

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Nonconvergent Dynamics and Cognitive Systems

Permalink

<https://escholarship.org/uc/item/1222d7z2>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 28(28)

ISSN

1069-7977

Authors

Harter, Derek
Kozma, Robert

Publication Date

2006

Peer reviewed

Nonconvergent Dynamics and Cognitive Systems

Derek Harter (Derek.Harter@TAMU-Commerce.edu)

Department of Computer Science; Texas A&M University
Commerce, TX 75429 USA

Robert Kozma (rkozma@memphis.edu)

Department of Computer Science; University of Memphis
Memphis, TN 38152 USA

Abstract

The conditions and mechanisms for producing general intelligent action in agents are the focus of intensive study of cognitive science. Many positions have been proposed, including symbolic and connectionist viewpoints. New points of view are beginning to emerge, such as embodied and dynamical cognition, but have not yet been fully solidified into a single comprehensive position. In this article we present one such new viewpoint that emphasizes the importance of nonconvergent dynamics to the production of general intelligent behavior. This approach represents a fourth generation of connectionist thought, and is informed from new results in neuroscience and computational neurodynamics. We formulate the necessary and sufficient conditions for the production of intelligent behavior in this approach to cognition and introduce one such model capable of meeting these conditions.

Introduction

What are the mechanisms by which biological organisms, including human beings, produce general intelligent actions in order to survive, reproduce and thrive in their environment? This, in some form, is the basic questions that lies at the heart of cognitive science. Various ideas have been put forward as possible answers to this question. Metaphors of cognition have been inspired from the advanced technologies of the times, including hydraulic, phone switchboard and computer metaphors [Von Neumann, 1958]. Inspiration has been sought from the realms of formal logic as the means of general intelligent action. Others have looked to abstracted models of how neural tissue functions as possibly holding the key insights to the production of intelligent behavior.

Intelligent behavior really lies on a continuum, from simple tropic behaviors of single celled organisms to tool and language use of human beings. We could imagine defining levels or categories of intelligence, each with an appropriate Turing-like behavioral test that could be used to determine inclusion of a system in a category. Such tests would necessarily not be of language use, but could test things like memory capacity (short, and long term), opportunistic vs. goal orientation and problem solving in unique situations. Such tests would give us the ability to state the necessary capacities and behaviors that need to be present to include a system in a certain level of intelligence.

Being able to define and detect levels of intelligence through outward behavioral characteristics is an impor-

tant piece in the study of intelligent behavior. However, since the rejection of behaviorism and the rise of AI and cognitive science, we have not only been interested in the necessary behavioral characteristics that let us detect intelligence, but also in the types of internal mechanisms and processes that might be necessary and sufficient to produce such observed behaviors. The opening up of the study of intelligence to include internal mechanisms has allowed us to attack the problem both from the outside in, and from the inside out.

Once we begin studying the possible mechanisms of intelligence, it is natural to ask if there is any simplest set or category of mechanism that is both sufficient and necessary for the production of general intelligent behavior. Given that there are different levels of intelligence, are different mechanisms needed to achieve these different levels, or can the same mechanisms be used, only expanded to do more? Previous attempts to define such conditions have focused on things like the ability of formal logic like symbol manipulations to perform tasks we usually think of as intelligent, like playing chess or planning a sequence of tasks to perform a goal [Newell, 1980, Newell, 1990]. Another major movement has focused on the power of simple non-linear processing units to remember, recognize and complete patterns [Werbos, 1974, Rumelhart et al., 1986].

Another way of discovering the constraints of intelligence, besides psychological experimentation, is to directly observe the workings of the only known systems that are capable of general intelligent behavior, biological brains. These types of direct measuring of neurological functioning through such methods as EEG recordings and brain imaging techniques have provided us with valuable further constraints on the possible sufficient and necessary dynamics involved in cognition. Such understanding led directly to early connectionist modeling results, and is leading us even further in new directions.

