

Models of Ontogenetic Development for Autonomous Adaptive Systems

Derek Harter (dharter@memphis.edu)

Department of Mathematical Sciences; University of Memphis
Memphis, TN 38152 USA

Robert Kozma (rkozma@memphis.edu)

Department of Mathematical Sciences; University of Memphis
Memphis, TN 38152 USA

Arthur C. Graesser (a-graesser@memphis.edu)

Department of Psychology; University of Memphis
Memphis, TN 38152 USA

Abstract

Biological organisms display an amazing ability during their ontogenetic development to adaptively develop solutions to the various problems of survival that their environments present to them. Dynamical and embodied models of cognition (Clark, 1997; Edelman & Tononi, 2000; Franklin, 1995; Freeman, 1999a, 1999b; Freeman & Kozma, 2000; Freeman, Kozma, & Werbos, 2000; Hendriks-Jansen, 1996; Kelso, 1995; Kozma & Freeman, 2001; Port & van Gelder, 1995; Skarda & Freeman, 1987; Thelen & Smith, 1994) are beginning to offer new insights into how the numerous, heterogeneous elements of neural structures may self-organize during the development of the organism in order to effectively form adaptive categories and increasingly sophisticated skills, strategies and goals. In this paper we present models of ontogenetic development built on neurologically inspired, bottom-up, dynamic approaches to embodied category formation such as those done by Freeman (1975, 1999b), Freeman and Kozma (2000), Kozma and Freeman (2001), Verschure (1998) and Edelman (1987, 1989). We believe that building on such mechanisms from an embodied dynamical perspective will produce autonomous agents that display greatly increased flexibility in their behavior. Such models will represent a better understanding of how the brains of biological organisms not only form perceptual categories of their environments during development, but also develop effective patterns of behavior through the dynamic self-organization of neurological patterns of activity.

Introduction

Biological organisms develop effective behaviors simply by perceiving and acting upon their environment in real time. Their learning is always guided by their basic needs. Through their experience with the environment, they begin to embody, anticipate and exploit the regularities of their ecological niche in the service of their intrinsic needs. Some models of learning and development for autonomous systems are beginning to display some of these properties (Almássy, Edelman, & Sporns, 1998; Edelman et al., 1992; Freeman & Kozma, 2000; Kozma & Freeman, 2001; Verschure, Kröse, & Pfeifer, 1992; Verschure, Wray, Sporns, Tononi, & Edelman, 1995). These abilities include the formation of embodied, organism significant categories through experience; the development of active searching and anticipation of relevant stimuli; the development of a repertoire of skills, or action loops, for the effective transformation of environmental problems and the exploitation of environmental regularities in the service of intrinsic needs.

In this paper we will present some of the most important properties of dynamical and embodied cognition. We will also discuss the properties of ontogenetic development of

skills, strategies and goals in biological organisms that make it a particularly powerful mechanism of learning. We will look at examples of existing systems that display properties of dynamical and embodied cognition. And finally we discuss our own plans for creating models of the ontogenetic development of behavior in autonomous adaptive systems.

Embodied Cognition

Embodied cognition is an emerging viewpoint in cognitive science that emphasizes many differing aspects from the standard cognitive hypothesis (Clark, 1997; Hendriks-Jansen, 1996; Pfeifer & Scheier, 1998; Varela, Thompson, & Rosch, 1993). In the standard view of cognition, the mind is the product of the manipulation of symbolic representations of the problem in order to produce solutions and generate intelligent behavior (Johnson-Laird, 1988; Newell & Simon, 1972, 1976; Newell, 1990). The environment is perceived and transduced into symbolic representations. These symbols encode the current state of the environment and the problem to be solved. They can be manipulated, independent of the environment, to discover solutions to the problem and produce intelligent behavior for the organism.

In an embodied view of cognition, intelligence in biological organisms does not arise through the static manipulation of amodal symbols and representations. Instead, organisms are seen to be embedded in their environments in fundamental ways. Through their real time experiences with their bodies and environments, they begin to embody the salient aspects of situations in ways that guide future perception and behavior towards improved performance. Experience with their ecological niche develops expectations of the environmental regularities that are of benefit to the intrinsic needs and desires of the organism. The organism actively learns to seek out expected stimuli that are relevant to the desires and needs of the organism at a particular moment.

