

Simulations of Embodied Evolving Semiosis: Emergent Semantics in Artificial Environments

LUIS MATEUS ROCHA[†] and CLIFF JOSLYN[‡]

Los Alamos National Laboratory

Los Alamos, NM 87545

[†] MS P990, e-mail: rocha@lanl.gov or rocha@santafe.edu, URL: <http://www.c3.lanl.gov/~rocha>
[‡] MS B265, e-mail: joslyn@lanl.gov, URL: <http://www.c3.lanl.gov/~joslyn>

In: *The Proceedings of the 1998 Conference on Virtual Worlds and Simulation 1998*. Landauer C. and K.L. Bellman (Eds). The Society for Computer Simulation International, pp. 233-238.

Abstract.

As we enter this amazing new world of artificial and virtual systems and environments in the context of human communities, we are interested in the development of systems and environments which have the capacity to grow and evolve their own meanings in the context of this community of interaction. In this paper we analyze the necessary conditions to achieve systems and environments with these properties: 1) a coupled interaction between a system and its environment; 2) an environment with sufficient initial richness and structure to allow for; 3) embodied emergent classification of that environment-system coupling; 4) which is subject to pragmatic selection.

1 ARTIFICIAL AND REAL ENVIRONMENTS AND SYSTEMS

1.1 Systems, Environments and Constraints

A systems analysis divides the universe between a "system" and an "environment" where the first is embedded. A system interacts with its environment by input-output relations: making distinctions in that environment as some kind of generalized "perception relation", making decisions about appropriate actions to take back into that environment, and then taking those actions. Of course, these distinctions can be drawn relatively arbitrarily, with system/environment boundaries drawn by investigators, or by the systems themselves, for pragmatic purposes. Generally, the system is regarded as somewhat autonomous, dynamic, and fast-scaled, connected to the static, slower-scaled environment through input/output channels of sensors for measurement into the system from the environment, and effectors for action into the

environment by the system. Also, both the system and the environment can be further decomposed into subsidiary environment/system couplings, for example the sensory inputs being the environment to an agent within the system, or the environment containing, or in the co-evolutionary limit just consisting of, many other "systems" which are all mutually interacting.

Both the system and the environment contain structures, or constraints on the range of possible values, configurations, or other variable properties [Joslyn, 1997]. These serve as sources of information between the environment and the system. As the system-environment coupling develops or evolves, these structures also change. Typically, constraints grow in the system as it learns, adapts, or evolves in response to the environment, the effects of its actions in the environment, and the effects of the actions of other systems on their collective environment.

1.2 Artificial and Real Systems and Environments

More specifically, we are interested in considering systems and/or environments which are "natural" or "real" in some sense, and those which are "artificial", "synthetic", or "virtual". The particular cases will help determine the range of kinds of informational constraints possible in the system-environment coupling. In all, we can consider four categories:

System	Environment	
Real	Real	Reality: organisms in the world
Artificial	Real	Robotics: artificial mechanism interact in the world which includes other such devices and organisms
Real	Artificial	Virtual reality, MUDs
Artificial	Artificial	Agents in artificial environments, e.g. artificial life experiments

Given these possibilities of real and artificial system-environment couplings, it is important to evaluate what kind of interactions they can evolve. More specifically, we need to ask what kinds of emergent behaviors and semantics can be evolved in them, or how the coupled constraints in the system and environment can coevolve semiotically. To effectively discuss this, we need to ground the debate in what we know about semantic emergence in natural and artificial systems, that is, in the discourse of complex systems research.

2. SEMANTIC EMERGENCE AND SELECTED SELF-ORGANIZATION

2.1 Emergent Classification

The attractor behavior of any dynamical system is dependent on the structural operations of the latter, e.g. the set of boolean functions and connections in a boolean network. In addition, 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). Not all possible distinctions in some environment can be “grasped” by the self-organizing system: it can only classify those aspects of its situated interaction with and environment which result in the maintenance of some internally stable state or attractor. Furthermore, the behavior of a situated agent in an environment is, in this sense, the result of how the agent classifies the environment. Thus, behavior is itself an emergent property of the lower level dynamics that implements the agent. The process of obtaining novel classifications of an environment, by an autonomous self-organizing system, can be referred to as *emergent classification*. 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.

