

LA-UR-

98-1921

CONF-980940--

Los Alamos National Laboratory is operated by the University of California for the United States Department of Energy under contract W-7405-ENG-36

TITLE: SYNTACTIC AUTONOMY

RECEIVED  
OCT 05 1998  
OSTI

AUTHOR(S): Luis Rocha

SUBMITTED TO: Proceedings 1998 Conference on Intelligent Systems and Semiotics  
Gaithersburg, MD  
Sept. 1998

DISTRIBUTION OF THIS DOCUMENT IS UNLIMITED

MASTER

By acceptance of this article, the publisher recognizes that the U.S. Government retains a nonexclusive royalty-free license to publish or reproduce the published form of this contribution or to allow others to do so, for U.S. Government purposes.

The Los Alamos National Laboratory requests that the publisher identify this article as work performed under the auspices of the U.S. Department of Energy.

Los Alamos

Los Alamos National Laboratory  
Los Alamos New Mexico 87545

## **DISCLAIMER**

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

## **DISCLAIMER**

**Portions of this document may be illegible  
in electronic image products. Images are  
produced from the best available original  
document.**

# Syntactic Autonomy<sup>1</sup>

Luis Mateus Rocha  
Los Alamos National Laboratory, MS P990

e-mail: [rocha@lanl.gov](mailto:rocha@lanl.gov) or [rocha@santafe.edu](mailto:rocha@santafe.edu)  
URL: <http://www.c3.lanl.gov/~rocha>

## Abstract.

The study of adapting and evolving autonomous agents should be based on a complex systems-theoretic framework which requires both self-organizing and symbolic dimensions. An inclusive framework based on the notions of semiotics and situated action is advanced to build models capable of representing, as well as evolving in their environments. Such undertaking is pursued by discussing the ways in which symbol and self-organization are irreducibly intertwined in evolutionary systems. With this semiotic view of self-organization and symbols, we re-think the notion of autonomy of evolving systems, and show that evolutionary systems are characterized by a particular type of syntactic autonomy. Recent developments in emergent computation in cellular automata are discussed as examples of the emergence of syntactic autonomy in computational environments. New experiments emphasizing this syntactic autonomy in cellular automata are presented.

**Keywords:** Evolutionary Systems, Self-Organization, Autonomy, Artificial Life, Semiotics, Emergence, Cellular Automata, Genetic Algorithms, Situated Action.

## 1. SITUATED SEMIOSIS

### 1.1 Self-Organization

Self-organization is seen as the process by which systems of many components tend to reach a particular state, a set of cycling states, or a small volume of their state space (attractor basins), with no external interference. This attractor behavior is often recognized at a different level of observation as the spontaneous formation of well organized structures, patterns, or behaviors, from random initial conditions (emergent behavior). The systems used to study this behavior computationally are referred to as dynamical systems or state-determined systems, since their

current state depends only on their previous state. They possess a large number of components or variables, and thus high-dimensional state spaces.

Computational self-organization is often used to model physical matter with systems such as boolean networks or cellular automata. The state-determined transition rules are interpreted as the laws of some physical or chemical system [Rocha and Joslyn, 1998]. It follows from the observed attractor behavior that there is a propensity for matter to self-organize (e.g. the work of Prigogine [1985] and Kauffman [1993]). In this sense, matter is described by the laws of physics and the emergent characteristics of self-organization. In the following, whenever the words matter and materiality are used, they should be understood as reflecting this notion of self-organization both in physical and computational environments.

### 1.2 Semantic Emergence

Self-organizing attractor values can be used to refer to observables accessible to the self-organizing system in its environment, and thus perform environmental *classifications* (e.g. classifying neural networks). The process of obtaining novel classifications of an environment by a self-organizing system, which can only be achieved by structural changes to its attractor landscape (e.g. weight changes in a neural network), can be referred to generally as *emergent classification* [for more details on this argument please refer to Rocha, 1996, 1997]. Emergent because it is the result of the local interaction of the basic components of the self-organizing system and not from a global controller.

---

<sup>1</sup>In: Proceedings of the Joint Conference on the Science and Technology of Intelligent Systems (ISIC/CIRA/ISAS 98). National Institute of Standards and Technology, Gaithersburg, MD.

