the classical orthodoxy. The selectionist/historicist/structuralist debate has been discussed at length by Gould (1989b), among others. It is becoming increasingly clear,

however, that the stark oppositions posited among these three frameworks are not only false oppositions, but are hindering progress in evolutionary theory. The purely structuralist theories don't explain how structures can be significantly changed in evolution, and the purely selectionist theories don't explain what intrinsic driving forces and constraints there are on the formation of biological structures. What is needed is a theory that incorporates both change and structure¹.

Dynamical Notions

Similarly, in cognitive science, dynamical approaches are theories of change and movement. Although different aspects of dynamical systems theory are emphasized in different dynamical approaches, a common theme is using dynamics as a language for describing continual temporal change in complex systems, something not easily captured in so-called computational approaches.

Dynamical approaches view the behavior of a temporally-changing system in a geometric way—in terms of “trajectories”, “attractors”, “bifurcations”, and so on. Historically, dynamical systems theory has been useful for understanding complex systems in which “self-organization” or “emergent behavior” appears. In many ways dynamics has a natural appeal for cognitive science, since it provides ways to conceptualize systems undergoing continual change, ways to characterize the relative stability of possible patterns of change as a function of system parameters, and ways to think about couplings between complex processes such as the brain, the body, and the environment.

However, there are some limitations to current dynamical approaches that seem difficult to overcome within a pure dynamics framework. Two major limitations, pointed out by Clark (1997, p. 101), are *scaling* and *style of explanation*.

First, current dynamical approaches, if they are to be quantitative, can deal only with low-dimensional analyses (e.g., Beer's analysis of a five-neuron neural network controller for a walking robot; Beer, 1995); it is not clear how the approaches currently being proposed will scale to higher-dimensional systems. Furthermore, it is not clear how the approaches being explored for motor abilities, simple perception, simple language processing, and the like will provide complete accounts of “higher-level” cognitive phenomena such as the recognition of and reasoning about abstract ideas (“representation-hungry” problems, to use Clark's term; Clark, 1997).

Second, while dynamical approaches provide useful high-level descriptions of behavior in

¹This formulation of the evolution debates was given to me by evolutionist Daniel McShea, personal communication. McShea's formulation was elaborated by Crutchfield (1994), who proposes a particular computation-theoretic notion of structure (“computational mechanics of nonlinear processes”) and a related mechanism for the transformation of structure (“hierarchical machine reconstruction”). Crutchfield suggests that a unified theory of these two processes might be termed “evolutionary mechanics”, which he proposes as a general theory of “emergence”.

geometrical terms, in general they don't on their own provide an understanding of how the underlying system gives rise to those aspects of behavior that are *functional* or *adaptive*. For example, we would like to distinguish adaptive from non-adaptive behavior, understand how two adaptive systems with very different dynamical portraits give rise to similar functional behavior, and understand the source of errors made by an adaptive system and how its function will be affected by various sorts of "lesions". We would also like to understand how *new* functional components give rise to improvements in the system. I will argue below that such accounts can be given in a dynamics framework, but only in ones in which functional, information-bearing, and information-processing components can be identified.

Computational Notions

There are many reasons to question computation as framework for understanding cognition. The von Neumann-style architecture that has dominated computer science for most of its history is quite different from the architecture of the brain. The former has centralized control, random access memory, and serial, deterministic processing. The latter consists of myriad relatively simple components with no (known) central control, limited interactions among components, spatial structure, massive parallelism, complex dynamics, and is permeated with noise, giving rise to stochastic processing. In this view, computational processes couldn't be more different than brain processes. As Beer wrote, "the organization that the very terms of the computational language presuppose is nowhere apparent" (Beer, 1995, p. 128).