2.2 Selected Self-Organization

A computational neural network by itself can classify an environment, but the processes (e.g. a back-propagation algorithm) that make it improve its classifying ability are external to the network. Similarly, Evolutionary strategies rely on internal random variation which must ultimately be externally selected. It is precisely the ability of such systems to adapt their structure in order to better classify a changing environment that leads to emergent classification of a particular environment. For a classifying self-organizing system to change its classification ability, structural changes must be performed to alter its attractor landscape. When the structure responsible for a given dynamics is changed, we obtain a new environmental classification (e.g. weight changes in a neural network). 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*.

Explicit control of a classifying system’s structure would amount to the choice of a particular dynamics for a certain task and can be referred to as *learning*. Under implicit control, the self-organizing system is subjected to some variation of its structure which may or may not be good enough to perform our task. Those self-organizing systems which are able to perform the task are thus *externally selected* by the environment to which they are structurally coupled. For selection to occur we must have some internal vehicle for classification — there must be different alternatives. The attractor landscape offers these alternatives. Selected self-organization relies on some system-environment coupling of structure. It also explicitly emphasizes a second dimension of an embodied semiosis of autonomous systems with their environments. If classification implies semantic emergence, selection implies pragmatic environmental influence.

2.3 Von Neumann, the Symbolic Advantage and the Parts Problem

Von Neumann’s [1966] model of self-replication is a systems-theoretic criteria of open-ended evolution [for a detailed discussion of this model see Rocha, 1996, 1998]. 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. 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 commands are not distributed over patterns of activation of the components of a dynamic system,

but instead localized on “inert” structures which can be used at any time — a sort of random access memory. For instance, in the genetic system (which Von Neumann’s model conceptually describes), most any sequence of nucleotides is possible, and its informational value is almost completely independent of the particular dynamic behavior of DNA or RNA. The 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 [Pattee, 1995a]. We can refer to this mechanism as *description based evolution* [Rocha, 1996, 1997, 1998] It introduces the third level of an evolving, embodied semiosis of autonomous systems with their environments: syntax.

Notice that according to Von Neumann’s own formulation, a symbol system utilized for the construction of self-reproducing systems is not an isolated artifact. Rather, in the context of construction, a symbol system entails a set of available parts. That is, construction blueprints are cast on a symbol system whose primitives are a finite set of parts. In the case of self-reproducing automata, these parts are “and”, “or” and other logical operators, and in the case of the genetic code the parts are the set of aminoacids (the symbols are codons or sets of 3 nucleotides). It is in this sense that open-ended evolution must be understood. A given material symbol system cannot represent everything, only what its primitive parts can construct. Natural selection is open-ended for any form that can be constructed through folding aminoacid chains.

The syntactic dimension of Von Neumann’s scheme enlarges the domain of classification in an open-ended fashion due to the description-construction dichotomy [Rocha, 1996, 1998], but it establishes another facet of embodiment with its Parts Problem. A particular materiality is tied to specific construction building blocks. The richer the parts, the smaller the required descriptions, but also the smaller the number of classifiable categories or constructed morphologies. For instance, Von Neumann used simple building blocks such as “and” and “or” gates to build his automaton, which in turn required a 29 state cellular automata lattice and very complicated descriptions. Arbib[1966] was able to simplify von Neumann’s model greatly by utilizing more complicated logical building blocks. Likewise, the genetic system does not need to describe all the chemical/dynamical characteristics of a “desired” protein, it merely needs to specify an aminoacid chain which will itself self-organize (fold) into a functional configuration with some reactive properties. In other words, a given materiality, that is, a given set of parts such as amino acids, provides intrinsic dynamic richness which does not

have to be specified by the symbol system on which construction commands are cast [Moreno, et al, 1994] making descriptions much smaller. Embodiment provides this kind of *material information compression*.

The other side of Embodiment, is that it also constrains the universe of possible constructions (universe of open-endedness). Living organisms are morphologically restricted to those forms that can be made out of aminoacid chains through the genetic code, while in principle, a formal symbol system, stripped as it is from any materiality, can describe anything whatsoever. Of course, this ‘in principle’ is seriously, and easily, constrained by computational limits, as formal descriptions are much larger than material ones. A complete formal description of a protein would have to include all of its physical characteristics from the atomic to the chemical level, while a gene needs only a description of an aminoacid sequence. A discussion of how to incorporate the notion of embodiment in computational models of evolutionary systems, in order to obtain some form of descriptorial information compression is presented in Rocha [1995, 1997].