Biological brains are awash in complex, nonconvergent dynamics. Such complex dynamics have usually been abstracted away in connectionist models, with the assumption that they are not necessary to the production of intelligent behavior. However, new ideas in nonlinear dynamical systems theories, both inside and outside of cognitive science, have begun to understand the possible important roles that aperiodic dynamics, such as chaos, may play in self-organizing systems.

Some researchers in dynamical cognition and neuro-

dynamics have speculated on the possibilities that more complex, chaotic like dynamics may play in the role of adaptive behavior [Skarda and Freeman, 1987, Freeman, 1999, Freeman et al., 2000, Kozma and Freeman, 1999, Kozma and Freeman, 2000, Kozma and Freeman, 2001]. Chaotic dynamics have been observed in the formation of perceptual states of the olfactory sense in rabbits [Skarda and Freeman, 1987]. Skarda and Freeman have speculated that chaos may play a fundamental role in the formation of perceptual meanings. Chaos provides the right blend of stability and flexibility needed by the system. Essentially, Skarda and Freeman believe that the normal background activity of neural systems is a chaotic state. In the perceptual systems, input from the sensors perturbs the neuronal ensembles from the chaotic background, and the result is that the system transitions into a new attractor that represents the meaning of the sensory input, given the context of the state of the organism and its environment. But the normal chaotic background state is not like noise. Noise cannot be easily stopped and started, whereas chaos can essentially switch immediately from one attractor to another. This type of dynamics may be a key property in the flexible production of behavior in biological organisms.

Theories and Conditions of Cognition

Symbolic Systems

Symbolic systems are often equated with the machine metaphor of mind. In this viewpoint of cognition, the brain is seen in some sense as a computer. The physical brain represents the hardware of the system, and the mind represents the software. The machine metaphor is a very attractive position for many reasons. It explains how the mind connects with and controls the body, the old mind-body problem, in a way that does not resort to a form of dualism.

The symbolic approach works well as a model of cognition, and is capable of modeling many impressive examples of intelligent behavior in AI. However, challenges to this viewpoint of cognition have appeared, both as practical criticisms of the performance of such systems and more philosophical challenges to the physical-symbol system hypothesis. For example, many tasks that seem almost effortless for biological brains, such as walking, moving, grasping, etc., have proved much more difficult for symbolic systems to address than more constrained (but seemingly more impressive) domains such as playing chess or diagnosing complicated diseases. However, symbolic systems that use probabilistic declarative structures, and are often referred to as gradient models, seem to have recently shown that such systems can behave more flexibly. For example recent successes at long-range navigation in the DARPA grand challenge contests demonstrate such systems can show impressive levels of adaptation and flexibility. Such probabilistic models are motivated by psychological findings that membership in human categories is often not black and white. People have ideas on the degree to which a certain example belongs in a category, and they have notions of the

prototypical member of a category. Probabilistic models show that thinking of intelligent behavior purely in terms of logical deduction and manipulation of symbols is probably too limiting a viewpoint, and thus brings into question the sufficiency of purely formal logical symbol manipulation. Moreover, probabilistic models are still somewhat deficient, from a neurological standpoint, of bridging the gap between high-level cognitive processes and low level neuronal dynamics.

Connectionist Systems

A connectionist view of cognition provides an alternative theory of mind to the symbolic approach. The connectionist approach to cognition has existed for as long as the symbolic approach. However, symbolic viewpoints of cognition have dominated the field of cognitive science until a resurgence of interest in connectionist models in the mid '80s.

The connectionist approach differs from the symbolic paradigm in almost all major dimensions. Connectionist models offer a subsymbolic paradigm, where representations are built from the changing contributions of processing units that represent features below the normal level of human symbolic features. Connectionist models emphasize parallel processing, while symbolic systems tend to process information in a serial fashion. Connectionist representations are distributed over many units, while cognitivist symbols are static localized structures. Connectionist models offer many attractive features when compared with standard symbolic approaches. They have a level of biological plausibility absent in symbolic models that allows for easier visualization of how brains might process information. Parallel distributed representations are robust, and flexible. They allow for pattern completion and generalization performance comparable to biological organisms. They are capable of adaptive learning. In short, connectionist models are an attractive alternative model of cognition.

