SSIE 480B/580B:
Evolutionary Systems and Artificial Life
Luis Rocha (Copyright)

Lecture Notes

Syllabus

Evaluation

Artificial Life computational project or specific topic paper to be discussed with instructor (70%) . Two small assignments (15% each): 1) self-organizing system; 2) Evolutionary System.

1. What is Life?

In order to understand Artificial Life we need to discuss the concept of *Life*:

> "life adj.— n.1. the general condition that distinguishes organisms from inorganic objects and dead organisms, being manifested by growth through metabolism, a means of reproduction, and internal regulation in response to the environment. 2. the animate existence or period of animate existence of an individual. 3. a corresponding state, existence, or principle of existence conceived of as belonging to the soul. 4. the general or universal condition of human existence. 5. any specified period of animate existence. 6. the period of existence, activity, or effectiveness of something inanimate, as a machine, lease, or play. 7. animation; liveliness; spirit: The party was full of life. 8. the force that makes or keeps something alive; the vivifying or quickening principle." [Random House Webster's Dictionary]

The definitions above fall into three main categories: (1) life as an organization distinct from inorganic matter (with an associated list of properties), (2) life as a certain kind of *animated behavior*, and (3) life as a special, incommensurable, quality — *vitalism*. Throughout this course we will see that all principles, and indeed all controversies, associated with the study of life fall into one of these categories or the differences between them. It is important to realize that for the most part of the history of humanity, this question was never an issue. Before the study of physics became important, everything was alive: the stars, the skies, the rivers and mountains, etc. There was no non-life, so the concept was of no importance. It is only when the deterministic mechanics of moving bodies become dominant that the question arises. If all matter follows simple physical laws, then what is indeed the difference between life and non-life, between biology and physics?

The third category above has been discarded as a viable scientific explanation, because for science nothing is in principle incommensurable. The traditional scientific approach has lead the study of living systems into a reductionist search for answers in the nitty-gritty of the biochemistry of living organisms. This alternative sees life as nothing more than complicated physics. However, the question remains unanswered since there are many ways to obtain some complicated dynamics, but of all of these, which ones can be classified as *alive*? What *kind of complexity* are we looking for? No one disputes that life is some sort of complex material arrangement, but when do we reach a necessary *threshold of complexity* after which matter is said to be living? Is it a discrete step, or is life a fuzzy concept? To study it must we synthesize organizations with the same threshold of complexity (first category above), or is it enough to simulate its animated behavior (second category above)?

Traditionally life has been identified with material organizations which observe certain lists of properties, e.g. *metabolism*, *adaptability*, *self-maintenance* (autonomy), *self-repair*, *growth*, *replicability*, *evolution*, etc. Most living organisms follow these lists, however, there are other material systems which obey only a subset of these rules, e.g. viruses, candle flames, the Earth, certain robots, etc. This often leads to the view that life is at best a fuzzy concept and at worst something we are, subjectively, trained to recognize — life is what we can eat — and is thus not an objective distinction. Objectively or subjectively, we do recognize some material organizations as being alive. It is perhaps in this ability to recognize and categorize events in our environments that an important difference between living and non-living systems lies.

Life requires the ability to both categorize and control events in its environment in order to survive. This is a common characteristic of all recognized life, together with the ability to store and transmit records of categorizations and controls. It is also the locus of the gap between physics and biology. The laws of physics are, by definition, independent of particular organisms, they are universal, inexorable — therefore immune to any control by an organism— and exist beyond observation. In contrast, biological systems achieve a degree of material control which allows them to categorize and control *relevant* aspects of their material surroundings. It is this "relevant" which gives life an extra attribute to mere physical action-reaction interactions. When an organization is able to recognize and act on aspects of its environment which are important *to its own* survival, we say that the mechanisms by which the organization recognizes and acts are *functional* in reference to the organization itself (*self-reference*). Physics is not concerned with function. A physical description of DNA is certainly possible, but will tell us nothing as to the function of some DNA molecule as a gene relevant to a particular organization. Only in *reference* to this larger organization does a piece of DNA *function* as a gene for some previously *categorized control* (e.g. an enzyme with some effect in an environment).

This issue could be rephrased in terms of the notion of *emergence*. Whatever organization exists after the complexity threshold for life is passed, we may say that it is emergent to the physical level because its attributes cannot be completely explained by the previous level. In particular, function, control, and categorization cannot be explained by physics alone. Notice, however, that emergence does not imply vitalism or dualism. When we say that certain characteristics cannot be explained by physics alone, we mean that they must be explained by different (Pattee's complementary) models for each level, as well as a model for the connection between both levels. In other words, though function, control, and categorization cannot be explained by physics alone, they must nonetheless follow physical laws. In particular, the *origin of life*, is a problem of emergence of categorization and control from a physical milieu.

The definition of emergence as an epistemological, explanatory requirement, is related to the notion of *emergence-relative-to-a-model* [Rosen, 1985; Cariani, 1989] or *intensional emergence* [Salthe 1991]. It refers to the impossibility of epistemological reduction of the properties of a system to its components [Clark, 1996]. It acknowledges that different properties of systems may require different, qualitatively unrelated, epistemological categories and models [Pattee, 1978; Rocha, 1997]. As an example, we can think of phase transitions such as that of water in its transition from liquid to gas. Water and its properties cannot be rephrased it terms of the properties of hydrogen and oxygen, it needs a qualitatively different model.

Artificial life concerns both the simulation and realization of life in some artificial environment, usually the computer. At least regarding the second of its goals, artificial life cannot escape the main issues raised above for biological life. In this course we will discuss what kind of artificial environments are required for simulations and realizations of life. The concepts in italics above will be discussed in more detail in the context of computational environments.

Further Readings and References:

Cariani, Peter [1989].*On the Design of Devices with Emergent Semantic Functions*. PhD.Dissertation. SUNY Binghamton.

Clark, Andy [1996]. "Happy couplings: emergence and explanatory interlock." In: *The Philosophy of Artificial Life*. M. Boden (ed.). Oxford University Press, pp. 262-281.

Pattee, Howard H. [1978]."The complementarity principle in biological and social structures." In: *Journal of Social and Biological Structures* Vol. 1, pp. 191-200.

Rocha, Luis M. [1997]. *Evidence Sets and Contextual Genetic Algorithms: Exploring Uncertainty, Context, and Embodiment in Cognitive and Biological Systems*. PhD. Dissertation. SUNY Binghamton.

Rosen, Robert [1985], *Anticipatory Systems*. Pergamon Press.

Salthe, Stanley N. [1991], "Varieties of Emergence". *World Futures* Vol. 32, pp.69-83

Schrödinger, Erwin [1944]. *What is Life?*. Cambridge University Press.

For next lecture read:

Langton, C. [1989], "Artificial Life" In *Artificial Life*. C. Langton (Ed.). Addison-Wesley. pp. 1-47.

Pattee, H. [1989], "Simulations, Realizations, and Theories of Life". In *Artificial Life*. C. Langton (Ed.). Addison-Wesley. pp. 63-77.

2. The logical Mechanisms of Life

2.1. *Life-As-It-Could-Be*: but, what is *non*-life-as-it-could-be?