There are many concepts associated with an embodied perspective of cognition. We will briefly present some of the more important concepts in the next sections.

Embodied Organisms are Complete Organisms

Biological organisms are currently the only examples capable of producing a full range of intelligent, adaptive behavior. Standard views of cognition place no special emphasis on the fact that these natural examples of cognition are **complete** organisms. In the standard view of cognition, it seems plausible that by connecting together many specialized subsystems that solve problems in limited, specialized domains, eventually a

complete intelligence will be produced.

From an embodied perspective, we are not likely to understand natural cognition from such a piecemeal approach to studying and building systems. Instead, we must examine and build complete cognitive systems. In this context, complete refers to systems that are autonomous and adaptive. Autonomous systems are those that have certain intrinsic needs, and that are able to produce behavior that is capable of satisfying those needs consistently over time. Pfeifer (Pfeifer, 1996; Pfeifer & Scheier, 1998) characterizes autonomy as the ability of the organism to maintain its critical, intrinsic values within a zone of viability. This is often referred to as “homeostasis”. Adaptivity refers to organisms that are capable of modifying their behavior so that they can more efficiently maintain their critical parameters in their zones of viability.

Studying complete cognitive systems is important for several reasons. Classical approaches to modeling cognition often tackle toy problems in limited domains. The hope is that the techniques developed can then be scaled up to the full problems of cognition. This approach to studying cognition has failed to produce clear insights into how such methods could eventually be scaled up. Embodied cognition, with its emphasis on complete systems, maintains that the answer is not to start with toy environments. Instead we should begin by studying simple, but complete, organisms, in more realistic environments (Brooks, 1990; Pfeifer & Scheier, 1998). Only complete organisms are capable of developing embodied representations and displaying intentional behavior.

Active, Action-Oriented Representations

Another important difference of embodied and classical perspectives concerns the nature of the representations developed and used by the organism. In a classical perspective, symbols are seen as passive structures that are syntactically manipulated to produce solutions. In an embodied perspective, representations are much more intimately tied to the intrinsic needs of the organism. Clark (1997) calls such structures *action-oriented representations*. Action-oriented representations are not passive representations of the state of the environment as it exists at some time. They are continuously updated from sensory information, and they continuously prescribe possibilities for action. Gibson (1979) has called this the concept of affordances, where the representations *afford* opportunities for action for the organism.

The World Represents Itself

Classical models of cognition often experience an exponential explosion of computational power as the environment increases in complexity. An embodied approach to cognition avoids this problem because it advocates the use of simple, cheap, action-oriented representations. From an embodied perspective, it is better to use cheap and active sensing to inform oneself of the state of the environment, rather than building complex representations of the environment. Brooks (1995) states this principle as “the world is its own best model”. Embodied cognition avoids the use of costly and detailed representations. Cheap, quick, active, specialized sensing of the environment is preferred. Instead of maintaining a complex representation of the state of the environment, we

simply direct specialized sensory apparatus to directly perceive the information required for behavior. This approach helps keep the need for computation from exploding in complex environments.

Emergence of Solutions through Collective Activity

A key concept of embodied cognition is the emergence of solutions from many parallel, distributed activities. In an embodied perspective, intelligence is seen as emerging from the parallel activity of many cooperating and competing processes. As in connectionist models, parallel emergence of solutions provides many benefits to the behavior of the system. Such emergent solutions are robust and resistant to damage; tolerant of noisy, incomplete data; satisfy general goals and yet are variable and context dependent. They are also fast, able to produce solutions easily in real time demanding environments. Unlike most classical connectionist modeling, embodied cognition views recurrent, non-linear interactions as a crucial property in the emergence of solutions.

Developing Within the Environment

The emergence of solutions through many parallel processes is not simply a product of the non-linear interactions of components in the organism’s brain. Intelligent behavior also emerges as the product of the interaction of simple behaviors with a complex environment. Simple, instinctive behaviors are seen as intelligent when they are coupled with local environmental cues (Braitenberg, 1984). Development of action-oriented representations aids in this process. Organisms learn simple actions that, when coupled with appropriate learned stimuli, yield intelligent, purposeful behavior.