The connectionist hypothesis might be stated as: large-scale parallelism of (relatively simple) non-linear processing units doing local processing and producing distributed representations are necessary and sufficient to the production of general intelligent behavior.

First Generation Clark [Clark, 2001] categorizes modern connectionism into three generations. The first-generation of connectionism, that began with the perceptron and the work of the cyberneticists [Rosenblatt, 1958, McCulloch and Pitts, 1943], was revived in the mid '80s with the PDP research groups work (among others) on parallel distributed processing [Rumelhart et al., 1986]. First-generation connectionist systems were typified by a multi-layer architecture (usually composed of two or three layers) with strictly feed-forward connections. Backpropagation learning rules have been especially successful in the proliferation of these models [Werbos, 1974]. Such architectures are very familiar to practitioners of AI and Neural Network research. These connectionist models of cognition are very attractive and important for many reasons. They are biologically plausible models with some of the flexibility of

pattern-recognition and generalization exhibited by biological organisms.

Second Generation Second-generation connectionism began to appear in the early '90s. Second-generation connectionism extends first-generation networks to begin to deal effectively with dynamic spatio-temporal events. First-generation networks displayed no real capacity to deal with time or order in the environment. Second-generation connectionist systems added recurrent connections to the networks in order to expand these capabilities [Elman, 1990, Elman, 1991]. Recurrent connections are connections that connect later layers in the network with earlier layers. So second-generation connectionist networks are no longer strictly feed-forward, they contain recurrent connections. The addition of recurrent connections allows for previous states of the network to affect decisions about the current input. In essence, recurrent connections provide a type of short term memory that allows for the categorization of patterns extended in time across the inputs of the network. This ability to deal with spatio-temporally extended patterns in time is an important addition to the capabilities of connectionist systems.

Third Generation Third-generation connectionism is the most recent extension of the connectionist paradigm. This generation of models is typified by even more complex dynamic and time involving properties. These models use more complex, and biologically inspired architectures, along with various recurrent and hard-coded connections. So, for example, rather than the simple multi-layer structure of first and second generations, third-generation networks may have many areas that represent and reflect architectures and subsystems of biological brains. Because of the increasing emphasis on dynamic and time properties, third-generation connectionism has also been called dynamic connectionism.

Nonlinear Dynamics and Cognitive Systems: The Fourth Generation

Biological brains exhibit aperiodic oscillations with a much more rich dynamical behavior than fixed-point and limit-cycle approximation allows. Early connectionist systems captured some of the flavor of neuronal functioning, but abstracted away much of this rich dynamical behavior in favor of simple fixed-point dynamics [Hopfield, 1982, Grossberg, 1980, Kohonen, 1972, Anderson et al., 1977]. Second and third generation systems recapture some of the more complex dynamics because of recurrent connections and specialized architectures, but many are still parameterized to ultimately settle down to fixed-point attractors. The question of what use, if any, aperiodic dynamics may play in cognition has largely been ignored, or its possible significance unrealized. The exploration of nonconvergent dynamics in cognitive processes may constitute the fourth generation of connectionist thought in its evolution towards capturing more of the dynamics and functioning of biological brains. In this section we will argue that, far from being unnecessary noise of no use in cognition, aperiodic

dynamics are necessary for general intelligent behavior.

Nonconvergent Dynamics for Perception

In their influential paper, Skarda and Freeman [Skarda and Freeman, 1987] argued that chaos, as an emergent property of intrinsically unstable neural masses, is very important to brain dynamics. In experiments carried out on the olfactory system of trained rabbits, Freeman was able to demonstrate the presence of chaotic dynamics in EEG recordings and mathematical models. In these experiments, Freeman and his associates conditioned rabbits to recognize smells, and to respond with particular behaviors for particular smells (e.g. to lick or chew). They performed EEG recordings of the activity in the olfactory bulb, before and after training for the smells.