> “Artificial Life [AL] is the study of man-made systems that exhibit behaviors characteristic of natural living systems. It complements the traditional biological sciences concerned with the *analysis* of living organisms by attempting to *synthesize* life-like behaviors within computers and other artificial media. By extending the empirical foundation upon which biology is based *beyond* the carbon-chain life that has evolved on Earth, Artificial Life can contribute to theoretical biology by locating *life-as-we-know-it* within the larger picture of *life-as-it-could-be*.” [Langton, 1989, page 1]

> “[AL] views life as a property of the *organization* of matter, rather than a property of the matter which is so organized. Whereas biology has largely concerned itself with the material basis of life, Artificial Life is concerned with the formal basis of life. [... It] starts at the bottom, viewing an organism as a large population of *simple* machines, and works upwards *synthetically* from there — constructing large aggregates of simple, rule-governed objects which interact with one another nonlinearly in the support of life-like, global dynamics. The ‘key’ concept in AL is *emergent behavior*.” [Langton, 1989, page 2]

> “Artificial Life is concerned with tuning the behaviors of such low-level machines that the behavior that emerges at the global level is essentially the same as some behavior exhibited by a natural living system. [...] Artificial Life is concerned with generating lifelike behavior.” [Langton, 1989, pp 4 and 5]

The previous quotes indicate the goals of Artificial Life according to Chris Langton: the search for complex, artificial, systems which observe some kind of life-like, animated, emergent behavior. There seems to be both a desire to obtain an artificial living organization, as well as obtaining some lifelike behavior. The first goal is more ambitious and related to the first definition of life introduced in lecture one, while the second goal is related to the second definition.

The methodology to reach either of these goals is also in line with the notion of emergence mentioned in lecture one: from the non-linear interaction of simple, mechanistic, components, we wish to observe the emergence of complicated, life-like, unpredictable, behavior. Natural living organisms are likewise composed of non-living components. As pointed out in lecture one, the problem of biology is precisely the emergence of life from non-living components. The material components follow, and are completely described, by physical laws, however, a physical explanation of the overall living system is incomplete. Similarly, in Artificial Life, we have formal components obeying a particular set of axioms, and from their interaction, some global behavior emerges which is not completely explained by the local formal rules. Clearly, the formal rules play the role of an artificial physics and the global behavior, <u>if recognized as life-like</u>, plays the role of an artificial biology in this bottom-up approach to complex behavior.

> “Of course, the principle assumption made in Artificial Life is that the ‘logical form’ of an organism can be separated from its material basis of construction, and that ‘aliveness’ will be found to be a property of the former, not of the latter.” [Langton, 1989, page 11]

The idea is that if we are able to find the basic principles of living organization, then the material substract used to realize life is irrelevant. By investigating these basic principles we start studying not only biological, carbon-based, life — life-as-we-know-it — but really the universal rules of life, or life-as-it-could-be. Several problems have been raised regarding this search for a universality without matter [Cariani, 1992; Moreno et al, 1994], which will not be discussed here. What needs to be made more explicit is the relationship between the two distinct goals of AL.

Looking at emergent behavior in formal complex systems in search of *interesting* behavior indicates a certain circularity. If AL is concerned with finding life-like behavior in artificial, universal, systems, we are ultimately binding life-as-could-be to the behavior of life-as-we-know-it by virtue of some subjective resemblance. This can hardly be accepted as the search for universal principles.

> "They say, 'Look, isn't this reminiscent of a biological or a physical phenomenon!' They jump in right away as if it's a decent model for the phenomenon, and usually of course it's just got some accidental features that make it look like something." [Jack Cowan as quoted in *Scientific American*, June 1995 issue, "From Complexity to Perplexity", by J. Horgan, page 104]

> "Artificial Life — and the entire field of complexity—seems to be based on a seductive syllogism: There are simple sets of mathematical rules that when followed by a computer give rise to extremely complicated patterns. The world also contains many extremely complicated patterns. Conclusion: Simple rules underlie many extremely complicated phenomena in the world. With the help of powerful computers, scientists can root those rules out." [J. Horgan, *Scientific American*, June 1995 issue, "From Complexity to Perplexity", page 107]

> "Artificial Life is basically a fact-free science". [John Maynard Smith as quoted in *Scientific American*, June 1995 issue, "From Complexity to Perplexity", by J. Horgan, page 107]

The problem is that Artificial Life must be compared to something, otherwise it becomes a factless manipulation of neat computer games with subjective resemblances to vague ideas of the behavior of real life. Again, we are faced with many possible emergent types of complex behaviors, this time formal, but what kinds of these behaviors can be classified as "life-as-could-be"?, what is the formal threshold of complexity needed? In the natural world we are, more or less, able to distinguish life from non-life, biology from physics, in the logical realm, we likewise need a formal criteria to distinguish logical life from logical non-life, artificial life from artificial physics. Only by establishing an artificial physics, from which an artificial biology can emerge, and a theory, or set of rules, distinguishing the two, can we aim at a proper science based on fact.

> "Artificial Life must be compared with a real or an artificial nonliving world. Life in an artificial world requires exploring what we mean by an alternative physical or mathematical reality." [Pattee, 1995]

The two goals of AL are usually described as *hard* and *soft AL* respectively. The first concerns the synthesization of artificial life from computational or material (situated robotics) components. The second is interested in obtaining life-like behavior and is essentially metaphorical. To be accepted as a scientific field, AL, more than imitating subjective behavior, should be concerned with the investigation of the rules that allow us to distinguish life from non-life and which can be experimentally replicated within a scientific discourse. Whether we are interested in hard or soft AL, our artifacts and models should always make explicit the set of rules which allow us to defend that some artificial organization is alive or observes some specific life-like behavior. Naturally, the requirements for hard AL are much stricter, as we are not merely interested in behavioral thresholds that can be compared to real biological systems with looser or stricter rules, but the actual realization of an artificial organization that must be agreed to be living in all of its aspects. Soft AL, may restrict itself to particular behavioral traits which need only to be simulated to a satisfactory degree. We will be looking into several alternatives for organizational requirements of life during the remaining of this course.

2.2. Simulations, Realizations, Systemhood, Thinghood, and Theories of Life

> "Boids are *not* birds; they are not even remotely like birds; they have no cohesive physical structure, but rather exist as information structures — processes — within a computer. But — and this is the critical 'but'— at the level of behaviors, *flocking Boids and flocking birds are two instances of the same phenomenon*: flocking." [Langton, 1989, page 32]
>
> "The 'artificial' in Artificial Life refers to the component parts, not the emergent processes. If the component parts are implemented correctly, the processes they support are genuine — every bit as genuine as the natural processes they imitate. [...] Artificial Life will therefore be genuine life —it will simply be made of different stuff than the life that has evolved on Earth." [Langton, 1989, page 33]
>
> "Simulations and realizations belong to different categories of modeling. Simulations are metaphorical models that symbolically 'stand for' something else. Realizations are literal, material models that implement functions. Therefore, accuracy in a simulation need have no relation to quality of function in a realization. Secondly, the criteria for good simulations and realizations of a system depend on our theory of the system. The criteria for good theories depend on more than mimicry, e.g., Turing Tests." [Pattee, 1989, page 63]

The bottom line is that a simulation, no matter how good it is, will never become a realization. Nonetheless, it may still be possible to obtain artificially living organisms (realizations) if, from a simulated environment, we are able to create genuine emergent evolution. Howard Pattee [1989] has proposed that if emergent artificial organisms are able to perform measurements, or in other words, categorize and control aspects of their (artificial) environment then they may be considered realizations. Some claim that computational environments do not allow for this creative form of emergence [see Cariani, 1992; Moreno, et all, 1994]. In any case, whatever artificial environment we may use, computational or material, we need a theory allowing us to distinguish life from non-life.