2.4 Situated Action

Traditionally, AI relied strongly on models of representation and direct perception of the world. It was mostly preoccupied with functional semantics. The control of its robotic artifacts, for instance, was solely based on the high-level symbol-manipulation of semantic categories. Artificial Life, mostly through the work of Brooks [1991], whose behavior language replaced the traditional high-level control of robots by a scheme of functional modularization by behavior generating modules, changed all this. Instead of a high-level computation of behavior, the bottom-up (emergentist) self-organization of simpler components produces a variety of behaviors depending on the interaction of a robot with its environment. “Situated” does not mean merely material, but interactive. The material (structural) coupling of the robot with its environment is the source of behavior, and not just the robot control system alone. In other words, the modeling of living and cognitive systems is moved to the dynamics of self-organization of a network of components and its interaction with an environment (selected self-organization).

It can be argued that the behavior modules utilized are still too high level and do not allow the sort of plasticity that living systems observe. Indeed, it is not always obvious how to physically compartmentalize behavior modules: a bird’s wing is both an airfoil and engine at the same time [Rosen, 1993]. The sort of behavioral decomposition pursued by Brooks may not offer yet the kind of entailment or network

causality found in living organisms [Rosen, 1991; Prem, 1995] which allows for genuine evolution of new behaviors [Cariani, 1992], however, it does mark a very important shift in the practice of AI: the transition from a central control to an emergentist, self-organizing, practice of autonomous agents. Cognition is no longer modeled as the creation of universal classifications of the world, but as the embodied, evolving, interaction of a self-organizing system with its environment. Whichever way situated robots solve a problem, it is done by the construction of their own classifications, given the set of low level components they have available, as they interact with their environment, and not by externally imposed rules.

3. LAWS AND RULES IN ARTIFICIAL MEDIA

As discussed above, what defines evolving systems is that they implement an embodied, evolving, semiosis (EES) [Rocha, 1997] with their environments. What is important here is to discuss whether artificial systems can observe an equivalent kind of organization. In summary, EES requires:

1. Material self-organization in situated interaction with an environment
2. Semantic emergence: classification based on structural perturbation of self-organizing dynamics by descriptions that define a material symbol system
3. Selection in an environment leading to open-ended evolution

Self-organization in natural systems is a result of inexorable laws of physics. Living organisms can generate an open-ended array of morphologies and modalities, but they can never change these laws. “Physical laws describe those events over which organisms have no control” [Pattee, 1995a, page24]. It is from these constant laws (and their initial conditions) that all levels of organization from life to cognition and social structure emerge. These levels of emergence typically produce their own principles of organization, which we can refer to as rules, but all of these cannot control or escape physical law and are “neither invariant nor universal like laws” [Pattee, 1995b, page 27]. The question of what kinds of rules can emerge from deterministic or statistical laws is at the core of the field of Artificial Life [Langton, 1989]. It is also pretty much the question of generating emergent semantics in artificial environments, given the discussion of selected self-organization in terms of emergent classification above. “Without principled restrictions this question will not inform philosophy or physics, and will only lead to disputes over nothing more than matters of taste in computational

architectures and science fiction.” [Pattee, 1995b, page 29] For artificial environments to be relevant for science in general, the same categories of laws/initial conditions and rules that we recognize in the natural world, need to be explicitly included in an artificial form.

3.1 Evaluating possible system-environment couplings

An important question, is how open-ended can semantic emergence be? To observe rich emergence in artificial environments we need to:

- a. SPECIFY THE DYNAMICS OF SELF-ORGANIZATION: specify laws and their initial conditions, which are responsible for the variety of the artificial environment (including agents) and the emergence of context-specific rules.
- b. OBSERVE EMERGENT OR SPECIFY CONSTRUCTED SEMANTICS: identify emergent or pre-programmed, but changeable, rules that generate agent behavior in tandem with environmental laws.
- c. PROVIDE A PRAGMATIC SELECTION CRITERIA: create or identify a mechanism of selection (pragmatics) so that the semantics identified in b is grounded in a given environment.