However, in the rush to rid the cognitive science world of symbolic computational notions, many of the proponents of dynamics have neglected the reasons why computation has been such an attractive framework for cognition for such a long time. First, computational notions have provided us with a new notion of "mechanism". In the history of science, the meaning of "mechanism" has been extended a number of times. For example, in the 17th and 18th centuries, a "model" of a scientific phenomenon was a mechanism described as a combination of the six "basic machines": the lever, the wheel and axle, the pulley, the inclined plane, the wedge, and the screw (Toulmin, 1993). Over time, what counted as a mechanism in science was gradually broadened, and in the 1930s, computation, in the form of Turing Machines, came to be thought of as a new type of mechanism, one that was capable of processing symbols. Further evolution of the notion of "mechanism" and "mechanistic explanation" can be expected (and has already occurred, e.g., in the understanding of the metabolism and self-reproduction of biological cells). The kind of explanation—in terms of function and adaptation—that I claimed above to be necessary requires uncovering mechanisms that explain how function arises and changes in complex systems like the brain and how information is processed. This is something that an extended computation theory—one that is relevant to complex systems—can offer; some steps in this direction will be described below.

Mechanisms explaining functionality are precisely what Marr was getting at in his "rep-

resentation and algorithm” level of description of complex information processing (Marr, 1982). This is the level at which equivalence classes of processes can be described, so that we can understand, for example, how two processes with quite different dynamics can use the same higher-level mechanism (“algorithm”) to accomplish a task and how higher-level structures (“internal representations”) give rise to functionality by carrying information.

In other words, computational theories in cognitive science are theories of structure, making claims about the information processing and functional structure of mental states (e.g., semantic networks, neural networks, schemata, Bayesian belief networks, fuzzy logic, theorem provers). Most of these theories assume that information processing consists of the manipulation of explicit, static symbols rather than the autonomous interaction of emergent, active ones (Hofstadter, 1985). Such theories typically cannot easily explain what driving forces and constraints there are on how the mental states in question can change, what trajectories they can take, their coupling with the body and the environment, and how high-level symbols can emerge from a lower-level substrate.

In short, dynamical approaches contribute a much needed characterization of *continual change* in cognitive systems and a much needed framework for describing complex couplings among brain, body, and environment. Computational approaches contribute notions of mechanism and equivalence classes of mechanisms that shed light on functional and adaptive behavior in complex systems. What we need is a rapprochement between computation and dynamics (between theories of structure and theories of change) that can provide both. Others have made similar points (e.g., Clark, 1997), but without concrete examples. I will argue below that complex systems research is now leading in this direction and can provide concrete examples that will help build our intuitions about how to achieve such a rapprochement.

A Complex Systems Perspective

Attempts at rapprochements between computation and dynamics are coming from many sectors, in particular from research on “complexity,” in which dynamics, computation, and adaptation are beginning to be viewed in a more unified framework. The goal of complex systems research is to explain, across disciplines, how complex and adaptive behavior can arise from systems composed of large numbers of relatively simple components with no central control and with complicated but limited interactions. Dynamics and computation have figured centrally in complex systems research, and a major effort in that field is to understand how sophisticated, functional information processing can arise from decentralized, dynamical substrates, and how that information processing can improve via processes of learning and evolution. This has resulted in considerable work on extending dynamical systems theory, computation theory, and evolution and learning theory to be relevant for such investigations. As noted by Crutchfield (1994), “The interplay between computation, dynamics, and induction emphasizes a trinity of conceptual tools required for studying the emergence of complexity.”

This interplay is exemplified by the work of Hofstadter and his colleagues on “active symbols” and mental fluidity (Hofstadter, 1995); Crutchfield and his colleagues’ work on the “computational mechanics of nonlinear processes” (Crutchfield, 1994; Crutchfield & Hanson, 1993); Moore and his colleagues’ work on understanding what dynamical systems can compute and extending computation theory to continuous-valued computation and two-dimensional languages (Moore, 1990, 1996; Lindgren, Moore, & Nordahl, 1997); Fontana and Buss’s work on self-organization and the development of hierarchies in an “algorithmic chemistry” (Fontana & Buss, 1996); and Crutchfield, Mitchell, Das, and others’ work on the evolution of emergent computation in cellular automata (Crutchfield & Mitchell, 1995; Das, Mitchell, & Crutchfield, 1994; Mitchell, Crutchfield, & Das, 1996). These are only a few examples. Here I will describe this last project as an idealized example of how computation and representation can emerge from a complex dynamical substrate.

Cellular automata (CAs) are spatially extended, discrete dynamical systems that capture some of the attributes of complex systems described above. A CA consists of a large number of simple components (“cells”), each with limited communication to other components and each following a simple transition rule. Like complex systems in nature, the microscopic simple components and rules can give rise to highly complicated and unpredictable macroscopic behavior. (For a recent review on CAs as viewed as dynamical systems and as computers, see Mitchell, in press.)