There are three levels that need to be addressed when dealing with the notion of emergent phenomena in self-organizing systems, in particular, of emergent classification. First, there is the material, dynamical, substrate (physical law or computational state-determinacy) which will be the causal basis for all other levels that we may further distinguish. Second, we have the attractor behavior of this dynamics. Finally, we have the (possible) utilization of the set of attractors as referents for some aspects of the interaction of the dynamical system itself with its environment (e.g. the pattern recognition abilities of neural networks). No physical or formal description of the dynamical system and its attractors alone can completely explain this “standing-for”, or semantic, dimension [Pattee, 1995a]. This argument is pursued in more detail in [Rocha 1996, 1997, 1998a].

### 1.3 Pragmatics: Selected Self-Organization and Situated Semantics

For a dynamic system to observe genuine emergence of new classifications, that is, to be able to accumulate useful variations, it must change its structure (that is, its components characteristics establishing a particular attractor landscape). One way or another, this structural change leading to efficient classification (not just random change), has only been achieved through some external influence on the self-organizing system. Artificial neural networks discriminate by changing the structure of their connections through an external learning procedure. Evolutionary strategies rely on internal random variation which must ultimately be externally selected. In other words, the self-organizing system must be structurally coupled [Maturana and Varela, 1987] to some external system which acts on structural changes of the first and induces some form of explicit or implicit selection of its dynamic representations: *selected self-organization* [Rocha, 1996, 1997, 1998a].

Now, for selection to occur we must have some internal vehicle for classification — there must be different alternatives. The attractor landscape of self-organizing systems offers these alternatives. One way of conceptualizing this, is to think of the attractor landscape as a distributed [Van Gelder, 1992] memory bank, where each attractor basin is seen as storing a given classification function. Therefore, semantic emergence in self-organizing systems depends on the existence of distributed memory [Rocha, 1996].

Selected self-organization explicitly emphasizes a second dimension of the semiosis of self-organizing systems in situation with their environments. If classification implies

semantic emergence, selection implies pragmatic environmental influence. In fact, these two dimensions of semiosis cannot be separated; the meaning of the classifications of a self-organizing system does not make sense until it is grounded in the feedback from the repercussions it triggers in its environment. The structural coupling, or situation, of a classifying, self-organizing, agent in its environment is the source of meaning. Indeed, selection does not act on memory tokens internal to a classifying system but on the repercussions those trigger in an environment. Situated Semantics is pragmatic. In this sense, meaning is not private to the agent but can only be understood in the context of the agent’s situation in an environment with its specific selective pressures.

### 1.4 Von Neumann and the Syntactic Advantage

Von Neumann’s [1966] model of self-replication is a systems-theoretic criteria for open-ended evolution [for a detailed discussion of this model see Rocha, 1996, 1997, 1998a]. Based on the notion of universal construction and description it provides a threshold of complexity after which systems that observe it can for ever more increase in complexity (open-ended evolution). However, unlike the situated semiosis of self-organizing systems described in 1.3, this model clearly does not rely on a distributed but on a local kind of memory. Descriptions entail a symbol system on which construction commands are cast. These descriptions are not distributed over patterns of activation of the components of a self-organizing system, but instead localized on “inert” structures which can be used at any time — a sort of random access memory.

By “inert” structures, I mean components with many dynamically equivalent states which can be used to set up an arbitrary semantic relation with the environment. For instance, in the genetic system (which Von Neumann’s model conceptually describes), most any sequence of nucleotides is equally possible, and its informational value (genetic information) is largely independent of the particular dynamic behavior of the DNA or RNA sequence. Genetic information is not expressed by the dynamics of nucleotide sequences, but is instead mediated through an arbitrary coding relation that translates such sequences into amino-acid sequences whose dynamic characteristics ultimately express genetic information into some environment. It is precisely the dynamic irrelevance of nucleotide sequences (“inertness”) that makes DNA and RNA ideal candidates for localized carriers of genetic information (descriptions) given an arbitrary genetic code [Pattee, 1995a; Umerez, 1995].