In particular systems and environments the laws and rules take different forms:

- ! REAL SYSTEMS IN REAL ENVIRONMENTS: systems with morphogenetically developed and biologically evolved neurological sensory and action modalities and structures (rules) in an environment defined by physical law which lead to geologically evolved structures, biologically co-evolved ecologies, etc.
- ! ARTIFICIAL SYSTEMS IN REAL ENVIRONMENTS: systems with internal lower-level physical rules leading to emergent, bottom-up behavior (situated robotics) or programmed rule-based (physical) behavior (traditional robotics) defining world-models and decision-making mechanisms, in an environment defined by physical law.
- ! REAL SYSTEMS IN ARTIFICIAL ENVIRONMENTS: systems with morphogenetically developed and biologically evolved neurological sensory and action modalities and structures (rules) in an environment defined by virtual physics, formal ontologies, external databases, user interfaces, etc. The key here is to define rich enough virtual laws which can couple with the modalities evolved in the real physical world, in order to observe true

semantic emergence with its own selective mechanisms.

- ! **ARTIFICIAL SYSTEMS IN ARTIFICIAL ENVIRONMENTS:** systems with internal lower-level virtual rules leading to emergent, bottom-up behavior (e.g. artificial life) or programmed rule-based (virtual) behavior (traditional AI) defining world-models and decision-making mechanisms, in an environment defined by virtual physics, formal ontologies, external databases, user interfaces, etc. Here besides rich artificial laws, we further require rich artificial rules.

Artificial Life may be a source of relevant examples for other completely computational media such as the World Wide Web. In artificial life, more than life-as-it-could-be [Langton, 1989], we further need physics-as-it-could-be [Pattee, 1995]. Good examples of this are the emergent computation experiments on Cellular Automata (CA) [Mitchell and Crutchfield, 1995] with Genetic Algorithms (GA's). There is an environment which requires a non-trivial task to be performed (laws). The CA provides the lower-level virtual rules that lead to emergent behavior in such an environment, provided a selection mechanism implemented by the GA (selected self-organization). The evolved, emergent, higher level structures (particle computation) define an emergent semantics with a primordial syntax to solve the task or environmental survival requirements. The same structures can be used to solve different tasks in different environments. This is an effective example of emergent semantics in artificial system/environment couplings.

The interplay between laws and rules is in particular system-environment couplings is, of course, highly particular to their specific natures. However, certain regularities about laws and rules in real and artificial systems and environments can be isolated from this complexity:

1. **Real Systems:** Organisms operate with rule-based action and sensory modalities which have been morphogenetically developed and biologically evolved in real environments and adhering to physical law.
2. **Real Environments:** Physical laws have led to particular geological structures, for example; rules are in the biologically co-evolved ecologies.
3. **Artificial Systems:** At their core, these are rule-based, although they have to interact with whatever the laws are of the environments in which they are embedded.

4. **Artificial Environments:** These are environments necessarily defined by the "laws" of their virtual physics, whether formal ontologies (Usschold and Gruninger, 1996), external databases, user interfaces, etc. Their rules, therefore, create a form of "virtual culture", for example the initial placement of objects in a MUD room.

3.2 Artificial Environments

Consider the range of artificial environments:

- ! **VIRTUAL REALITY:** Intended to provide an accurate model of the real environment. Representations are usually continuous, spatio-temporal dimensions with Euclidean metrics; interfaces are graphical; a virtual physics (laws) of the model world is included.
- ! **MUDS:** Intended to provide a logical representation of a synthetic pseudo-real environment. Representations are discrete, formal, linguistic, high dimensional with no metrics; interfaces are textual; both laws (connections among rooms, properties of objects) and rules (changeable structures, properties, etc) are initially set up. The more the rules are pre-defined, the least room for genuine semantic emergence there will be. The more these rules are left to be a result of self-organization, the more semantic emergence one can expect to see.
- ! **WEBS:** Intended to provide a graph-theoretical representation of an abstract, logical space. Representations are primarily textual, with minimal graphical elements; very high dimensional, graph-theoretical metrics (e.g. path lengths between pages); environment is almost entirely lawful (web topologies fixed by authors). Very little room for semantic emergence.
- ! **EXPERT SYSTEMS:** Intended to provide logical representations of a knowledge system. Representations are textual; no metrics; initial environment is based on law-like formal ontologies, potential for user-developed rules as extensions which can potentially allow for semantic emergence if there is selected self-organization of these rules.