Our project focused on one-dimensional, binary-state CAs with 7-bit neighborhoods. Such a CA consists of a one-dimensional lattice of N cells, each of which can be in state 0 or 1 at a given time step. Each cell communicates with three neighbors on either side. At each time step, each cell decides, based on its own state and those of its six neighbors, whether to remain in its current state or to change state. Each cell obeys the same *transition rule* which can be expressed as a look-up table giving the action to take for each possible configuration of 7 cells. The look-up table has $2^7 = 128$ entries. At each time step, all the cells update in parallel, with the edges wrapping around.

My colleagues and I carried out a study of how an evolutionary process (modeled by a genetic algorithm) could design this type of cellular automaton to perform sophisticated computations (Crutchfield & Mitchell, 1995; Das, Mitchell, & Crutchfield, 1994; Mitchell, Crutchfield, & Das, 1996). Though this project was not meant to be a cognitive model, the results have turned out to have some relevance, I believe, for the computation/dynamics debate in cognitive science.

In particular, we defined a “density-classification” task for cellular automata that requires the cells to perform collective computation. The task is to decide whether the initial configuration (IC) of states in the lattice contains a majority of 1s or of 0s. If it contains a majority of 1s (0s), the task is to iterate to a fixed-point of all 1s (0s). This task is trivial for a system with central control and random-access memory: all the system has to do is count up the number of 1s in the lattice and divide by the lattice size. But CAs have no central control or globally accessible memory—each cell can only communicate locally with

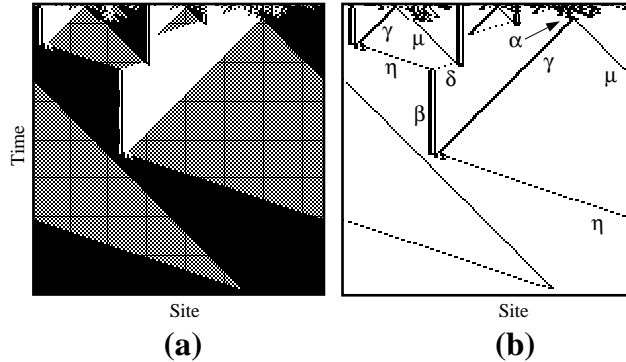


Figure 1: (a) Space-time diagram of ϕ_{100} , one of the best-performing CAs discovered by the GA in our experiments. The 149-cell one-dimensional lattice is arrayed along the horizontal, with time increasing down the page. Cells in state 0 are colored white; cells in state 1 are colored black. (The “greyish” area is a checkerboard pattern of alternating 0s and 1s.) In this diagram, ϕ_{100} starts with an initial configuration with 51% 1s, and by 150 time steps reaches the correct classification of all 1s. (b) The same diagram with the three regular domains filtered out, revealing the particles. (Adapted from Mitchell, Crutchfield, & Das, 1996.)

its neighbors. How can a collection of cells limited to local communication cooperate on a global scale to determine a property of the entire lattice (the initial majority state)²?

We used a genetic algorithm (GA) to evolve cellular automaton transition rules to perform the density-classification task (see Mitchell, Crutchfield, & Das, 1996 for details of the algorithm and the experiments we performed). Figure 1(a) gives a diagram illustrating the space-time behavior of ϕ_{100} , one of the best-performing CAs evolved by the GA. The 149-cell one-dimensional lattice is arrayed along the horizontal, with time increasing down the page. Cells in state 0 are colored white; cells in state 1 are colored black. Here ϕ_{100} starts with an initial configuration with 51% 1s, and by 150 time steps reaches the correct classification of all 1s.

We can estimate how well a given CA performs the density-classification task by testing it on a sample of initial configurations to see how many times it reaches a correct classification. ϕ_{100} reached a correct classification on approximately 80% of the IC samples drawn from a uniform random distribution—these the hardest cases since they almost always have density of 1s very close to 0.5. (ϕ_{100} ’s performance was almost as high as that of the best-known human-designed CA for this task; Das, Mitchell, & Crutchfield, 1994.) However, it is not immediately clear from ϕ_{100} ’s space-time behavior *how* it performs the density-classification task and why it obtains 80% performance.