Related to this issue, and in the context of systems science, is the search of those properties of the world which can be abstracted from their specific material substrate: systemhood from thinghood. Systems science is concerned with the study of *systemhood* properties, but there may be systems from which systemhood cannot be completely abstracted from *thinghood*. Life, and complexity in general, is sometimes proposed as one of those systems [see Rosen, 1986, 1991; Moreno et al, 1994; Pattee, 1995]. The difficulty for systems science, or complexity theory, lies precisely in the choice of the appropriate level of abstraction. If we abstract enough, most things will look alike, leading to a theory of factless, reminiscent analogies, exposed by Cowan and Maynard-Smith above. If we, on the other hand, abstract too little, all fields of inquiry tend to fall into more and more specific niches with little communication amongst them. In the context of life, we do not want to be tied uniquely to carbon-based life, or *life-as-we-know-it*, but we also do not want *life-as-could-be* to be anything at all. The challenge lies precisely on finding the right amounts of systemhood and thinghood, as well as the interactions between the two, necessary for a good theory of life, real or artificial.

In philosophy of biology, this problem is posed between structuralism and functionalism. Structuralism often leads into a reductionist explanation of meaningless, to a higher level, physical interactions, while functionalism may disregard important characteristics of matter. A given function may be implemented in many ways, however, not all of these structures will have the same evolutionary potential, which may even be independent of how well the function is implemented. Life is ultimately linked to some physics from which it emerges: this physics may be artificial, but the living organisms which may inhabitate it will survive depending on how well they are able to harness their environment and not only on the abstract function they may implement. Thus to study a particular life form, we cannot disregard neither its particular physical environment nor its function.

To systematically study life should concern not solely matter or form alone, but precisely the interactions between the two.

In the next lectures we will study several formal mechanisms which observe certain characteristics associated with life. With all of them we must ask the following question: "granted, life follows these characteristics, but are there other characteristics not captured in this model? Are those important? Must a theory of life include them, or are we restricting life too much?" Somewhere between "any animation goes" and "carbon-chauvinism" we must find a satisfying criteria for recognizing life from non-life. Like with anything else, the chosen theory, or theories, will be the one capable of satisfying some group's consensus, after which facts can be built upon. It should be noted that scientific consensus is not merely built upon conversational interactions, but also on the replicability of experiments available to all observers, which allows the establishment of a larger and larger consensual understanding of the world. It also binds scientific discourse to the laws of a particular world (natural or artificial), as opposed to pure linguistic interaction between participants in a given discourse.

Further Readings and References

Cariani, P. [1992], "Emergence and Artificial Life" In *Artificial Life II*. C. Langton (Ed.). Addison-Wesley. pp. 775-797.

Langton, C. [1989], "Artificial Life" In *Artificial Life*. C. Langton (Ed.). Addison-Wesley. pp. 1-47.

Klir, G. [1991], *Facets of Systems Science*. Plenum Press. (On Constructivism pp. 12-13)

Maturana, H. And F. Varela [1987], *The Tree of Knowledge*. Shambhala Publications.

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, H. [1989], "Simulations, Realizations, and Theories of Life". In *Artificial Life*. C. Langton (Ed.). Addison-Wesley. pp. 63-77.

Pattee, H. [1995], "Artificial Life needs a real Epistemology". In Advances in Artificial Life. F. Moran, A Moreno, J.J. Merelo, P. Chacon (Eds.). Springer-Verlag. (In press)

Rosen, R. [1986], "Some Comments on Systems and System Theory". In *Int. Journal of General Systems*. Vol. 13, No.1.

Rosen, Robert [1991]. *Life Itself: A Comprehensive Inquiry into the Nature, Origin, and Fabrication of Life.* Columbia University Press.

Sober, E. [1992], "Learning from Functionalism — Prospects for Strong Artificial Life". In *Artificial Life II*. C. Langton (Ed.). Addison-Wesley. pp. 749-765.

Zeleny, M., G. Klir, and K. Hufford [1989], "Precipitation Membranes, Osmotic Growths and Sythetic Biology". In *Artificial Life*. C. Langton (Ed.). Addison-Wesley. pp. 125-139.

For next lecture read:

Chapter I of Emmeche's [1991], *The Garden in the Machine: The Emerging Science of Artificial Life*. Princeton University Press.

3. Self-Organization and Emergent Complex Behavior

> “[...] a self-organizing system is a system that tends to improve its performance in the course of time by making its elements better organized for achieving the goal. This formulations includes the special case in which the goal is to achieve a high degree of organization (order) of relevant entities from low degree of their organization (disorder, chaos).” [Klir, 1991, page 156]

We should start the study of self-organization and complex behavior with the thought that “Complexity exists, in some murky sense, in the eye of the beholder” [Horgan, page 106]. George Klir [1991], in line with Ashby [1962] thinks of self-organizing systems as a case of goal-oriented systems “that are capable, with no explicit outside help, of improving their performance while pursuing their goals” [Klir, 1991, page 165] and which must be evaluated with some performance function. However, this goal (e.g. order, complex behavior) is established in relation to some observer interested in a particular behavior. This is the reason why the study of emergent, interesting, complex behavior is a tricky affair. We can all agree on the simple local rules causing the emergent global behavior, but the latter is a more subjective affair since it is not explicitly programmed or described in physical terms. It is instead an observed behavior, relevant for some observer with some goal like understanding life or building some sort of computational engine [see Forrest, 1991]. Kauffman[1993] calls for a new statistical mechanics to understand the behavior of self-organizing computational structures, in other words, new higher-level parameters (with physical analogues such as temperature) must be developed to understand self-organizing, emergent, behavior regarding some observer’s interest.

> “Complexologists often employ ‘interesting’ as a synonym for ‘complex’. But what government agency would supply funds for research on a ‘unified theory of interesting things’?” [Horgan, 1995, page 106]

What is usually referred to as self-organizing behavior is the spontaneous formation of well organized structures, patterns, or behaviors, from random initial conditions. The systems used to study this behavior are referred to as dynamical systems: state-determined systems. They possess a large number of elements or variables, and thus very large state spaces. However, when started with some initial conditions they tend to converge to small areas of this space (attractor basins) which can be interpreted as a form of self-organization. The existence of attractors is identified with the dissipation of some form of energy (friction), therefore, like living systems, dissipative structures can only be maintained by a constant flux of energy through them, and are therefore not in equilibrium. These attractors may be chaotic in which case the emergent behavior becomes too disorganized to grasp (disorganized complexity), though still self-organizing since chaotic attractors tend to be restricted to small volumes of their state-space (e.g. chaotic in a small subset of dimensions of the stat-space). The behavior of interest is often found in the transition between order and chaos — edge of chaos— and classified as a kind of organized complexity [Weaver, 1948]. Another parallel to living systems here is that such dynamical structures are not devised to exhibit this behavior, they develop spontaneously from random initial conditions (note: not from special initial conditions). This behavior — many parts working together to achieve a higher order — is also known as *synergetics* [Haken, 1977].

Since such formal dynamical systems are usually used to model real dynamical systems such as chemical networks of reactions, non-equilibrium thermodynamic behavior [Nicolis and Prigogine, 1977], or even mineral osmotic growths [Leduc, 1911; Zeleny, Klir, and Hufford, 1989], the conclusion is that in nature, there is a tendency for spontaneous self-organization which is therefore universal [Kauffman, 1993]. Further, only matter out of equilibrium (with dissipation) can achieve self-organization, which may be quite complex (strange attractors, etc.) [Prigogine, 1985]. This undeniable tendency for the spontaneous formation of complex physical patterns, is also frequently used to propose that life (an autonomous dissipative organization maintained through metabolism) is more general than usually accepted and that even mineral structures can be in this sense alive [Zeleny, Klir, and Hufford, 1989]. This process of self-organization is also often interpreted as the evolution of order from chaos. However, notice that this evolution is limited in its complexity level to the attractors dynamic systems converge to. A given dynamic system, unless its parameters are changed (structural perturbation), cannot escape its own attractor landscape and it is therefore constrained in its evolutionary potential. This limitation will become more apparent when we approach the problem of self-replication.