Von Neumann showed that there is an advantage of local memory over purely dynamic, or distributed, memory in self-replication because if we do not have symbolic descriptions directing self-replication, then an organism must replicate through self-inspection of its parts. Clearly, as systems grow in complexity, self-inspection becomes more and more difficult [Pattee, 1995a]. The existence of a language, a symbol system, allows a much more sophisticated form of communication. Functional, dynamic structures do not need to replicate themselves, they are simply constructed from non-functional (dynamically inert) descriptions. For instance, for an enzyme to replicate itself, it would need to have this intrinsic property of self-replication "by default", or it would have to be able to assemble itself from a pool of existing parts. But for this, it would have to "unfold" so that its internal portions could be reconstituted for the copy to be produced [Pattee, 1995a]. With the genetic code, however, none of these complicated gimmicks are necessary: functional molecules can be simply folded from inert messages. This method is by far more general since any functional molecule can be produced from a description, not merely those that either happen to be able to self-reproduce, or those that can unfold and fold at will to be reproduced from available parts.

The genetic symbol system, with its utilization of inert structures, opens up a whole new universe of functionality which is not available for purely dynamical self-replication. In this sense, it can evolve functions in an open-ended fashion. It also introduces the third level of a semiosis of classifying systems in situation with their environments: syntax – as defined by a construction code. Arguments for the idea of language as a provider of such an enabling syntax for cognitive systems have been pursued elsewhere [Henry and Rocha, 1996; Rocha, 1997, 1998b].

### 1.5 Why do we need syntax?

It can always be argued that the random access memory the genetic system establishes, is nothing but complicated dynamics, and the syntactic dimension is just the result of our subjective observation. But similar arguments can always be pursued to discourage any kind of emergence. Indeed, the notion of self-organization also requires an emergentist argument as pursued in sections 1.1 and 1.2. The dynamic/self-organizing level results from the necessity of complementary modes of description to describe our (ultimately subjective) observation. So why stop there? The genetic dimension has established a new hierarchical level in evolutionary systems which allows a greater level of control of the purely self-organizing bio-

chemical dynamics. Failing to recognize this emergent symbolic level, would prevent the distinction between self-organizing systems such as autocatalytic networks [Kauffman, 1993], from living systems whose replication by genetic memory is much more efficient than template-based replication.

In evolutionary systems this is at the core of the feud between those who claim that natural selection is the sole explanation for evolution and those who stress that other aspects of evolutionary systems, such as developmental constraints, also play an important role. It is no wonder then that the first group stresses the symbolic description, the gene, as the sole driving force of evolution [Dawkins 1976, Dennett, 1995], while the second group likes to think of the propensities of matter or historical contingencies as being of at least equal importance in evolution [Gould, 1989, Salthe 1985, 1993, Kauffman 1993]. In pragmatic terms, however, most evolutionary theorists, one way or another, acknowledge that all these factors play important roles [Eldridge, 1995].

Since all of these aspects of evolutionary systems co-exist, we need inclusive theories and models that incorporate both symbolic and dynamic characteristics [Pattee, 1995b; Rocha, 1996, 1997; Mitchell, 1998]. Classifying systems exist that are purely dynamic; they observe the selected self-organization with distributed memory discussed in 1.3 that is capable of semantic emergence in a selective environment (pragmatics). But the introduction of the syntactic level as prescribed by Von Neumann defines a richer (open-ended) classifying function available to systems capable of a full situated semiosis (semantics, pragmatics, and syntax) with their environments.

## 2. SYNTACTIC AUTONOMY

### 2.1 Semiotic Codes

Semiotics concerns the study of signs/symbols in three basic dimensions: *syntactics* (rule-based operations between signs within the sign system), *semantics* (relationship between signs and the world external to the sign system), and *pragmatics* (evaluation of the sign system regarding the goals of their users) [Morris, 1946]. When Von Neumann's universal constructor interprets a description to construct some automaton, a *semiotic code* [Umerez, 1995] is utilized to map instructions into actions to be performed in some environment to construct the described automaton. When the copier copies a description, only its *syntactic* aspects are replicated. Semiotics leads us to think of symbols not simply as abstract memory tokens, but as material tools [Prem,

1998] for a situated open-ended semiosis of classifying systems with their environments, which requires the definition of components that interact and self-organize with the laws of their environment<sup>2</sup>. Thus, a situated semiotic code presupposes a set of components (e.g. parts and processes) for which the instructions are said to "stand for". Descriptions are not universal as they refer to some building blocks which cannot be changed without altering the significance of the descriptions.