The EEG recordings revealed that in fact, chaotic dynamics (as shown by the observed strange attractors) represented the normal state when the animal was attentive, in the absence of a stimulus. These patterns underwent a dramatic (nonlinear) transition when a familiar stimulus was presented and the animal displayed recognition of a previously stored memory (through a behavioral response). The pattern of activity changed, very rapidly, in response to the stimulus in both space and time. The new dynamical pattern was much more regular and ordered (very much like a limit cycle, though still chaotic of a low dimensional order). The spatial pattern of this activity represented a well defined structure that was unique for each type of odor that was perceptually significant to the animal (e.g. conditioned to recognize). After recognition, all of the EEG waves are firing in phase, with a common frequency (which Freeman called the carrier wave). The pattern of recognition is encoded in the heights (amplitude modulations) of the individual areas. The amplitude patterns, though regular, are not exact limit cycles and exhibit low dimensional chaos. In other words, different learned stimuli were stored as a spatio-temporal pattern of neural activity, and the strange attractor characteristic of the attention state (before recognition) was replaced by a new, more ordered attractor related to the recognition process. Each (strange) attractor was thus shown to be linked to the behavior the system settles into when it is under the influence of a particular familiar input odorant.

Freeman suggests that "an act of perception consists of an explosive leap of the dynamic system from the basin of one (high dimensional, in the attentive state) chaotic attractor to another (low dimensional state of recognition) [Freeman, 1991]. These results suggest that the brain maintains many chaotic attractors, one for each odorant an animal or human being can discriminate. Freeman and Skarda speculate on many reasons why these chaotic dynamics may be advantageous for perceptual categorization. For one, chaotic activity continually produces novel activity patterns which can provide a source of flexibility in the individual. But since chaos is a ordered state, such flexibility is under control.

As Kelso remarks, inherent fluctuations continuously probe the system, allowing it to feel its stability and pro-

viding opportunities to discover new patterns. Another advantage of chaos is that it allows for very rapid switching between attractors, which random activity is not able to do. Excellent examples of synchronization and desynchronization of motor behavior between coupled individuals/oscillators are given in [Kelso, 1995], where sensory/cognitive coupling provided the modulatory effect that induced the transitions between metastable states.

Sufficient Conditions of Nonconvergent Dynamical Viewpoint

Aperiodic dynamics play a significant role in the organization of perceptual mechanisms in biological organisms. The presence of self-organizing critical states have also been detected in other brain systems. These observations have led to the hypothesis that such dynamics are ubiquitous in brains, and are necessary to the flexible organization of biological behavior. Symbolic systems provide little insight into how they may be connected with an environment and generatively construct knowledge about the world they experience. Looking at symbolic systems as models of biological cognition, they are also silent on why such aperiodic dynamics appear in biological brains. Classical connectionist systems have yet to explore the uses of aperiodic dynamics in memory and action.

These observations of the possible significance of nonconvergent dynamics in brains has led us to speculate on the sufficient conditions they suggest. Specifically:

- Complex, nonconvergent dynamics are sufficient to the production of general intelligent behavior.
- An embodied system with appropriate environmental/sensory coupling and internal structural systems for handling the “what”, “where”, “why” and “how” functions of the agent are sufficient to the production of general intelligent behavior.
- The exploitation of nonconvergent dynamics by and within such an appropriately embodied system are necessary and sufficient for producing general intelligent behavior.

In essence we have proposed two conditions for the production of general intelligent behavior. Aperiodic dynamics characteristic of critical states are necessary for the flexible self-organization of memory and behavior. The dynamics of the brain are strongly coupled with their environment. The interaction of brain dynamics with the environmental system produces behavior. We will explore these issues further in the next section, where we describe one such model of cognition.