3.1. Edge of Chaos

Another interesting aspect of the behavior of dynamical systems concerns the concept of *bifurcation*. When the parameters of a dynamic system are changed gradually its trajectories and attractors typically change gradually, however, for certain parameter values sudden changes in the dynamic behavior can occur (e.g. Benard Cells). It is at this critical point that complicated spatio-temporal organization may occur (e.g. oscillation with constant period). Close to bifurcations the system also becomes increasingly more sensitive to parameter and initial condition changes. It is often proposed that bifurcations offer a selection mechanism [Prigogine, 1985] due to this sensitivity, as organizations may respond very differently to very small changes in their parameters, e.g. a flower's decision to bloom.

However, if the parameter space is divided by many bifurcations, the system becomes increasingly sensitive to initial conditions and small parameter changes; in this sense its behavior becomes chaotic. It is usually argued that the most useful behavior lies instead in between full order and chaos. Langton [1990, 1992] has shown (for two-dimensional cellular automata) that it is in this range of behavior that dynamical systems can carry the most complicated computations. Computation here is used an a loose sense, and means that information exchange between elements of these systems is maximized in this range, or in other words, Langton showed that the highest degrees of correlation between the states of his cellular automata occur at this stage. The same idea has been proposed by others, including Prigogine who "speaking anthropomorphically [proposed that], matter at equilibrium is 'blind,' it only 'sees' at very short distances, while matter out of equilibrium develops a sensitivity to the outside world that is a sensitivity to distant events." [Prigogine, 1985, page 484].

Kauffman [1993, page 232] further hypothesizes that "living systems exist in the [ordered] regime near the edge of chaos, and natural selection achieves and sustains such a poised state". This hypothesis is based on Packard's [1988] work showing that when natural selection algorithms are applied to dynamic systems, with the goal of achieving higher discriminative power, the parameters are changed generally to lead these systems into this transitional area between order and chaos. This idea is very intuitive, since chaotic dynamical systems are too sensitive to parameter changes (structural perturbation), that is, a single mutation leads the system into another completely different behavior (sensitive to damage). By contrast, ordered systems are more resilient to damage, and a small parameter change will usually result in a small behavior change which is ideal for smooth adaptation (hill-climbing) in correlated fitness landscapes. However, even though very ordered systems can adapt by accumulation of useful successful variations (because damage does not propagate widely), they may not be able 'step out' of certain areas of their fitness landscapes. It is here that systems at the edge of chaos enter the scene; they are not as sensitive to damage as chaotic systems, but still they are more sensitive than fully ordered systems, and thus, some mutations will accumulate (by causing minor structural changes) and some others will cause major changes in the dynamics allowing more distant searches in fitness spaces. These characteristics of simultaneous mutation buffering (to small changes) and dramatic alteration of behavior (in response to larger changes) is ideal for evolvability [Conrad, 1983, 1990].

Perhaps these concepts can be better grasped in the context of classifier networks. As stressed in the first lecture, the ability to discriminate (categorize) relevant events in an environment is an important characteristic of life. Dynamic systems such as boolean networks have the ability to discriminate inputs. Generally, the attractors of their dynamics are used to represent events in their environments: depending on inputs, the network will converge to different attractors. However, for any classification to have survival value, it must relate its own constructed states (attractors) to relevant events in its environment, thus, similar events in the world should correspond to the same attractor basin. Chaotic systems clearly do not have this property due to their sensitivity to initial conditions. Ordered systems follow this basic heuristic. If on the edge of chaos, ordered systems may in addition allow for higher information exchange and perhaps more 'clever' (evolvable) categorization mechanism.

> " Organisms and other entities which interact with their worlds are likely to couple to those worlds in such a way that smooth classification occurs, and the world is seen as relatively stable. Then the 'knower' should not be chaotic, nor should its classification, the 'known', be. It is a reasonable guess that both the knowing system and the known world are in the [ordered] regime, perhaps near the edge of chaos." [Kauffman, 1993, page 234]

3.2 G-Type/P-Type Distinction and Emergent Behavior

> "In the context of Artificial Life, we need to generalize the notions of genotype and phenotype, so that we may apply them in non-biological situations. [...] The GTYPE, essentially, is the specification for a set of machines, while the PTYPE is the behavior that results as the machines interact with one another in the context of a specific environment. This is the bottom-up approach to the generation of behavior." [Langton, 1989, pp. 22-23]

Langton's definition of GTYPE and PTYPE is not so much a generalization of the genotype/phenotype distinction in biology, as it is a framework to conceptualize emergent behavior in Artificial Life. It states the requirement of a minimum of two levels of description for models of emergence. The first specifically describes the level of the rules of dynamics (e.g. laws of physics or cellular automata rules). The second is the description of whatever global behavior one decides to observe (e.g. self-organization, function, patterns, etc.). However, this distinction fails to generalize the biological genotype/phenotype distinction, since the genotype does not define the laws that allow the phenotype to self-organize (protein folding), those are simply the chemical laws of the constituents of proteins (aminoacid chains). The genotype merely offers the initial conditions for such a process of self-organization, in this sense it can be seen more as data than as a program for a phenotype [Atlan and Koppel, 1990]. This problem can be easily recognized when we realize that the biological genotype, depending on the level of description chosen, can be seen both as a GTYPE and as a PTYPE. The GTYPE of the genotype being the chemical rules of interaction between the components of nucleic acids, and the PTYPE being its self-organization into DNA strands and its subsequent utilization as a genetic information carrier. Likewise, the biological phenotype may also be granted a GTYPE and PTYPE description. The former being the rules of interaction of protein constituents, and the latter being the functions associated with the specific phenotype (e.g. catalytic behavior).

Historical papers and books in self-organization include: Farley and Clark [1954], Yovits and Cameron [1960], von Foerster and Zopf [1962], Yovits, Jacobi, and Goldstein [1962], Ashby [1962], Nicolis and Prigogine [1977], and more recently Kauffman [1993].

Further Readings and References:

Ashby, W.R. [1960], Design for a Brain. Wiley.

Ashby, W.R. [1962]."Principles of the self-organizing system." In: *Principles of Self-Organization*. H. von Foerster and G.W. Zopf (Eds.). Pergamon Press. Reprinted in Klir [1991 pp. 521-536.

Atlan, H. And M. Koppel [1990]. "The cellular computer DNA: program or data". In: *Bulletin of Mathematical Biology*, Vol. 52, No. 3, pp.335-348.

Conrad. M. [1983], *Adpatability*. Plenum Press.

Conrad, M. [1990], "The geometry of evolutions". In *BioSystems* Vol. 24, pp. 61-81.

Farley, B.G. and W.A. Clark [1954], "Simulations of self-organizing systems by digital computer. IRE Transactions, IT-Vol 4 No. 3, pp 76-84.

Forrest, S. (Ed.) [1990]. *Emergent Computation.* MIT Press/ North-Holland. Special Issue of *Physica D*. Vol. 42.

Haken, H. [1977], *Synergetics*. Springer-Verlag.

Horgan J. [1995], *Scientific American*, June 1995 issue, "From Complexity to Perplexity"

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

Klir, George J. [1991]. *Facets of Systems Science.* Plenum Press.

Langton, C. [1989], "Artificial Life" In *Artificial Life*. C. Langton (Ed.). Addison-Wesley. pp. 1-47.

Langton, C. [1990]. "Computation at the edge of chaos: phase transitions and emergent computation". In Forrest [1990] pp12-37.