We can see that a self-reproducing organism following this scheme is an entanglement of *symbolic controls* and *component constraints* which is closed on its semantics only through its repercussions in an environment. Pattee [1982, 1995a] calls such a principle of self-organization *semantic closure*. Perhaps a better description would be to refer to it as *semiotic closure* since this principle explicitly recognizes the three semiotic dimensions of semantics, pragmatics and syntax [Rocha, 1998c].

The implications of the component (enabling and restraining) constraints for systems observing a semiotic closure in situation with their environments have been investigated conceptually and experimentally in Rocha [1996, 1997, 1998a, 1998c]. The study of genetic systems with richer syntactics, in particular the modeling of the RNA editing system, have also been explored in Rocha [1995, 1997, 1998c]. Here I examine the emergence of syntax in systems in selected self-organization with their environments, particularly, with respect to the notion of autonomy.

## 2.2 Reference and Syntactic Autonomy

A classifying self-organizing system is autonomous if all processes that establish and sustain its dynamics are internally produced and re-produced over and over again [Maturana and Varela, 1987]. These are the systems capable of *self-reference* (including hurricanes) [Salthe, 1998]. But how autonomous are the systems that follow some form of situated semiosis with their environments? Given the arguments for Selected Self-Organization, we know that it is the environment which ultimately selects the dynamic configurations of classifying systems. The structural coupling between system and environment [Rocha and Joslyn, 1998] on which situated semiosis is

based requires this structural openness [Maturana, 1987; Pask, 1992], *other-reference* [Salthe, 1998; Hoffmeyer, 1998]], or external scaffolding [Clark, 1997]. Semantics is therefore defined only by the situated, pragmatic, conjugation of system and environment, which indicates that even though the organization of the dynamic components of self-organizing classifying systems is autonomous, these systems are not semantically autonomous. But is there any kind of semiotic autonomy in evolutionary systems?

Biological systems have developed a system of structural perturbation of their self-organization clearly based on a (genetic) code that essentially implements Von Neumann's scheme of inert symbolic descriptions (section 1.4). It is undeniable that this syntactic code is completely specified within organisms since its reading and constructing machinery is found within each cell: an autonomous code defined by specific syntactic rules. Even though environmental conditions clearly affect what is decoded in different circumstances [Rocha, 1995, 1998c], the code itself remains fixed. The ability to generate such a powerful system of assembly of self-organizing encoded components for the construction of evolving classifying systems [Rocha, 1997, 1998c], is the one defining characteristic of all known life forms, which somehow produced an *autonomous syntax* for a more efficient situated semiosis with the environment.

A consequence of this argument is that the concept of autonomy alone is not enough to characterize living organisms, unless by that we mean, in addition to material autonomy (organizational closure), also *syntactic autonomy*. In other words, situated semiosis is based on organizational closure (self-organization, self-reference, etc), semantic openness by virtue of a situated coupling to an environment (other-reference), and syntactic autonomy (syntactic stability or inert codes). Hoffmeyer [1998] pursues a similar argument to insist that it is the stable integration of self-reference and other-reference (established by the syntactic autonomy of the Von Neumann code, I argue here) which establishes the minimum requirement for an *umwelt* [Uexküll 1982 (1940)], or evolving personal categorizations of an environment, and thereby sets living systems apart from all their non-living predecessors.

<sup>2</sup>In the natural world, self-organization results from physical law while in artificial environments it results from some set of pre-defined state-determined rules acting as laws of an artificial environment. See Rocha and Joslyn, 1998 for an extended discussion of laws and rules in artificial environments.

Regarding cognitive systems, it is possible that human language established a system of structural perturbation of self-organizing processes similar to the genetic scheme [Henry and Rocha, 1996; Rocha, 1997, 1998b], and that somehow, the brain has evolved another type of coded semiotic closure with its environment. Language may be

a syntactic tool that allows cognition the ability of open-ended conceptual variety. For this reason, the study of the emergence of syntactic autonomies is relevant for both evolutionary systems research and cognitive science. Next a model is discussed which may give some insights into the problem of the origin of syntactic autonomy.