Hippocampal Simulation

Experimental Architecture

In this section we give an example application of nonconvergent dynamics using Freeman’s K-sets in order to simulate the formation of cognitive maps in the hippocampus using aperiodic attractors. Using an autonomous




Figure 1: Agent morphology (bottom left) and environmental setup for hippocampal simulations.

agent, we demonstrate the formation of cognitive maps as the agent explores the environment.

In this experiment, we used the Khepera virtual environment simulator [Michel, 1996]. Figure 1 (bottom left) shows the morphology of the Khepera agent. The Khepera robot is a simple agent that contains 8 infra-red and 8 light sensors. It has two independently controlled wheels that allow it to move forward, backward, and turn left and right in place. The environment for this experiment is shown in figure 1. In the environment we place 8 light sources, which will be used as salient environmental locations (i.e. they can be thought of as good food sources for the agent in the environment). The light sources are detectable to the agent at a distance, and the range where the food source is detectable is indicated in Figure 1. In addition to the 8 salient environmental locations, there are 4 landmarks. The landmarks are always detectable to the agent, and it knows the distance and direction to each of the 4 landmarks as part of its sensory information.

The architecture of the simulated hippocampus is shown in Figure 2. The portions of the architecture that form the cognitive map of the environment are simulated by a KA-III [Harter and Kozma, 2004, Harter and Kozma, 2005]. These are the CA1, CA2 and CA3 areas, and are based on biological evidence of the structure of the biological hippocampus. Each of the CA areas contains an 8x8 array of oscillatory units (for a total of 64 units in each CA region). Each CA area is connected to the other 2. The interconnection of these 3 CA regions via inhibitory and excitatory feedback forms a KA-III unit. The connections between CA regions will be changed via Hebbian modification.

Orientation beacons are fed into the hippocampal simulation through the DG region (Figure 2, left). The DG again contains an 8x8 matrix of KA-II units. Orientation signals from the 4 landmarks are fed into the DG units. Each of the 4 landmarks has 8 units associated with the




Figure 2: Architecture of the computational model of hippocampal simulations

direction to the landmark, and 8 units associated with the distance. Directions are broken into 8 cardinal units, North, NorthEast, East, SouthEast, South, SouthWest, West and NorthWest. Units are sensitive to the direction of a particular landmark, though we use a graded response with a normal distribution, instead of a simple 1 unit is active and the others being inactive []. Similarly there are 8 cardinal distance values VeryClose, Close, MediumClose, Medium, MediumFar, Far, VeryFar, Distant. Again a graded response with normal distribution is applied to the units. The DG area connects with the CA3 area, and the connections between these areas are also subject to Hebbian modification.

Method

We use two types of learning in the simulation, Hebbian modification and habituation [Kozma and Freeman, 2001]. Hebbian modification only occurs when the robot is within a certain range of a light source. Therefore the light sources provide a certain valence signal that acts as a stimulus to learn environmentally salient locations. When the robot is not within proximity to a light source, no reinforcement signal is produced. During these times habituation of the stimulus occurs. This has the effect of lessening the response of the simulated hippocampus to unimportant regions in the environment.

The expected effect of this stimulation is to form 2 distinct types of dynamical patterns in the CA regions. When the agent is out of range of an environmentally salient location, the dynamics should be in the high-dimensional chaotic state, receptive to input but not indicative of recognizing a salient event. When in range of a light source, the system should transition to a low dimensional attractor, indicative of recognition of the important location. Further, the spatial amplitude modulation patterns in the CA regions upon such recognition should form 8 unique patterns, one for each of the recognized regions.