Langton, C. [1992], "Life at the edge of chaos". In *Artificial Life II*. C. Langton (Ed.) Pp 41-91. Addison-Wesley

Leduc, S. [1911], *The Mechanisms of Life*. Rebman.

Nicolis, G. and I. Prigogine [1977]. *Self-Organization in Nonequilibrium Systems: From Dissipative Structures to Order through Fluctuations*. Wiley-Interscience.

Packard, N. [1988] "Adaptation toward the edge of chaos". In. *Complexity in Biologic Modelling*. S. Kelso and M. Shlesinger.

Prigogine, I. [1985]. "New Perspectives on Complexity". In the Sciences and Praxis of Complexity. The United Nations University pp.107-118. Reprinted in Klir [1991] pp.484-492

von Foerster, H. and G.W. Zopf (Eds.) [1962], *Principles of Self-Organization*. Pergamon Press.

Weaver,W. [1948], "Science and Complexity", In *American Scientist*, Vol. 36. Reprinted in Klir [1991] pp. 449-456.

Yovits, M.C. and S. Cameron (Eds.) [1960], *Self-Organizing Systems*. Pergamon Press.

Yovits, M.C., G.T. Jacobi, and G.D. Goldstein (Eds.) [1962], *Self-Organizing Systems — 1962.* Spartan Books.

Zeleny, M., G. Klir, and K. Hufford [1989], "Precipitation Membranes, Osmotic Growths and Sythetic Biology". In *Artificial Life*. C. Langton (Ed.). Addison-Wesley. pp. 125-139.

For Next Lecture Read:

Chapter II of Emmeche's [1991], *The Garden in the Machine: The Emerging Science of Artificial Life*. Princeton University Press.

4. Self - Organization and Cellular Automata

Generally, 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, with no external interference. All the mechanisms dictating its behavior are internal to the system: self-organization as opposed to externally imposed organization. Thus, it is reasonable to further demand that for a system to observe self-organizing behavior, its order cannot be imposed by special initial conditions, which would amount to a special creation. Therefore, to guarantee that a system is self-organizing, we start it with random initial conditions and see if it attains the desired order, or attractor behavior.

So far, we have seen two different types of computational systems said to be self-organizing in this sense: the discrete logistic equation and Kauffman's NK-Boolean nets. We have seen that the former observes several ranges of ordered behavior according to its parameter a. For a # 3, the system will converge to a single point steady state (independently of its initial value). For 3 # a # 4 the system enters a series of bifurcations, meaning that it starts converging to a two-state limit cycle which progressively doubles the number of states in its cycle as a increases. Close to $a = 4$, the cycles become chaotic. That is, in the chaotic range, the slightest change in the initial value, will lead to a completely different trajectory (though similarly chaotic). The system goes from being independent to strongly dependent of initial conditions, though, in each range, the attractor behavior of the equation is the same for random initial conditions. Thus, we can see the logistic equation as self-organizing.

But there is another aspect of the logistic equation that should be understood. In all of its ranges of behavior, from full order to full chaos, the system is (fairly) reversible. That is, I can always obtain a specific initial condition which caused some behavior, by formally running the system backwards. This means the system is deterministic in both temporal directions. Formally, this means the state transition function is invertible. (This is actually only true, if we decide to work on the lower half of its state space, since the logistic equation is a quadratic function, it has always two possible solutions for the previous value of the current state, these values are symmetric about the middle point of its state space). Some, Howard Pattee for instance, resist calling this kind of reversible systems self-organizing. They reason that if a system is self-organizing, when ran backwards it should be self-*dis*organizing, that is, it should lead to random initial conditions, or to an incomplete knowledge of possible initial states. Pattee reserves the term self-organization to those irreversible systems whose behaviors must be evaluated statistically. The logistic map shows "hints" of this backwards self-disorganization, but we can still work out effectively its backwards trajectory to an initial condition by restricting the quadratic solutions to half of its state space.

Random Boolean Networks are much more complicated than this. They are completely deterministic since a certain state will always lead to the same next state (state-determinacy), however, we cannot usually know exactly what the predecessor of a current state was. Systems like this are usually studied with statistical tools. Even though the rules that dictate the next state of its components are simple and deterministic, the overall behavior of the system is generally too complicated to predict and statistical analysis has to be performed. For instance, Kauffman has shown that when K=2 (number of inputs to each node), his networks will have on average $\sqrt{N}$ basins of attraction with a length of $\sqrt{N}$ states; if the output of one node is switched to the other boolean value (perturbation), the trajectory returns to that cycle 85% of the time, while on the remaining 15% of the time it will "jump" into a different basin of attraction.

Cellular automata (CA) fall into this same category of deterministic, irreversible, self-organization. We will discuss Wolfram's four statistical classes reached by all one-dimensional CA from random initial conditions, Langton's further refinement of these classes, and Conway's game of Life. Like the NK-networks, CA self-organize exclusively in accordance to their local rules. This is usually interpreted in boolean networks as the simulation of some closed abstract dynamics (e.g. chemical reactions, genomic networks with epistasis, etc), but in CA the local rules are often viewed as the simulation of some artificial physics in an artificial topological space, while the patterns of cellular activation (state cycles) are seen as emergent phenomena. In particular, when coherent patterns are observed which behave like life in some way (motion, self-reproduction, etc), it is often argued that it represents the emergence of artificial life from artificial matter. The simple local rules stand for an artificial physics with *micro-causality* (a cell's state is solely dependent on its spatial neighbors and on their previous value), and the emergent patterns for artificial life. One obvious problem with this interpretation is that in the real world the local rules leading to some physical causality generate all that we see around us, living and non-living. In the CA world, arguments are often made for the emergence of artificial life, but not for an explicit criteria to distinguish artificial life from artificial non-life also generated by the artificial physics.

CA further observe the three ranges of behavior exhibited by boolean networks: ordered, chaotic, and intermediate. We will discuss Langton's results indicating that CA will perform computations more effectively in the edge of chaos, which is based on the definition of his ? parameter. Later on, once evolutionary algorithms are introduced, the evolution of CA rules for the solution of non-trivial tasks is discussed.

Further Readings and References:

Forrest, S. (Ed.) [1990]. *Emergent Computation.* MIT Press/ North-Holland. Special Issue of *Physica D*. Vol. 42.

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

Langton, C. [1990]. "Computation at the edge of chaos: phase transitions and emergent computation". In Forrest [1990] pp12-37.

Packard, N. [1988]. "Adaptation to the edge of chaos". In *Complexity in Biological Modeling*, S.Kelso and M. Shlesinger (eds.).

Rucker R. And J. Walker [1989]. *CA Lab: Rudy Rucker's Cellular Automata Laboratory*. Autodesk, 1989 (book and software)

Sigmund, K. [1993], *Games of Life: Explorations in Ecology, Evolution, and Behavior*. Oxford University Press

Toffoli T. and N. Margolus [1987]. *Cellular Automata Machines*. MIT Press

Wolfram, S. (Ed.) [1986], *Theory and Applications of Cellular Automata*, World Scientific Press.

For Next Lecture Read:

Chapter III of Emmeche's [1991], *The Garden in the Machine: The Emerging Science of Artificial Life*. Princeton University Press.