Clark (1997) says that embodied minds use extensive external scaffolding. The ecological niche of the organism provides many consistent cues for intelligent behavior. Most intelligent behavior in natural organisms involves the fast recognition and exploitation of such opportunities, not in complex planning and reasoning. Also, most organisms tend to offload complex planning and reasoning tasks onto the environment. They do this by allowing the state of the environment to represent the progression of the problem solving task. One example, given by Rumelhart, McClelland, and The PDP Research Group (1986), is in the behavior of people when multiplying large numbers. Most people can instantly recognize and produce the answer to simple, single digit multiplication problems, of the type $7 \times 7 = 49$. However, when given the task of multiplying large numbers together, say 4356×1897 , they invariably resort to pencil and paper, or even a calculator. People do not compute large chains of complicated reasoning and logic. Instead they offload the representation of the progress of the task onto the environment by maintaining the state of the problem solving task with environmental cues. In this case, people make marks on paper (the environment) to keep track of their problem solving progress, while reducing the problems to those simple ones that they can directly recognize and solve. Embodied cognition sees this type of external scaffolding not as simply useful, but as a prevalent and pervasive method used by cognitive systems to reduce computational complexity and perform problem solving tasks in real time.

Better Imperfect than Late

Biological cognition is exemplified by fast pattern completion. It has evolved to produce behavior in real time. The behavior does not necessarily have to be perfect, so long as it is good enough for the continued survival of the organism (at least until the next crisis occurs). Organisms are continually presented with threats and dangers that must be handled immediately in order to ensure their survival. Such requirements do not favor solutions that take large amounts of time. Natural cognition seems to be built upon a foundation of fast pattern recognition and behavior generation keyed to threats and opportunities for action. The embodied cognitive viewpoint recognizes this fundamental feature of natural cognitive systems. According to Port and van Gelder:

”The cognitive system is not a discrete sequential manipulator of static representational structures; rather, it is a structure of mutually and simultaneously influencing *change*. Its processes do not take place in the arbitrary, discrete time of computer steps; rather, they unfold in the *real* time of ongoing change in the environment, the body, and the nervous system. (Port & van Gelder, 1995, pg. 3)”

The Dynamics of Development

The ontogenetic development of behavior provides a powerful mechanism by which organisms learn to organize effective patterns of behavior for performing the necessary tasks of survival. There are many properties of this type of development. It is fundamentally a self-organizing process, in which the constraints of body and environment guide the system towards discovering certain patterns of behavior. Development of behavior in organisms is not so much a process of finding complex chains of effective behaviors, but in finding salient perceptual cues and effective manipulations that simplify and transform the task environment into problems that are directly recognizable and solvable. Problem solving in natural cognitive systems is more often the application of many transformations until the problem is sufficiently simplified to be directly solved. Clark (1997) calls such phenomena action loops. Kirsh and Maglio (1994) call actions that are primarily performed to transform and simplify the task environment epistemic actions.

Problem solving behavior in biological organisms does not tend to be encoded as static, procedural steps. Instead, organisms develop a wide repertoire of action loops and epistemic actions. Development of behavior takes the form of learning more and better action loops for the effective manipulation and transformation of problems. As an organism's repertoire of action loops grows, they become better able to deal with a wide variety of subtle differences in the problems they need to solve. Their solutions become both robust and efficient with experience in problem solving in the environment.

Development of Embodied Cognition

Thelen and Smith (1994), Thelen (1995) envision the development of behavior in cognitive systems as an ontogenetic landscape of stable and unstable attractors and repellers. As the body of the organism changes, new opportunities for behavior are created and destroyed. Development is seen as

a reduction of the degrees of freedom of the system as useful patterns for solving problems are discovered. As stable solutions to problems develop, these in turn change the ontogenetic landscape, opening up new opportunities for some behaviors, and closing off opportunities for others. Development is the discovery of stable patterns of behavior, given the current constraints of the body and the environment.

Natural cognitive systems display both physical and behavioral development. Physical changes in a maturing organism are continually reshaping the ontogenetic landscape, destabilizing previously stable solutions, and forcing the system into finding new patterns of behavior. Natural cognitive systems also display this flexibility in the development of behavior for problem solving. Sequences of behaviors are not learned so much as behaviors that change the state of the environment and thus cue the next behavior in the sequence.