Of course, there is also a vast diversity in the range of possible artificial environments. The broad classes of artificial environments each present their systems (whether real or artificial) with different representational modalities and forms

of reference. The possible system-environment couplings will necessarily be constrained by these modalities.

4. CONCLUSIONS

To address emergent semantics in artificial environments, in particular web environments, we need the following, partially equivalent, statements to be observed.

- ' Understand the modalities of web-based virtual environments as a NEW medium, related to, but unlike, any of the others considered above (VR, MUDs, webs, application programs).
- ' Imbue them with a sufficiently rich "virtual physics".
- ' Allow the definition of system-specific, low or high level, rules of behavior.
- ' Set up a mechanism of explicit or implicit selection.
- ' Environments and systems should be complex enough to allow the mechanisms of selected self-organization to occur in the situated interaction of agents in artificial or real environments. This requires the evolving, embodied, semiosis previously discussed.
- ' Environments cannot be entirely lawful, so that users can interact with both the virtual "physical" world and each other as actors in that environment.
- ' There should be sufficient interaction to have these agents constructing, preferably in a bottom-up emergent fashion, a new emergent, shared semantics by virtue of their situated interactions with the "world" and each other.

References:

- Arbib, M.A. [1966]. "A simple self-reproducing universal automaton." *Information and Control*. Vol. 9, pp. 177-189.
- Brooks, R.A. [1991]. "Intelligence without reason." In: *Proceedings of the 12th International Conference on Artificial Intelligence, San Mateo 1992*. Morgan Kaufmann.
- Cariani, Peter [1992]. "Emergence and Artificial Life." In: *Artificial Life II*. C. Langton et al. (eds.). SFI Series in the Sciences Complexity. Addison-Wesley, pp. 775-797.
- Joslyn, 1997: "Semiotic Aspects of Control and Modeling Relations in Complex Systems", in: *Control Mechanisms for Complex Systems*, ed. Michael Coombs, Addison-Wesley, Redwood City CA
- Langton, C.G. [1989]. "Artificial life." In: *Artificial Life*. C. Langton (ed.). Addison-Wesley.
- Maturana, H. and F. Varela [1987]. *The Tree of Knowledge: The Biological Roots of Human Understanding*. New Science Library.
- Moreno, A., A. Etxeberria, and J. Umerez [1994]. "Universality without matter?." In: *Artificial Life IV*. R. Brooks and P. Maes (eds.). MIT Press, pp. 406-410.
- 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]. "Evolving self-reference: matter, symbols, and semantic closure." *Communication and Cognition - Artificial Intelligence*. Vol. 12, nos. 1-2 (Rocha[1995a]), pp. 9-27.
- Pattee, Howard H. [1995b]. "Artificial life needs a real epistemology." In: *Advances in Artificial Life*. F. Moran, et al. Springer-Verlag. pp 23-38.
- Prem, Erich [1995]. "Grounding and the entailment structure in robots and Artificial Life." In: *Advances in Artificial Life*. F. Moran, A. Moreno, J.J. Merelo, and P. Chacon (eds.). Springer-Verlag, pp. 39-52.
- Rocha, Luis M. [1995]. "Contextual genetic algorithms: evolving developmental rules." In: *Advances in Artificial Life*. J. Moran, A. Moreno, J.J. Merelo, and P. Chacon (Eds.). Springer, 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. [1998]. "Selected self-organization and the semiotics of evolutionary systems." In: *Evolutionary Systems* S. Salthe and G. Van de Vijver. Kluwer. (in press).
- Rosen, Robert [1991]. *Life Itself*. Columbia University Press.
- Rosen, Robert [1993]. "Bionics revisited." In: *The Machine as a Metaphor and Tool*. H. Haken, A. Karlqvist, and U. Svedin (eds.). Springer-Verlag, pp. 87-100.
- Uschold, Mike and Gruninger, Michael [1996] "Ontologies: Principles, Methods and

Applications", *Knowledge Engineering Review*,
vol. 11:2, pp. 93-136.
von Neumann, J. [1966]. *The Theory of Self-Reproducing
Automata*. University of Illinois Press.