5. Von Neumann and Natural Selection

5.1 Von Neumann's Self-Reproduction Scheme

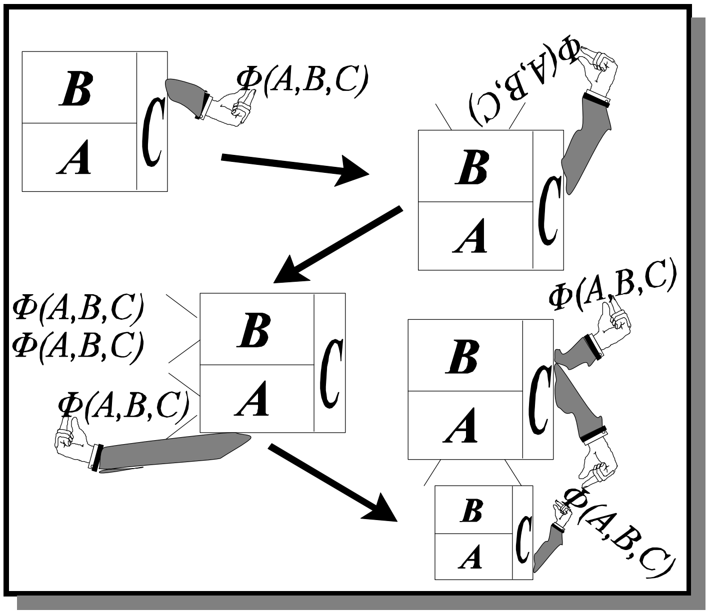

Von Neumann thought of his logical model of self-reproduction as an answer to the observation that, unlike machines, biological organisms have an ability to self-replicate while increasing their complexity without limit. Mechanical artefacts are instead produced via more complicated factories (as opposed to self-production) and can only degenerate in their complexity. He was searching for a complexity threshold beyond which systems may self-reproduce (no outside control) while possibly increasing their complexity.

Von Neumann concluded that this threshold entails a memory-stored description F (*X*) that can be interpreted by a universal constructor automaton *A* to produce any automaton *X*; if a description of *A*, F (*A*), is fed to *A* itself, then a new copy of *A* is obtained. However, to avoid a logical paradox of self-reference, the description, which cannot describe itself, must be both copied and translated into the described automaton. This way, in addition to the universal constructor, an automaton *B* capable of copying any description, F (*X*), is included in the self-replication scheme. A third automaton *C* is also included to effect all the manipulation of descriptions necessary — a sort of operative system. To sum it up, the self-replicating system contains the set of automata (*A* + *B* + *C*) and a description F (*A* + *B* + *C*); the description is fed to B which copies it three times (assuming destruction of the original); one of these copies is then fed to *A* which produces another automaton (*A* + *B* + *C*); the second copy is then handled separately to the new automaton which together with this description is also able to self-reproduce; the third copy is kept so that the self-reproducing capability may be maintained (it is also assumed that *A* destroys utilized descriptions). Notice that the description, or program, is used in two different ways: it is both *translated* and *copied*. In the first role, it controls the construction of an automaton by causing a sequence of activities (*active* role of description). In the second role, it is simply copied (*passive* role of description). In other words, the *interpreted* description controls construction, and the *uninterpreted* description is copied separately, passing along its stored information (memory) to the next generation.

The notion of description-based self-reproduction implies a self-referential linguistic mechanism. A description must be cast on some symbol system while it must also be implemented by some physical structure. Since many realizations of the same symbol system are possible, viewing descriptions only as physical systems cannot explain the symbolic nature of the control of construction. When *A* interprets a description to construct some automaton, a *semantic* code is utilized to map instructions into construction commands to be performed. When *B* copies a description, only its *syntactic* aspects are replicated. Now, the language of this semantic code presupposes a set of primitives (e.g. parts and processes) for which the instructions are said to "stand for". Descriptions are not universal insofar as they refer to these building blocks which cannot be changed without altering the significance of the descriptions. The building blocks ultimately produce the dynamics, behavior, and/or functionality of the overall system, and may be material or computational (standing for some artificial materiality). In the genetic system, the parts are amino acids. Computational parts might be for example the building blocks of neural networks coded by genetic algorithms and L-Systems. We can see that a self-reproducing organism following this scheme is an entanglement of symbolic controls and dynamic constraints which is closed on its semantics. Howard Pattee calls such a principle of self-organization *semantic closure*.

Semantic closure requires the parts problem discussed above to be explicitly taken into account. The evolvability of a self-reproducing system, to be discussed below, is dependent on the dynamic parts used by the semantic code. If the parts are very simple, then the descriptions will have to be very complicated, whereas if the parts possess rich dynamic properties, the descriptions can be simpler since they will take for granted a lot of the dynamics that otherwise would have to be specified. In the genetic system, genes do not have to specify the functional characteristics of the proteins produced, but simply the string of amino acids that will produce that functionality "for free" [Moreno et al, 1994]. Furthermore, there is a trade-off between programmability and evolvability [Conrad, 1983, 1990] which renders some self-reproducing systems no evolutionary potential whatsoever. If descriptions require high programmability they will be very sensitive to damage (e.g. Langton's self-reproducing loops). Low programmability grants self-reproducing systems the ability to change without destroying their own organization, though it also reduces the space of possible evolvable configurations. In computational realms this implies that we should move towards models that include both programmable and self-organizing components [Rocha, 1997]. We will discuss such systems later in this course.

5.2 Open-ended emergent evolution and natural selection

Perhaps the most important consequence of the requirement of memory-based descriptions in Von Neumann's self-reproduction scheme is its opening the possibility for open-ended emergent evolution. As Von Neumann [1966] discussed, if the description of the self-reproducing automata is changed (mutated), in a way as to not affect the basic functioning of $(A + B + C)$ — that is, if the semantic closure in not destroyed — then, the new automaton $(A + B + C)$` will be slightly different from its parent. Von Neumann used a new automaton $D$ to be included in the self-replicating organism, whose function does not disturb the basic performance of $(A + B + C)$; if there is a mutation in the $D$ part of the description, say $D$`, then the system $(A + B + C + D)$ + F $(A + B + C + D$`) will produce $(A + B + C + D$`) + F $(A + B + C + D$`). Von Neumann [1966, page 86] further proposed that non-trivial self-reproduction should include this "ability to undergo inheritable mutations as well as the ability to make another organism like the original", to distinguish it from "naive" self-reproduction like growing crystals.

Notice that changes in $(A + B + C + D$`) are not heritable, only changes in the description, F $(A + B + C + D$`), are inherited by the automaton's offspring and are thus relevant for evolution. This ability to transmit mutations is precisely at the core of the principle of natural selection of modern Darwinism. Through variation (mutation) populations of different organisms are produced; the statistical bias these mutations impose on reproduction rates of organisms will create survival differentials (fitness) on the population which define natural selection. In principle, if the language of description is rich enough (its material constraints are dynamically rich), an endless variety of organisms can be evolved: open-ended emergent evolution.

The threshold of complexity proposed by Von Neumann is taken by some (e.g. Pattee, Rosen, Cariani, Kampis see Rocha[1995]) as another category of self-organization which is capable of creative organization and selection from outside. Notice that the self-organizing systems we have been studying so far (random nets, Cellular Automata) are said to self-organize when they converge to small areas of their state spaces or attractors. This sort of evolution is constrained by the complexity of the attractor landscape of organisms seen as dynamical systems. It cannot evolve a truly novel dynamics. In Cariani's [1991] terms, it cannot evolve new functionalities (such as sensors). Non-trivial self-replicating systems rely instead on *memory-based selected self-organization* [Rocha, 1996, 1996; Henry and Rocha, 1996] which can be seen as a type complex adaptive systems that observe a principle of organization referred to as *embodied, evolving semiosis* [Rocha, 1997].

Further Readings and References:

Cariani, Peter [1991]."Some Epistemological Implications of Devices which Construct their own Sensors and Effectors." In: *Towards a Practice of Autonomous Systems,* F. Varela and P. Bourgine (Eds.). MIT Press. pp 484-493.