Self-Organization of Behavior

Theories of the self-organization of patterns in nonequilibrium systems provide new insights into the creativity and flexibility displayed by biological organisms (Kelso, 1995). Many of the desirable properties of development in biological organisms make sense only in view of non-linear dynamics. According to Kelso:

“The thesis here is that the human brain is *fundamentally* a pattern-forming self-organized system governed by nonlinear dynamical laws. Rather than compute, our brain dwells (at least for short times) in metastable states: it is poised on the brink of instability where it can switch flexibly and quickly. By living near criticality, the brain is able to anticipate the future, not simply react to the present. (Kelso, 1995, pg. 26)”

The development of problem solving behavior in biological organisms displays these important properties. Solutions are developed that are flexible, efficient and quick. Such systems are not simply reactive, they learn to anticipate and actively seek out future stimuli.

Bottom Up Neurological Models of Categorization and Action

Some systems have been developed that display properties of dynamic and embodied cognition as discussed above. In this section we present four interesting examples of research that display dynamic, self-organizing category formation and development of behavior. These are all examples of systems that have been built using neurologically inspired, intermediate level neural dynamics.

Distributed Adaptive Control

Distributed Adaptive Control, or DAC (Pfeifer & Verschure, 1992; Pfeifer & Scheier, 1998; Verschure et al., 1992; Verschure & Pfeifer, 1993; Verschure, 1998; Verschure & Voegtlin, 1999) is an example of a model of learning based on large scale neural dynamics. At its heart, DAC is a model of classical conditioning, or the learned association of a response to a conditioned stimuli. In the DAC model, there are three levels of control: reactive, adaptive and reflective control.

The reactive level is prewired in the model, and represents the intrinsic values of the autonomous agent. In the case of DAC, the robot instinctively turns away from things when it bumps into them. This represents the value of avoiding damage from collisions with the environment. In addition to a collision sensor, a special sensor for target acquisition is present. DAC is hardwired to move towards the target when it is detected by the target sensor.

The next level is the adaptive control layer. In this layer representations of the states of long range sensors are slowly associated with events that happen in the reactive control layer. So, for example, the system will learn to avoid collisions by associating the profiles of objects sensed with the long range sensor to collisions and the subsequent activation of avoidance behavior. DAC is also capable of learning and exploiting the regularities of the ecological niche it finds itself in. So, if targets are always found behind openings in walls, DAC is capable of learning this association and begins to search out such openings since they tend to lead to finding the targets in the environment.

The final layer of DAC is the Reflective control layer. At this level sequences of actions are formed and remembered through developing sequential representations. This level represents the addition of long term memory to the basic mechanisms of adaptive learning.

DARWIN

DARWIN (Almásson et al., 1998; Edelman, 1987, 1989; Edelman et al., 1992; Edelman & Tononi, 2000; Sporns, Almásson, & Edelman, 1999; Verschure et al., 1995) is another neurologically inspired model that is capable of learning and developing representations simply by interacting within its environment. At the heart of Edelman's DARWIN systems is the classification couple. In a classification couple, two maps of neuronal groups receive input from separate sensors. The two maps are wired together with many reentrant connections. As a result of reentrant coupling and the change of synaptic strengths, corresponding classification patterns begin to be associated and mutually activate one another in the maps. Thus, for example, the feel (tactile map) and shape (visual map) of an object become functionally correlated through repeated experience with the objects in the environment. The correlated patterns of activity in the maps represent coordinated properties of objects encountered within the environment.

DARWIN III is capable of self-organizing categories of objects that it encounters in its environment, and of learning appropriate behavior patterns. DARWIN is capable of learning to track moving objects in its environment and also of directing its manipulator in a targeted manner in order to manipulate its environment. DARWIN III is also capable of adaptive learning of behavior, like DAC. It learns to associate visual properties of desirable and undesirable objects, to the feel of the object. As it gains experience in the environment, it no longer needs to touch a bad object in order to avoid it. It has formed associations between the visual and tactile maps, and it begins to avoid undesirable objects upon seeing them.

KIII: Mesoscopic Dynamics

The discovery that brain dynamics operate in chaotic domains has profound implications for the study of higher brain func-

tion (Skarda & Freeman, 1987). A chaotic system has the capacity to create novel and unexpected patterns of activity. It can jump instantly from one mode of behavior to another, which manifests the fact that it has a collection of attractors, each with its basin, and that it can move from one to another in an itinerant trajectory. It retains in its pathway across its basins a history, which fades into its past, just as its predictability into its future decreases. Transitions between chaotic states constitute the dynamics that we need to understand how brains perform such remarkable feats as abstraction of the essentials of figures from complex, unknown and unpredictable backgrounds, generalization over examples of recurring objects never twice appearing the same, reliable assignment to classes that lead to appropriate actions, and constant up-dating by learning.