The agent is allowed to roam in the environment, using a low level mechanisms to produce efficient, but random wandering. The agent roams for some time, 10,000 time

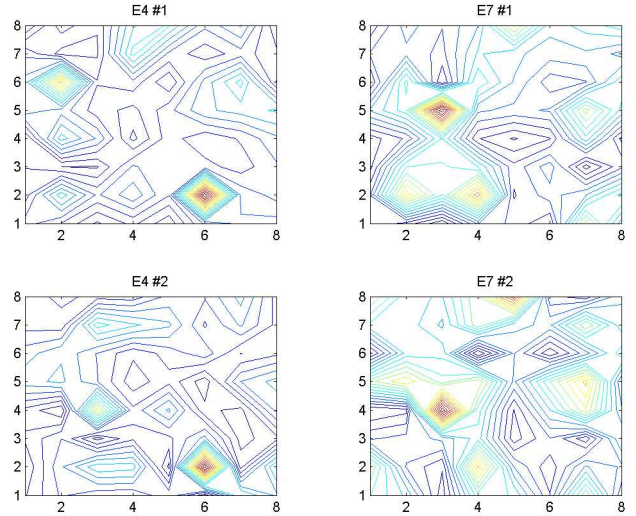


Figure 3: Example of AM Pattern formed in the CA3 hippocampal region. In this figure we show a pattern from two different locations within an environmentally salient region (Top and Bottom). We show AM patterns from environment regions E4 and E7. Similar AM patterns are organized and exhibited when the agent is in the same environmental region.

steps in our simulations. In our simulation 10 time steps approximates 1 second of real world running time, therefore the totaled simulated time of an experiment is 1000 seconds.

Results

Here we examine the amplitude modulation (AM) patterns produced by the hippocampal simulation. Figure 3 shows examples of the AM patterns formed in the CA3 hippocampal matrix for 2 different locations within environmental regions 2, 4, 6 and 8 respectively. The AM patterns shown are from the CA3 hippocampal region. This region has 8x8 units, for a total of 64 time series. We measure the standard deviation of each of the 64 units for a 50ms time window, and plot the results as an 8x8 contour map of the deviations of each of the units in the area. The AM pattern contour plots, therefore, give you an idea of which units are more highly stimulated (higher amplitudes in their activity) and which are less so. As Figure 3 shows, the AM patterns are more similar to those produced from locations within the same environmental region.

As a more complete test of the formation of unique AM patterns, we feed the robot with input from randomly selected locations, within the environmental food areas. AM patterns were collected for the randomly selected regions and compared to one another by calculating the euclidean distance between each pattern. This testing showed that, in fact, the patterns produced within a region are consistently more similar to one another, than those produced in another environmental region.

Conclusion

The hippocampal simulation described here forms distinct AM patterns for the 8 salient environmental regions. These patterns are aperiodic spatio-temporal activity in the CA regions. The characteristic activity peaks in the AM patterns are examples of so called 'place cell' formation. Here we see high activity among certain regions correlated with being in a particular environmental location. For example, looking at the AM pattern for location 7 (Figure 3, right) you notice 3 peaks of activity among the units in the region. It is possible to interpret these peaks as being correlated with environmental locations, and therefore typical examples of the place cell.

The self-organization of spatio-temporal patterns in nonlinear systems are essential to cognitive mechanisms in biological brains. We need to better understand how such mechanisms operate in order to build better models of cognition and smarter autonomous agents. This paper has demonstrated one such self-organizational mechanism for the creation of AM patterns in a cognitive map of an agents environment.

Acknowledgments

This work was supported by NASA Intelligent Systems Research Grant NCC-2-1244.