Conrad. M. [1983], *Adpatability*. Plenum Press.

Conrad, M. [1990], "The geometry of evolutions". In *BioSystems* Vol. 24, pp. 61-81.

Henry, C. and L.M.Rocha [1996]."Language theory: consensual selection of dynamics." In: *Cybernetics and Systems'96*. R.Trappl (Ed.).Austrian Society for Cybernetic Studies. pp. 477-482. To be reprinted in *Cybernetics and Systems*.

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." In: *Cognition and Brain Theory* Vol. 5, no. 4, pp 325-341.

Pollack, R. [1994]. *Signs of Life: The Language and Meanings of DNA.* Houghton Mifflin.

Rocha, Luis M.(Ed.) [1995]. *special issue in Self-Reference in Biological and Cognitive Systems*. *Communication and Cognition - AI* Vol. 12, nos. 1-2 .

Rocha, Luis M. [1996]."Selected self-organization: from eigenbehavior to descriptions." In: *Proceedings of the Intern. Seminar on Evolving Systems, Vienna 1995*. S. Salthe and G. Van de Vijver (Eds.). MIT Press. (in press).

Rocha, Luis M. [1996]."Eigenbehavior and symbols." In: *Systems Research* Vol. 12, No. 3 (In Press).

Rocha, Luis M. [1997]. *Evidence Sets and Contextual Genetic Algorithms: Exploring Uncertainty, Context, and Embodiment in Cognitive and Biological Systems*. PhD. Dissertation. SUNY Binghamton.

von Neumann, John [1966]. *The Theory of Self-Reproducing Automata.* Arthur Burks (Ed.) University of Illinois Press.

For Next Lecture Read:

Chapter V of Emmeche's [1991], *The Garden in the Machine: The Emerging Science of Artificial Life*. Princeton University Press.

6. Evolutionary Computation

> "How does evolution produce increasingly fit organisms in environments which are highly uncertain for individual organisms? How does an organism use its experience to modify its behavior in beneficial ways (i.e. how does it learn or 'adapt under sensory guidance')? How can computers be programmed so that problem-solving capabilities are built up by specifying '*what* is to be done' rather than '*how* to do it'?" [Holland, 1975, page 1]

These were some of the questions concerning John Holland when he thought of Genetic Algorithms (GA's) in the 1960's. Basically, all these problems were shown to be reduced to a problem of optimizing a multiparameter function necessary for solving a particular problem. Nature's "problem" is to create organisms that reproduce more (are more fit) in a particular environment: the environment dictates the selective pressures, and the solutions to these pressures are organisms themselves. In the language of optimization, the solutions to a particular problem (say, an engineering problem), will be selected according to how well they solve that problem. GA's are inspired by natural selection as the solutions to our problem are not algebraically calculated, but rather found by a population of solution alternatives which is altered in each time step of the algorithm in order to increase the probability of having better solutions in the population. In other words, GA's (or other Evolutionary Strategies (ES) such as Evolutionary Programming (EP)), explore the multi-parameter space of solution alternatives for a particular problem, by means of a population of encoded strings (standing for alternatives) which undergo variation (crossover and mutation) and are reproduced in a way as to lead the population to ever more promising regions of this search space (selection).

6.1. Evolutionary Strategies and Self-Organization

The underlying idea of computational ES is the separation of solutions for a particular problem (e.g. a machine) from *descriptions* of those solutions (memory). GA's work on these descriptions and not on the solutions themselves, that is, variation is applied to descriptions, while the respective solutions are evaluated, and the whole (description-solution) selected according to this evaluation. Such machine/description separation follows aspects of von Neumann's self-reproducing scheme which is able to increase the complexity of the machines described. However, the form of organization attained by GA's is not self-organizing in the sense of a boolean network of cellular automata. Even though the solutions are obtained from the interaction of a population of elements, and in this sense following the general rules usually observed by computationally emergent systems (e.g. Langton [1988], Mitchell [1992]), they do not ***self***-organize since they rely on the selective pressures of some fitness function. The order so attained is not a result of the internal dynamics of a collection of interacting elements (like a random net), but is instead dictated by the *external* selection criteria. In this sense, ES follow a memory-based selective organization scheme.

Self-organization is instead equated with those behaviors of organizations that are unavoidable and solely dependent on their internal dynamics. This is usually thought of in terms of the attractor behavior of state-determined dynamic systems. ES rely on different concepts: first, with the description-solution dichotomy the concept of memory is introduced (state-determined, self-organizing, systems are memoryless); second, the transition rules of ES are not state-determined — variation is stochastic; third, as already discussed, selection is external to the populations of descriptions. This way, we can hardly say that a population of memories is interacting with any sort of *"self-dynamics":* the solutions reached by a GA do not self-organize but are a result of stochastic (population) variation and external selection.

Systems which are able to develop the mechanisms to harness this variation based on an internally defined description-solution dichotomy may follow the kind of selective based self-organizing principle described as semantic closure in section 5. However, (computational) GA's are not closed in this sense, the coding relation is externally imposed and not evolved within the system itself. For all the reasons above it is therefore natural to think of ES as completely distinct from self-organization. It is perhaps useful to think that ES are modeling a very different aspect of biological systems that is related to natural selection. Self-organizing systems model the abstract, internal, characteristics of matter, while ES model the existence of, external, selective pressures on populations of varying memory based descriptions of some system.

6.2. development and morphogenesis: self-organization and selection come together

Many of the new developments of GA's have to do with the inclusion of a developmental stage between genotype and phenotype, in other words, the creation of some artificial morphogenesis. Basically, the idea has been to encode rules that will themselves self-organize to produce a phenotype, rather than the direct encoding of the phenotype itself. As discussed in class, these rules often use L-System grammars which dictate production system programs [Wilson, 1988] leading to some phenotype. The most important advantage of this intermediate stage, as explored by Kitano [1990], Gruau [1993], Belew [1992] and others, is the ability to code for much larger structures than a direct encoding allows. In practical terms, they have solved some of the scalability problems of encoding (e.g.) neural networks in GA's, by reducing the search space dramatically.

The L-system grammars are higher level descriptions of self-organizing developmental processes. However, these first approaches used solely context-free, state-determined, L-System grammars, compromising epistasis (or mutual, non-linear, influence of genetic descriptions amongst each other) in the simulation of self-organizing development. Dellaert and Beer [1994] and Kitano [1994], for instance, used Kauffman's boolean networks to simulate genetic epistasis and self-organization. In other words, the GA will code for rules which will construct boolean networks whose nodes stand for aspects of the phenotypes we wish to evolve on some physical simulation. In Dellaert and Beer's model, the nodes stand for cell mitosis and other characteristics. This way, the solutions of the GA are self-organizing systems whose attractor behavior dictates pre-defined phenotypic traits.

These approaches in effect offer an emergent morphology, that is, they code for rules which will themselves self-organize into some phenotype (instead of strict programming of morphology). The indirect encoding further allows the search to occur in a reduced space, amplified through development. An interesting side effect has to do with the appearance of modularity traits on the evolved phenotypes [Wagner, 1995]. In Rocha [1995, 1997] I have proposed a scheme where the contextual elements of development might be themselves evolved. Instead of boolean networks I utilize operations between fuzzy sets to simulate a material phenotype. This scheme represents a general purpose GA which searches the same search space for any size of phenotypic behaviors.