The KIII model (Freeman & Kozma, 2000; Kozma & Freeman, 2001) consists of various sub-units; i.e., the KO, KI, and KII sets. The KO set is a basic processing unit, and its dynamics is described by a 2nd order ordinary differential equation. By coupling a number of excitatory and inhibitory KO sets, KI(e) and KI(i) sets are formed. Interaction of interconnected KI(e) and KI(i) sets forms the KII unit. Examples of KII sets in the olfactory system are the olfactory bulb, anterior olfactory nucleus and prepyriform cortex. Coupling KII sets with feed-forward and feedback connections, one arrives at the KIII system.

KIII shows very good performance in learning input data and it can generalize efficiently in various classification problems. KIII has a high dimensional chaotic attractor in the basal state. It can be destabilized by sensory stimuli and switched to a lower dimensional attractor wing that represents a previously learned memory pattern.

Basic Intentional System: The Limbic System

We consider biological organisms to be behaving intelligently when they act in ways that will enhance their current and future survival. The behavior exhibited by biological organisms is often very creative and flexible. Yet such behavior is always directed towards the satisfaction of the basic needs of the organism. Freeman (1999a, 1999b) describes such behavior as intentional behavior. Intentionality provides a key concept that links the neurodynamics of brains to goal-directed behavior.

One of the primary acts of intentional behavior is in directing sensory observation in expectation of information to guide future actions. Both the formation of expectations and the real time dynamic interaction of the organism with the environment are important principles of intentional behavior. Freeman's view of the mechanisms of intentionality is one of nonlinear dynamic interaction of heterogeneous neural elements on many levels and time scales. The neurodynamic architecture of the brain forms many recurrent loops between brain and brain, brain and body, and organism and environment. But the basic architecture of intentional behavior can be found in the simplest and phylogenetically oldest parts of biological brains: the limbic system.

Conclusion and Future Directions

In this paper we have presented an overview of the dynamical and embodied cognitive hypothesis. We have also given

an overview of some systems that display category formation and developmental learning of the type we are interested in. We have begun work on our own models of the ontogenetic development of behavior in autonomous systems (Harter, Kozma, & Franklin, 2001a, 2001b; Harter, 2001; Kozma, Harter, & Franklin, 2001). Our own models emphasize the development of action-oriented representations that afford opportunities for action-loop like interactions between the agent and the environment. Such models are based upon the formation of embodied categories from chaotic non-linear dynamics.

We begin with bottom-up neurological models that are capable of chaotic non-linear dynamics (Freeman & Kozma, 2000; Kozma & Freeman, 2001). These neurologically inspired models are neither low nor high level simulations of neurological function, but instead capture behavior of the mesoscopic dynamics of brain function (Freeman & Kozma, 2000). These models of neurological function are capable of the dynamic formation of categories. These dynamic categories can be thought of as models of embodied category formation. We are planning to expand such mechanisms to not only form perceptual categories, but develop and display action-loop like skills in the context of the problem domain. Our goals are to see how far such mechanisms can go in developing problem solving behaviors, and to what extent these behaviors mimic those seen in natural cognitive systems.

Eventually we plan to build simplified models of complete limbic systems. We hope that these models will be capable of displaying forms of true intentional behavior in autonomous adaptive systems. Such models should display some of the characteristic flexibility of the problem solving behavior that develops in natural cognitive agents. We are developing agents in cognitively demanding real time task environments. Beginning with some virtual environments, like the game of Tetris (Kirsh & Maglio, 1992, 1994), we are developing bottom-up neurological models that are capable of category formation and the development of behavior in such environments. We hope to eventually move to more complex environments, and real autonomous robots.

Acknowledgments

Portions of this work were funded by NSF grant SBR-9720314.

References

Almásson, N., Edelman, G. M., & Sporns, O. (1998). Behavioral constraints in the development of neuronal properties: A cortical model embedded in a real world device. *Cerebral Cortex*, 8, 346-361.

Braitenberg, V. (1984). *Vehicles: Experiments in synthetic psychology*. Cambridge, MA: The MIT Press.