### 3. SYNTACTIC AUTONOMY IN COMPUTATIONAL ENVIRONMENTS

#### 3.1 Emergent Particle Computation

A very interesting problem that genetic algorithms (GA's) have been used successfully in, is the evolution of Cellular Automata (CA) rules for the solution of non-trivial tasks<sup>3</sup>. Certain CA rules are capable of solving global tasks assigned to their lattices, even though their transition rules are local (each cell computes its next value given the current value of the cells in its immediate neighborhood). One such tasks is usually referred to as the *density task*: given a randomly initialized lattice configuration (IC), the CA should converge to a global state where all its cells are turned "ON" if there is a majority of "ON" cells in the IC, and to an all "OFF" state otherwise. This rule is not trivial because the local rules of the component cells do not have access to the entire lattice, but can only act on the state of their immediate neighborhood.

Crutchfield and Mitchell [1995] used a GA to evolve the CA rules for such a task. The GA found a number of fairly interesting rules, but a few of the runs evolved very interesting rules (with high fitness) which create an intricate system of lattice communication. Basically, groups of adjacent cells propagate certain patterns across the lattice, which as they interact with other such patterns "decide" on the appropriate solutions for the lattice as a whole. An intricate system of signaling patterns and its communication syntax has been identified, and can be said to establish the emergence of embedded-particle computation in evolved CA's [Ibid; Hordijk, Crutchfield, Mitchell, 1996]. The emergent signals (or embedded particles) refer to the borders of the different patterns that develop in the space-time diagrams. If the areas inside these patterns are removed, their boundaries can be identified as a system of signals with a definite syntax, or emergent logic grammar. This syntax is based on a small number of signals,  $\alpha$ ,  $\beta$ ,  $\delta$ ,  $\gamma$ ,  $\eta$ , and  $\mu$ , and a small

number of rules such as:  $\alpha + \delta \rightarrow \mu$ , meaning that when signals  $\alpha$  and  $\delta$  collide, the  $\mu$  signal results. Please refer to the references above for more details.

These experiments are very interesting because from the interaction of self-organization (CA's) and selection (GA) a very simple semantics emerges from the selective pragmatics of the GA: the CA rule either classifies its initial lattice configurations correctly or incorrectly. Now, most CA rules evolved with this set up show very simple space-time patterns: they try to solve the problem by block-expansion, that is, when large neighborhoods of either "ON" or "OFF" states exist in the initial configuration, they are expanded. These block expansion rules solved the task in typical dynamical fashion: by taking into account only local information.

Instead, the system of particle computation uses signals that are capable of integrating distant global information to solve the task. These CA rules rely on a system of personal (to the CA rule) signals used to communicate across the lattice and compute the answer to the task: an autonomous sign system that grants great selective advantage to the rules capable of developing it. The particle computation system truly introduces a qualitatively different way of solving the task: through the emergence of autonomous syntax, which allows certain rules to gain access to global lattice information. Obviously, such a system does not possess the rich self-reproduction scheme of Von Neumann, but it does show how the emergence of autonomous syntax grants simple dynamical systems the ability to move from trivial to non-trivial classification of their interaction with an environment.

#### 3.2 Increasing Arbitrariness: Logical Tasks

The signals of the emergent particle computation system in CA's, even though being a small set of discrete entities, are not full-fledged symbols in the senses described in section 1, because they do not possess the degree of arbitrariness required of pure symbols: the syntax is specific to the task solved. However, very similar signals and grammars can be evolved to solve different tasks, e.g. the synchronization task [ Hordijk, Crutchfield, and Mitchell, 1996]. In other words, this class of CA's can develop similar signals to solve different problems.

To increase the arbitrariness of the emergent syntax of these rules, we can evolve rules that are good at solving several tasks. I have conducted some experiments to evolve CA rules with radius 3 which can solve both the density task and some related logical tasks. To implement

---

<sup>3</sup> This work has been mainly pursued by the Adaptive Computation Group of the Santa Fe Institute. I am indebted to Melanie Mitchell and Wim Hordijk for many illuminating discussions.

logical tasks, we divide the CA lattice in two halves (the center cell is not used). The first half is interpreted as the first bit, and the second half as the second bit. A bit is "ON" if there is a majority of "ON" cells in its half, and "OFF" otherwise. Notice that since the boundary conditions of the lattice are periodic, this lattice has two boundaries between the two halves or bits. The cells on the neighborhood of these boundaries compute their values from cells in both halves, which in most cases makes the computation on these boundaries unreliable. However, since we are looking for global communication across the lattice, we expect the local errors at the boundaries not to be too relevant for the global computation, especially as lattices grow in size.