²As was described in Mitchell, Crutchfield, & Das (1996), the simple “local-majority vote” CA cannot perform the task because it can process information only locally and cannot transmit information about local segments of the initial configuration to different parts of the lattice.

A purely dynamical approach to understanding ϕ_{100} 's behavior, omitting all kinds of “information processing and representation talk”, would consider the system’s time-varying global state to be the 149-dimensional vector encoding the current configuration and a system trajectory to be the sequence of configurations the CA goes through starting with a particular initial configuration. In principle, attractors could be identified and stability properties of those attractors could be determined (such an analysis of a similar CA, proving that the only attractors are the all-0s and all-1s fixed points, was given by Gonzaga de Sá and Maes, 1992). However, such an analysis would miss two essential properties of this CA: first, that with respect to density-classification performance, the important action goes on during the *transient* period leading up to a fixed point, and second, that the transient configurations have *internal structure* that cannot be identified by defining the system’s global states to be 149-bit vectors. Proponents of dynamical approaches would no doubt argue that reduced-dimensional descriptions of the global state could be found, and they are right; I will argue below that useful dimension-reduction in this case requires “information-processing and representation talk”.

Our analysis of ϕ_{100} 's behavior builds on the “computational mechanics of cellular automata” framework of Crutchfield and Hanson (1993), which decomposes CA space-time behavior roughly into “pattern bases” and “particles”. Very briefly, these pattern bases—called “regular domains”—are regions of space-time consisting of strings in the same regular language; in other words, they are regions that are computationally homogeneous and simple to describe. For example, in Figure 1(a), there are three regular domains (black, white and checkerboard), corresponding to the regular languages 0^* , 1^* , and $(01)^*$. Particles are the localized boundaries between those domains. They are revealed in Figure 1(b), in which the three regular domains have been filtered out. For convenience, some of the particles have been labeled with Greek letters.

In computational mechanics, particles are identified as information carriers, and collisions between particles are identified as the loci of information processing. In our case, particles are an information-based method for reducing dimensionality in explaining ϕ_{100} 's behavior: since the three regular domains are simple (where “simple” means “computationally simple” in the sense of simple regular languages), we can deduce that none of those regions alone carries the information needed to globally determine the relative density of 0s and 1s. It is the particles that are doing the important work.

Focusing on the level of particles allows us to understand *how* ϕ_{100} classifies initial densities of 1s. Roughly, over short times, ϕ_{100} maps local high-density regions to all 1s and local low-density regions to all 0s. When an all-1s region on the left meets an all-0s region on the right, a vertical boundary is created and propagated with zero velocity. When an all-0s region on the left meets an all-1s region on the left, a checkerboard region (alternating 1s and 0s) is created and propagated with velocity 1 in opposite directions. When one side of the propagating checkerboard region collides with the black-white boundary, the inner region (e.g., each of the white regions in Figure 1(a)) is cut off and the outer region is allowed to propagate. For example, in Figure 1(a), the large inner white region is smaller than the large

outer black region, and thus the propagating checkerboard pattern reaches the black-white boundary on the white side before it reaches it on the black side; the former is cut off, and the latter is allowed to propagate. In this way many cells collaborate, using local interactions and global geometry, to determine the relative sizes of low- and high-density regions much larger than the neighborhood size. As can be seen in Figure 1(a), this type of collaboration occurs at several spatial scales.

This imprecise description can be made precise and rigorous by phrasing it in terms of particles, particle velocities, and interactions. The microscopic level of behavior— ϕ_{100} 's transition rule on 7-bit neighborhoods and the detailed space-time configurations that result from that rule—gives rise to an emergent macroscopic level that can be thought of as a “particle physics” for this tiny world. In Hordijk, Crutchfield, and Mitchell (1996) we showed how particle-level descriptions on their own can be used to accurately predict the computationally relevant behavior of their corresponding CAs, such as classification performance, mean time to classification, and so on.