Brooks, R. A. (1990). Elephants don't play chess. *Robotics and Autonomous Systems*, 6, 3-15.

Brooks, R. A. (1995). Intelligence without reason. In L. Steels & R. Brooks (Eds.), (pp. 25-81). Hillsdale, NJ: Lawrence Erlbaum Associates, Inc.

Clark, A. (1997). *Being there: Putting brain, body, and world together again*. Cambridge, MA: The MIT Press.

Edelman, G. M. (1987). *Neural darwinism: The theory of neuronal group selection*. New York, NY: Basic Books.

Edelman, G. M. (1989). *The remembered present: A biological theory of consciousness*. New York, NY: Basic Books.

Edelman, G. M., Reeke, G. N., Gall, W. E., Tononi, G., Williams, D., & Sporns, O. (1992). Synthetic neural modeling applied to a real-world artifact. *Proceedings of the National Academy of Science*, 89, 7267-7271.

Edelman, G. M., & Tononi, G. (2000). *A universe of consciousness: How matter becomes imagination*. New York, NY: Basic Books.

Franklin, S. P. (1995). *Artificial minds*. Cambridge, MA: The MIT Press.

Freeman, W. J. (1975). *Mass action in the nervous system*. New York, NY: Academic Press.

Freeman, W. J. (1999a). Consciousness, intentionality and causality. In R. Núñez & W. J. Freeman (Eds.), (pp. 143-172). Bowling Green, OH: Imprint Academic.

Freeman, W. J. (1999b). *How brains make up their minds*. London: Weidenfeld & Nicolson.

Freeman, W. J., & Kozma, R. (2000). Local-global interactions and the role of mesoscopic (intermediate-range) elements in brain dynamics. *Behavioral and Brain Sciences*, 23(3), 401.

Freeman, W. J., Kozma, R., & Werbos, P. J. (2000). Bio-complexity: Adaptive behavior in complex stochastic dynamical systems. *BioSystems*.

Gibson, J. J. (1979). *The ecological approach to visual perception*. Houghton Mifflin.

Harter, D. (2001). Ontogenetic development of skills, strategies and goals for autonomously behaving systems. In *Proceedings of the 5th world multi-conference on systemics, cybernetics and informatics (sci 2001)*. Orlando, FL. (Submitted)

Harter, D., Kozma, R., & Franklin, S. P. (2001a). Models of ontogenetic development: The dynamics of learning. In *Proceedings of the 2001 learning workshop*. Snowbird, UT. (Submitted)

Harter, D., Kozma, R., & Franklin, S. P. (2001b). Ontogenetic development of skills, strategies and goals for autonomously behaving systems. In *Proceedings of the fifth international conference on cognitive and neural systems (cns2001)*. Boston, MA. (Submitted)

Hendriks-Jansen, H. (1996). *Catching ourselves in the act: Situated activity, interactive emergence, evolution and human thought*. Cambridge, MA: The MIT Press.

Johnson-Laird, P. N. (1988). *The computer and the mind: An introduction to cognitive science*. Cambridge, MA: Harvard University Press.