We can now define such logical tasks as the AND and the OR task, according to the values of the bits. For the AND (OR) task, for all values of the bits the lattice should converge to an all "OFF" ("ON") state, except when both bits are "ON" ("OFF"). These tasks are both related to the density task because when the density of both halves is below (over) 0.5, both bits are "OFF" ("ON"), leading to a desired final lattice with all cells "OFF" ("ON"). They differ for the cases when the two halves of the lattice have opposing densities. In other words, these tasks should perform the density task in each half, and then integrate the results, with the AND (OR) task biased by "OFF" ("ON") information on either half.

Several rules were evolved with a GA whose initial population was composed of some of the best rules evolved so far for the density task, and whose fitness functions was derived from presenting each rule with 100 different initial lattices, 50 to be analyzed by the density task, and the other 50 by either the AND or the OR task. The 50 rules to be presented to the density task have their density of "ON's" uniformly distributed over the unit interval (just as the experiments described in 3.1). The 50 rules presented to the AND (OR) task are biased to a uniform distribution of lattices leading to at least one bit "OFF" ("ON") 50% of the time, and both bits "ON" ("OFF") the other 50%. If we were to use an unbiased generation of lattices, only 25% of the time would the case of both bits "ON" ("OFF") be generated, making rules that always tend to "OFF" ("ON") always too favorable.

From these experiments, several rules were evolved that can solve both the density task and one of the logical tasks very well. Also, as expected, and unlike the density task, the performance of the logical tasks actually increases with the lattice size as the boundary errors described earlier loose relevance. Details of these experiments will be shown at this conference.

The relevance of these experiments is that they show that there is a family of particle computation rules which with a few mutations can develop a system of particle computation that can solve two different, yet related, tasks. That is, the particle computation systems provides a self-organizing system with the ability to adapt to a new environment that requires the solution of two similar tasks. In other words, it has the ability to evolve into a system that with the same syntax can effectively solve a related class of problems and not just one single task. In this case the class of tasks includes the density task and some logical task that is coherent with the density task. The ability to solve more than one task increases the arbitrariness of the emergent syntax of these rules, as the same syntactic rules of particle computation are used to compute different tasks.

#### 4. EMERGENCE OF ARTIFICIAL SYMBOLS

These particle-computation CA rules possess the intertwined semantics and pragmatics of selected self-organization (CA rules evolved with a GA), plus a primordial autonomous syntax (the emergent grammar of the particles) in an artificial environment. In this sense they are a case of a purely computational situated semiosis as described in section 2, which represents a truly exciting new development in evolutionary systems research. These experiments provide an abstract model of how signs can emerge from purely dynamical interactions evolved under an artificial situated semiosis. These experiments seem to indicate that it is possible to evolve symbols from artificial matter, in other words, that it is possible to study syntactic autonomy, so important to distinguish living from non-living systems as discussed in section 2, in computational environments.

#### References

- Clark, Andy [1997]. *Being There: Putting Brain, Body, and World Together Again*. MIT Press.
- Crutchfield, J.P. and M. Mitchell [1995]. "The evolution of emergent computation." *Proc. Nat. Acad. Sciences, Computer Sciences*, Vol. 92, pp. 10742-10746.
- Dawkins, R. [1976]. *The Selfish Gene*. Oxford University.
- Dennet, D.C. [1995]. *Darwin's Dangerous Idea: Evolution and The Meanings of Life*. Simon and Schuster.
- Eldridge, N. [1995]. "A battle of words." In: *The Third Culture: Beyond The Scientific Revolution*. J. Brockman. Simon and Schuster, pp.119-128.
- Gould, S.J. [1989]. *Wonderful Life*. Norton.
- Henry, C. and L.M. Rocha [1996]. "Language theory: consensual selection of dynamics." *Cybernetics and Systems*. Vol. 27, pp 541-553.

Hoffmeyer, Jesper [1998]. "The emergence of externalism". To appear.

Hordijk, W., J.P. Crutchfield, and M. Mitchell [1996]. "Embedded-particle computation in evolved cellular automata." In: *Proceedings of Physics and Computation 96*.