This work is relevant for the computation/dynamics debate in cognitive science in that it gives an idealized example of how a non-traditional form of representations and information-processing can emerge from a dynamical substrate. Particles are idealized examples of emergent representations, and particle interactions are idealized examples of emergent information processing. Particles are emergent because they are nowhere explicitly encoded in the microscopic CA rules, and yet have been shown to be of fundamental relevance to a CA's ultimate performance (and thus to its survival in the GA evolution). Particles are representations in that they carry compressed information about the “environment” (here the IC) encoded in their velocities and relative phases. Particle interactions are the loci at which this information is combined and used in decision-making. For example, the α and β particles encode different types of ambiguity in the IC (large black and white regions bordering one another). α decays into γ and μ . γ carries the information that it borders a white region and μ carries the information that it borders a black region. These two particles, by having the same velocity, also carry the mutual information of having come from the same ambiguous region in the IC. When (as in the figure) γ collides with β before μ does, the information contained in β and γ is combined to deduce that the large initial white region was smaller than the large initial black region it bordered. This new information is encoded in a newly created particle η , whose job is to catch up with and annihilate the μ (and itself).

Thus, in this very simple system, particles accomplish, in an idealized way, two of the main things representation needs to accomplish: compressing information about the environment and communicating that information to other parts of the system.

The “particle-logic” story above sounds very computational, but it is certainly of a non-traditional kind. For one thing, representations in the form of particles are not static, passive, or symbolic—they encode information dynamically, actively, and numerically in terms of their velocities and other dynamical attributes. Furthermore, they are not explicitly defined—they emerge from a lower-level dynamical substrate. And finally, there is no central

executive processing the information encoded in particles; their collective dynamics is what effects information processing in the system.

In short, here is a system in which both dynamical and computational notions are necessary for a full account in functional and mechanistic terms. Without the particle level description, we would not understand what makes one CA have higher performance than another, what mistakes are made by a given CA, or how two CAs with quite different microscopic dynamics can implement the same “strategy” for performing the task. Particles provide us with a non-traditional version of Marr’s representation and algorithm level of description; they allow us to discover equivalence classes of mechanisms among CAs with quite different dynamics. They also allow us to understand how innovation in the evolution of these systems takes place (Das, Mitchell, & Crutchfield, 1994).

The purpose of giving this example is to show that for some complex adaptive systems, even idealized ones like evolving cellular automata, a full understanding will require rapprochements between “computation talk” and “dynamics talk”. My claim is that the same will be true of cognitive phenomena. The CA example and the notion of particles and particle interactions are not meant to be a model of a cognitive system; rather, they act as an “intuition pump” (Dennett, 1991) to help us make sense of difficult ideas in a concrete rather than abstract way. I believe that many concrete, and progressively realistic, examples of such systems will be necessary for us to make sense of the terms of the computation/dynamics debate and to effect its eventual resolution.

Acknowledgments

This research was supported by the Santa Fe Institute, NSF grant IRI-9320200, and DOE grant DE-FG03-94ER25231. Thanks to Randy Beer, Andy Clark, Jim Crutchfield, Rajarshi Das, Peter Hrabér, Wim Hordijk, and Erik van Nimwegen for contributions to this work and many discussions about these issues.

1 References

- Ashby, W. R. (1952). *Design for a Brain*. New York: Wiley.
- Beer, R. D. (1995). Computational and dynamical languages for autonomous agents. In Port & van Gelder (1995), 121–147.
- Clark, A. (1997). *Being There: Putting Brain, Body, and World Together Again*. Cambridge, MA: MIT Press.
- Crutchfield, J. P. (1994). Is anything ever new? Considering emergence. In G. Cowan, D. Pines, & D. Meltzer (Eds.), *Complexity: Metaphors, Models, and Reality*, 515–537. Reading, MA: Addison-Wesley. Available at <http://www.santafe.edu/~jpc>.