The most important aspect of these GA's with emergent morphologies is the utilization in the same model of an external selection engine (the GA) coupled to a particular self-organizing dynamics (e.g. boolean networks) standing for some materiality. Such schemes bring together, computationally, the two most important aspects of evolutionary systems: self-organization and selection. These models belong to a category of self-organization referred to as *Selected Self-Organization* which is based on local memory [Rocha, 1996a, 1996b]. Selected Self-Organization with distributed memory is also possible in autocatalytic structures, though its evolutionary potential is much smaller than the local memory kind. The reason lies in Von Neumann's notion of non-trivial self-reproduction. The introduction of symbolic descriptions allows a much more sophisticated form of communication: structures are constructed from static descriptions and do not have to reproduce through some complicated process of self-inspection. In other words, descriptions can construct any kind of structure, while self-inspection relies on only those structures that happen to be able to make copies of themselves. As an example, a non-genetic protein-based life form, would have to rely only on those proteins that could make direct copies of themselves [These issues are treated in detail in Rocha, 1996a, 1996b, 1997].

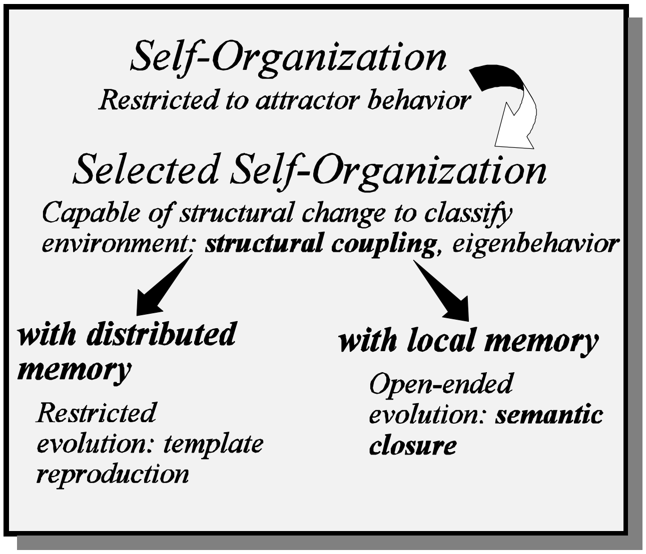

Taxonomy of Self-Organization

Further Readings and References:

Belew, R.K. [1993]."Interposing an Ontogenic Model Between Genetic Algorithms and Neural Networks." In: *Advances in neural information processing (NIPS5)*. J. Cowan (Ed.). Morgan Kaufmann.

Dellaert, F. and R.D. Beer [1994]."Toward an evolvable model of development for autonomous agent synthesis." In: *Artificial Life IV: Proceedings of the Fourth International Workshop on the Synthesis and Simulation of Living Systems*. R. Brooks and P. Maes (Eds.). MIT Press.

Goldberg, D. E. [1989]. *Genetic Algorithms in Search, Optimization, and Machine Learning.* Addison-Wesley.

Gruau, Frédéric [1993]."Genetic Sythesis of Modular Neural Networks." In: *Proceedings of the fifth international conference on Genetic Alogorithms*. S. Forrest. Morgan Kauffman, pp 318-325.

Holland, John H. [1975]. *Adaptation in Natural and Artificial Systems.* University of Michigan Press.

Kitano, H. [1990]."Designing Networks using Genetic Algorithms with Graph Generation System." In: *Complex Systems* Vol. 4, pp 461-476.

Kitano, Hiroaki [1994]."Evolution of Metabolism for Morphogenesis." In: *Artificial Life IV: proceedings of the fourth international workshop on the synthesis and simulation of living systems*. R. Brooks and P. Maes (Eds.). MIT Press.

Mitchell, M. [1992]. "Genetic algorithms". In Lectures in Complex Systems. L. Nadel and D. Stein (Eds.). SFI Studies in the Science of Complexity Vol. V, Addison-Wesley. pp 3-87.

Rocha, Luis M. [1995]."Contextual Genetic Algorithms: Evolving Developmental Rules." In: *Advances in Artificial Life*. F. Moran, A. Moreno, J.J. Merelo, and P. Chacon (Eds.). Series: Lecture Notes in Artificial Intelligence, Springer-Verlag. pp. 368-382.

Rocha, Luis M. [1996a]."Selected self-organization: from eigenbehavior to descriptions." In: *Proceedings of the Intern. Seminar on Evolving Systems, Vienna 1995*. S. Salthe and G. Van de Vijver (Eds.). MIT Press. (in press).

Rocha, Luis M. [1996b]."Eigenbehavior and symbols." In: *Systems Research* Vol. 12, No. 3 (In Press).

Rocha, Luis M. [1997]. *Evidence Sets and Contextual Genetic Algorithms: Exploring Uncertainty, Context, and Embodiment in Cognitive and Biological Systems*. PhD. Dissertation. SUNY Binghamton.

Wagner, Gunter [1995] "Adaptation and the modular design of organisms". In: *Advances in Artificial Life*. F. Moran, A. Moreno, J.J. Merelo, and P. Chacon (Eds.). Series: Lecture Notes in Artificial Intelligence, Springer-Verlag. pp. 317-328.

7. Artificial Life: Self-Organizing and Evolutionary Systems

Throughout this course emphasis was put on identifying the most important tools utilized in the field of Artificial Life. We started with *self-organizing systems*, exemplified with the logistics equation, random boolean networks, cellular automata (e.g. Conway's game of Life), and all characterized in terms of *dynamics systems theory*. Later, with the *von Neumann self-reproduction scheme*, I argued that state-determined (purely dynamic) systems are not able to offer *open-ended evolution*, that is, to increase their complexity with genuine emergence of new functionalities. Dynamic systems are restricted to the complexity of their attractor landscape.

For this purpose, systems inspired by von Neumann's scheme, which demand the separation between the description of a machine from the machine itself, and therefore introduce the concept of *memory* and external selection, were introduced. Such systems offer a model of the mechanisms utilized by natural selection, and are accordingly known as *evolutionary systems* (or evolutionary strategies) — e.g. genetic algorithms and evolutionary programming. We can also refer to the mechanisms utilized to model the kind of evolution that natural selection offers as *memory based selective strategies*: selection acting on memory elements in order to change the dynamic structure they encode.

I further emphasized hybrid systems which try to model both the self-organizing and selective mechanisms of biological systems, and can therefore offer a more complete understanding of evolutionary systems. I referred to these systems as getting close to the category of *local memory based selective self-organization*, or semantic closure. In practice I showed those approaches aiming at the introduction of non-deterministic, self-organizing, developmental steps between genotype and phenotype such as the evolution of boolean/neural networks encoded through L-System rules in a genetic algorithm. Also discussed were models capable of emergent computation by coupling genetic algorithms to cellular automata in order to have the latter solve non-trivial tasks.

The understanding of the relative importance the two basic categories of organization in artificial systems introduces a very powerful way to study the relative importance of self-organization and natural selection in biological systems themselves. In other words, by creating different forms of *life-as-it-could-be* with different degrees of both these categories, we may shed some light on the credit assignment problem of biology: how much of evolution is a result of natural selection and how much is a result of the self-organizing characteristics of its specific materiality.

I was able to introduce many of the usual applications of Artificial Life, from bugs and boids, to evolutionary robots and social evolution. Each of these applications can be a universe of investigation in itself, so emphasis was instead put on the basic categories of organization and their respective simulation tools referred above. In one way or another, all of these applications utilize in different degrees such tools described throughout the course. For instance, evolutionary robots may use a genetic algorithm to evolve a boolean network for its control system allowing it to solve some maze. To the extent that its control system was evolved and uses self-organizing mechanisms, we can say that such control system was subjected to a memory based selective type of self-organization. Naturally, the robot itself (its moving parts and sensors) were not evolved but engineered; the complete evolution of a robot through self-organization and selection represents probably the most ambitious long-term goal of Artificial Life, showing us how far behind we still are from getting there.