Kauffman, S. [1993]. *The Origins of Order: Self-Organization and Selection in Evolution*. Oxford university Press.

Maturana, H. and F. Varela [1987]. *The Tree of Knowledge: The Biological Roots of Human Understanding*. New Science Library.

Mitchell, Melanie [1996]. "Computation in cellular automata: a selected review." In: *Nonstandard Computation*. H.G. Schuster and T. Gramms (eds.). VCH Verlagsgesellschaft.

Morris, C.W. [1946]. *Signs, Language, and Behavior*. G. Braziller Publishers.

Pask, Gordon [1992]. "Different kinds of cybernetics." In: *New Perspectives on Cybernetics*. G. van de Vijver (ed.). Kluwer, pp.11-31.

Pattee, Howard H. [1982]. "Cell psychology: an evolutionary approach to the symbol-matter problem." *Cognition and Brain Theory*. Vol. 5, no. 4, pp. 191-200.

Pattee, Howard H. [1995a]. "Artificial life needs a real epistemology." In: *Advances in Artificial Life*. F. Moran, et al. Springer-Verlag, pp 23-38.

Pattee, Howard H. [1995b]. "Evolving self-reference: matter, symbols, and semantic closure." *Communication and Cognition - Artificial Intelligence*. Vol. 12, nos. 1-2, pp. 9-27.

Prigogine, I. [1985]. "New perspectives on complexity." In: *The Science and Praxis of Complexity*. United Nations Library.

Rocha, Luis M. and Cliff Joslyn [1998]. "Models of Embodied, Evolving, Semiosis in Artificial Environments." In: *Proceedings of the Virtual Worlds and Simulation Conference*. C. Landauer and K.L. Bellman (Eds.). The Society for Computer Simulation, pp. 233-238.

Rocha, Luis M. [1995]. "Contextual genetic algorithms: evolving developmental rules." In: *Lecture Notes on Artificial Intelligence: Advances in Artificial Life*. J. Moran, A. Moreno, J.J. Merelo, and P. Chacon (Eds.). Springer-Verlag, pp. 368-382.

Rocha, Luis M. [1996]. "Eigenbehavior and symbols." *Systems Research*. Vol. 13, No. 3, pp. 371-384.

Rocha, Luis M. [1997]. *Evidence Sets and Contextual Genetic Algorithms: Exploring Uncertainty, Context and Embodiment in Cognitive and biological Systems*. PhD. Dissertation. State University of New York at Binghamton.

Rocha, Luis M. [1998a]. "Selected self-organization and the Semiotics of Evolutionary Systems." In: *Evolutionary Systems*. S. Salthe and G. Van de Vijver. Kluwer Academic Publishers. (in press).

Rocha, Luis [1998b]. "Categorizing Databases, Evidence Sets, and Evolutionary Constructivism." *International Journal of Human -Computer Studies*. (In Press).

Rocha, Luis M. [1998c]. "Semiotic Closure and Contextual Genetic Algorithms: Artificial Matter and Symbols" *BioSystems*. (Submitted).

Salthe, Stanley N. [1985]. *Evolving Hierarchical Systems*. Columbia University Press.

Salthe, Stanley N. [1993]. *Development and Evolution: Complexity and Changes in Biology*. MIT Press.

Salthe, Stanley [1998]. "The role of selection theory in understanding evolutionary systems." In: *Evolutionary Systems*. S. Salthe and G. Van de Vijver. Kluwer Academic Publishers (In Press).

Uexküll, J.V. [1982 [1940]]. "The theory of meaning." *Semiotica*. Vol. 42, no. 1, pp. 25-87.

Umerez, Jon [1995]. "Semantic Closure: a guiding notion to ground Artificial Life." In: *Advances in Artificial Life*. F. Moran, A. Moreno, J.J. Merelo, and P. Chacon (eds.). Springer-Verlag, pp. 77-94.

van Gelder, Tim [1991]. "What is the 'D' in 'PDP': a survey of the concept of distribution." In: *Philosophy and Connectionist Theory*. W. Ramsey et al (Eds.). Lawrence Erlbaum.

von Neumann, J. [1966]. *The Theory of Self-Reproducing Automata*. University of Illinois Press.