- Crutchfield, J. P., & Hanson, J. E. (1993). Turbulent pattern bases for cellular automata. *Physica D* 69: 279–301. Available at <http://www.santafe.edu/~jpc>.
- Crutchfield, J. P. & Mitchell, M. (1995). The evolution of emergent computation. *Proceedings of the National Academy of Sciences, USA*, 92 (23): 10742. Available at <http://www.santafe.edu/pro>
- Das, R., Mitchell, M, & Crutchfield, J. P. (1994). A genetic algorithm discovers particle-based computation in cellular automata. In Y. Davidor, H.-P. Schwefel, & R. Männer (editors), *Parallel Problem Solving from Nature—PPSN III*. Berlin: Springer-Verlag. Available at <http://www.santafe.edu/projects/evca>.
- Dennett, D. C. (1991). *Consciousness Explained*. Boston: Little, Brown and Co.
- Fontana, W., & Buss, L. (1996). The barrier of objects: From dynamical systems to bounded organizations. In J. Casti & A. Karlqvist (Eds.), *Boundaries and Barriers: On the Limits to Scientific Knowledge*. Reading, MA: Addison-Wesley.
- Freeman, W. J., & Skarda, C. A. (1990). Representations: Who needs them? In J. L. McGaugh, J. L. Weinberger, & G. Lynch (Eds.), *Brain Organization and Memory: Cells, Systems, and Circuits*, 375–380. New York: Guildford Press.
- Gonzaga de Sá, P., & Maes, C. (1992). The Gacs-Kurdyumov-Levin automaton revisited. *Journal of Statistical Physics*, 67 (3/4), 507–522.
- Goodwin, B. C. (1990). Structuralism in biology. *Sci. Prog.*, 74: 227.
- Gould, S. J. (1989a). *Wonderful Life: The Burgess Shale and the Nature of History*. New York: W. W. Norton.
- Gould, S. J. (1989b). A developmental constraint in *Cerion*, with comments on the definition and interpretation of constraints in evolution. *Evolution*, 43(3): 516–539.
- Hofstadter, D. (1985). Waking up from the Boolean dream, or, subcognition as computation. Chapter 26 in D. Hofstadter, *Metamagical Themas*. New York: Basic Books.
- Hofstadter, D. (1995). *Fluid Concepts and Creative Analogies*. New York: Basic Books.
- Hordijk, W., Crutchfield, J. P., & Mitchell, M. (1996). Embedded particle computation in evolved cellular automata. In *Proceedings of the Conference on Physics and Computation—PhysComp96*, Boston, MA. Available at <http://www.santafe.edu/projects/evca>.
- Kauffman, S. A. (1993). *The Origins of Order: Self-Organization and Selection in Evolution*. New York: Oxford University Press.
- Lindgren, K., Moore, C., & Nordahl, M. (1997). *Complexity of two-dimensional patterns*. (Working Paper 97-03-023.) Santa Fe, NM: Santa Fe Institute.
- Marr, D. (1982). *Vision: A Computational Investigation Into the Human Representation and Processing of Visual Information*. San Francisco: W.H. Freeman.
- Mitchell, M. (In press). Computation in cellular automata: A selected review. To appear

- in H. G. Schuster & T. Gramss (editors), *Nonstandard Computation*. Weinheim: VCH Verlagsgesellschaft. Available at <http://www.santafe.edu/~mm>.
- Mitchell, M., Crutchfield, J. P., & Das, R. (1996). Evolving cellular automata to perform computations: A review of recent work. In *Proceedings of the First International Conference on Evolutionary Computation and its Applications (EvCA '96)*. Moscow, Russia: Russian Academy of Sciences. Available at <http://www.santafe.edu/projects/evca>.
- Moore, C. (1990). Undecidability and unpredictability in dynamical systems *Physical Review Letters*, *64*, 2354–2357.
- Moore, C. (1996). Recursion theory on the reals and continuous-time computation. *Theoretical Computer Science*, *162*, 23–44.
- Newell, A., & Simon, H. (1976). Computer science as empirical enquiry: Symbols and Search. *Communications of the Association for Computing Machinery*, *19*(3), 113–126.
- Thelen, E., & Smith, L. B. (1994). *A Dynamic Systems Approach to the Development of Cognition and Action*. Cambridge, MA: MIT Press.
- Toulmin, S. (1993). From clocks to chaos: Humanizing the mechanistic world-view. In H. Haken, A. Karlqvist, & U. Svedin (Eds.), *The Machine as Metaphor and Tool*, 139–153. Berlin: Springer-Verlag.
- van Gelder, T., & Port, R. F. (1995). It's about time: An overview of the dynamical approach to cognition. In R. F. Port & T. van Gelder (Eds.), *Mind as Motion: Explorations in the Dynamics of Cognition*, 1–43. Cambridge, MA: MIT Press.