References

- [Anderson et al., 1977] Anderson, J. A., Silverstein, J. W., Ritz, S. A., and Jones, R. S. (1977). Distinctive features, categorical perception, and probability learning: Some applications of a neural model. *Psychological Review*, 84:413–451.
- [Clark, 2001] Clark, A. (2001). *Mindware: An Introduction to the Philosophy of Cognitive Science*. Oxford University Press, Oxford, NY.
- [Elman, 1990] Elman, J. L. (1990). Finding structure in time. *Cognitive Science*, 14:179–211.
- [Elman, 1991] Elman, J. L. (1991). Distributed representations, simple recurrent networks, and grammatical structure. *Machine Learning*, 7:195–225.
- [Freeman, 1991] Freeman, W. J. (1991). The physiology of perception. *Scientific American*, 264(2):78–85.
- [Freeman, 1999] Freeman, W. J. (1999). *How Brains Make Up Their Minds*. Weidenfeld & Nicolson, London.
- [Freeman et al., 2000] Freeman, W. J., Kozma, R., and Werbos, P. J. (2000). Biocomplexity: Adaptive behavior in complex stochastic dynamical systems. *BioSystems*, 59:109–123.
- [Grossberg, 1980] Grossberg, S. (1980). How does a brain build a cognitive code? *Psychological Review*, 87:1–51.
- [Harter and Kozma, 2004] Harter, D. and Kozma, R. (2004). Aperiodic dynamics and the self-organization of cognitive maps in autonomous agents. In *Proceedings of 17th International Florida Artificial Intelligence Research Society Conference (FLAIRS)*, pages 424–429, Miami Beach, FL.
- [Harter and Kozma, 2005] Harter, D. and Kozma, R. (2005). Chaotic neurodynamics for autonomous agents. *IEEE Transactions on Neural Networks*. in press.
- [Hopfield, 1982] Hopfield, J. J. (1982). Neuronal networks and physical systems with emergent collective computational abilities. *Proceedings of the National Academy of Science*, 81:3058–3092.
- [Kelso, 1995] Kelso, J. A. S. (1995). *Dynamic Patterns: The Self-organization of Brain and Behavior*. The MIT Press, Cambridge, MA.
- [Kohonen, 1972] Kohonen, T. (1972). Correlation matrix memories. *IEEE Transaction on Computers*, C-21:353–359.
- [Kozma and Freeman, 1999] Kozma, R. and Freeman, W. J. (1999). A possible mechanism for intermittent oscillations in the KIII model of dynamic memories - the case study of olfaction. In *Proceedings IJCNN 1999*, pages 52–57.
- [Kozma and Freeman, 2000] Kozma, R. and Freeman, W. J. (2000). Encoding and recall of noisy data as chaotic spatio-temporal memory patterns in the style of the brains. In *Proceedings of the IEEE/INNS/ENNS International Joint Conference on Neural Networks (IJCNN'00)*, pages 5033–5038, Como, Italy.
- [Kozma and Freeman, 2001] Kozma, R. and 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):1607–1629.
- [McCulloch and Pitts, 1943] McCulloch, W. S. and Pitts, W. (1943). A logical calculus of the ideas immanent in nervous activity. *Bulletin of Mathematical Biophysics*, 5:115–133.
- [Michel, 1996] Michel, O. (1996). Khepera simulator package version 2.0. Downloaded from WWW at <http://wwwi3s.unice.fr/om/khep-sim.html>. Freeware mobile robot simulator written at the University of Nice Sophia-Antipolis.
- [Newell, 1980] Newell, A. (1980). Physical symbol systems. *Cognitive Science*, 4:135–183.
- [Newell, 1990] Newell, A. (1990). *Unified Theories of Cognition*. Harvard University Press, Cambridge, MA.
- [Rosenblatt, 1958] Rosenblatt, F. (1958). The perceptron: A probabilistic model for information storage and organization in the brain. *Psychological Review*, 65:386–408.
- [Rumelhart et al., 1986] Rumelhart, D. E., McClelland, J. L., and The PDP Research Group (1986). *Parallel Distributed Processing: Explorations in the Microstructure of Cognition*. MIT Press, Cambridge, MA.
- [Skarda and Freeman, 1987] Skarda, C. A. and Freeman, W. J. (1987). How brains make chaos in order to make sense of the world. *Behavioral and Brain Sciences*, 10:161–195.
- [Von Neumann, 1958] Von Neumann, J. (1958). *The Computer and the Brain*. Yale University Press, New Haven, CT.
- [Werbos, 1974] Werbos, P. J. (1974). *Beyond Regression: New Tools for Prediction and Analysis in the Behavioral Sciences*. Ph.D. thesis, Harvard, Cambridge, MA.