- Kelso, J. A. S. (1995). *Dynamic patterns: The self-organization of brain and behavior*. Cambridge, MA: The MIT Press.
- Kirsh, D., & Maglio, P. (1992). Reaction and reflection in tetris. In J. Hendler (Ed.), *Artificial intelligence planning systems: Proceedings of the first annual international conference (aips92)*. San Mateo, CA: Morgan Kaufman.
- Kirsh, D., & Maglio, P. (1994). On distinguishing epistemic from pragmatic action. *Cognitive Science*, 18, 513–549.
- Kozma, R., & Freeman, W. J. (2001). Chaotic resonance - methods and applications for robust classification of noisy and variable patterns. *International Journal of Bifurcation and Chaos*, 11(6).
- Kozma, R., Harter, D., & Franklin, S. P. (2001). Self-organizing ontogenetic development for autonomous adaptive systems (SODAS). In *Proceedings of the international joint conference on neural networks (ijcnn 2001)*. Washington, DC. (Submitted)
- Maes, P., Mataric, M. J., Meyer, J.-A., Pollack, J., & Wilson, S. W. (Eds.). (1996). *From animals to animats 4: Proceedings of the fourth international conference on simulation of adaptive behavior*. Cambridge, MA: The MIT Press.
- Meyer, J.-A., Roitblat, H. L., & Wilson, S. W. (Eds.). (1993). *From animals to animats 2: Proceedings of the second international conference on simulation of adaptive behavior*. Cambridge, MA: The MIT Press.
- Newell, A. (1990). *Unified theories of cognition*. Cambridge, MA: Harvard University Press.
- Newell, A., & Simon, H. A. (1972). *Human problem solving*. Englewood Cliffs, NJ: Prentice-Hall.
- Newell, A., & Simon, H. A. (1976). Computer science as empirical inquiry: Symbols and search. *Communications of the Association for Computing Machinery*, 19, 113-126.
- Núñez, R., & Freeman, W. J. (Eds.). (1999). *Reclaiming cognition: The primacy of action, intention and emotion*. Bowling Green, OH: Imprint Academic.
- Pfeifer, R. (1996). Building "fungus eaters": Design principles of autonomous agents. In P. Maes, M. J. Mataric, J.-A. Meyer, J. Pollack, & S. W. Wilson (Eds.), (pp. 3–12). Cambridge, MA: The MIT Press.
- Pfeifer, R., & Scheier, C. (1998). *Understanding intelligence*. Cambridge, MA: The MIT Press.
- Pfeifer, R., & Verschure, P. F. M. J. (1992). Distributive adaptive control: A paradigm for designing autonomous agents. In F. J. Varela & P. Bourguine (Eds.), (pp. 21–30). Cambridge, MA: The MIT Press.
- Port, R. F., & van Gelder, T. (Eds.). (1995). *Mind as motion: Explorations in the dynamics of cognition*. Cambridge, MA: The MIT Press.
- Rumelhart, D. E., McClelland, J. L., & The PDP Research Group. (1986). *Parallel distributed processing: Explorations in the microstructure of cognition*. Cambridge, MA: MIT Press.
- Skarda, C. A., & Freeman, W. J. (1987). How brains make chaos in order to make sense of the world. *Behavioral and Brain Sciences*, 10, 161–195.
- Sporns, O., Almásy, N., & Edelman, G. M. (1999). Plasticity in value systems and its role in adaptive behavior. *Adaptive Behavior*, 7(3-4).
- Steels, L., & Brooks, R. (Eds.). (1995). *The artificial life route to artificial intelligence: Building embodied, situated agents*. Hillsdale, NJ: Lawrence Erlbaum Associates, Inc.
- Thelen, E. (1995). Time-scale dynamics and the development of an embodied cognition. In R. F. Port & T. van Gelder (Eds.), (pp. 69–100). Cambridge, MA: The MIT Press.
- Thelen, E., & Smith, L. B. (1994). *A dynamic systems approach to the development of cognition and action*. Cambridge, MA: The MIT Press.
- Varela, F. J., & Bourguine, P. (Eds.). (1992). *Toward a practice of autonomous systems: Proceedings of the first european conference on artificial life*. Cambridge, MA: The MIT Press.
- Varela, F. J., Thompson, E., & Rosch, E. (1993). *The embodied mind: Cognitive science and human experience*. Cambridge, MA: The MIT Press.
- Verschure, P. F. M. J. (1998). Distributed adaptive control: Explorations in robotics and the biology of learning. *Informatik/Informatique*, 1, 25–29.
- Verschure, P. F. M. J., Króse, B., & Pfeifer, R. (1992). Distributed adaptive control: The self-organization of behavior. *Robotics and Autonomous Systems*, 9, 181–196.
- Verschure, P. F. M. J., & Pfeifer, R. (1993). Categorization, representations, and the dynamics of system-environment interaction: A case study in autonomous systems. In J.-A. Meyer, H. L. Roitblat, & S. W. Wilson (Eds.), (pp. 210–217). Cambridge, MA: The MIT Press.
- Verschure, P. F. M. J., & Voegtlin, T. (1999). A bottom-up approach towards the acquisition, retention, and expression of sequential representations: Distributed adaptive control III. *Neural Networks*, 11, 1531-1549.
- Verschure, P. F. M. J., Wray, J., Sporns, O., Tononi, G., & Edelman, G. M. (1995). Multilevel analysis of classical conditioning in a behaving real world artifact. *Robotics and Autonomous Systems*, 16, 247–265.