TR－H－033

Artificial Life Primer

Jari VAARIO

1993. 9. 29

ATR人間情報通信研究所

〒619-02 京都府相楽郡精華町光台2-2　☎07749-5-1011

ATR Human Information Processing Research Laboratories

2-2, Hikaridai, Seika-cho, Soraku-gun, Kyoto 619-02 Japan

Telephone: +81-7749-5-1011
Facsimile: +81-7749-5-1008

Artificial Life Primer

Jari Vaario
ATR Human Information Processing Research Laboratories
2-2 Hikari-dai, Seika-cho, Soraku-gun
Kyoto 619-02 JAPAN
Phone: +81-7749-5-1004
FAX: +81-7749-5-1008
E-mail: jari@hip.atr.co.jp

This paper gives a short overview of Artificial Life research. The paper answers the questions of how Artificial Life started, how far we have gone, and where we are heading. The paper emphasizes the emergence process and the dynamic embodiment with the environment as the new concepts introduced by Artificial Life. In the paper the major research works in Artificial Life are introduced by categorizing them into wetware synthesis, software synthesis, hardware synthesis and philosophical issues.

1 Introduction

It has always been a dream of man to build an artificial system that obeys his orders. A good story about this is the Golem, a story of an artificial man, and what went wrong with him (or should we say it?). Or, the old movie of Frankenstein is monster, that finally attacked its creator. A modern version of these stories is the Spielberg's "Terminator" movie, where machines have overtaken the humans — well, at least almost. Or, the very latest movie by him, "Jurassic Park" (based on the book of Michael Crichton), where the dinosaurs are rebuilt from the genetic codes that have been preserved for tens of millions of years. All these are still science fiction, but they give an idea of what Artificial Life (ALife) *could* be. However, the reality is still much, much more rudimentary, and much, much less fearsome.

These popular views of ALife should stimulate some discussions of the ultimate goal of Alife. And after all, the ultimate goal of ALife is to find an answer to the question: "What is life?". Instead of taking the analytic approach to study biological life, ALife takes a *synthetic* approach to build something that could be called life. In the case that we really succeed in this we should keep our mind on the above science fiction examples.

However, this introductory paper does not try to moralize further the goal of ALife, but instead to introduce how far we have come, and where we are heading.

In some sense ALife could be extended to cover the whole history of engineering, but we take here a narrower perspective to cover only the work where the emphasis is on life.

In the following we describe first the early works of the imitations of animals. These works gave birth to the automaton, and furthermore to the current computer technology. Now the computers have given birth to the artificial life models. This process might furthermore reinforce itself to create new computational models, that could be used to create new ALife models, and so on. Thus ALife research could be a part of the evolution of computer technology, and it could play an important role in the creation of future artifacts (whether they are computational machines, robots, or something for which we do not have words yet). This gives us a practical goal for ALife: *to find a mechanism for an evolutionary process to be used in the automatic design and creation of artifacts.*

2 Perspective on Artificial Life

Physical implementations of life can be found in the history of technology [6]. The most often mentioned example is the mechanical duck by Jacques de Vaucanson (1709-1782): "an artificial duck made of gilded copper who drinks, eats, quacks, splashes about on the water, and digest his food like a living duck."[1]

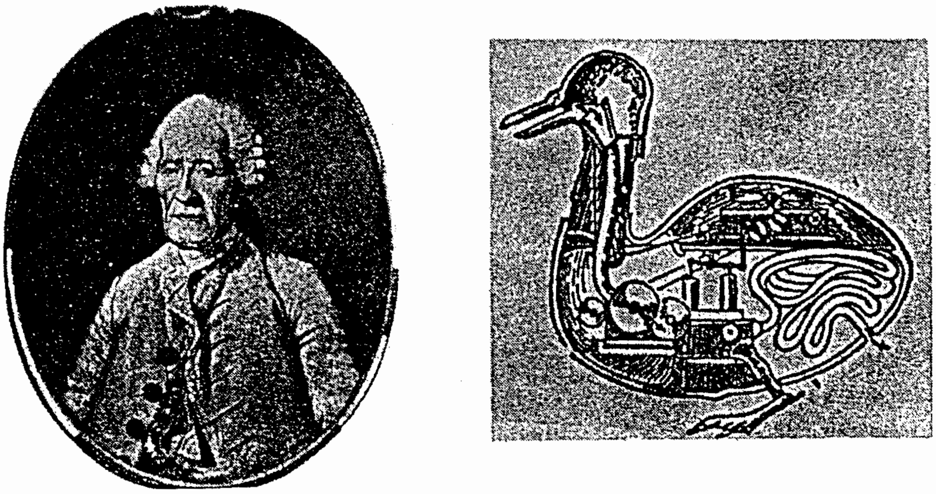

Figure 1: Jacques de Vaucanson (1709-1782) and his famous duck that was capable to quack, waddle, and had "inestial function". [11, page 189]

These mechanical constructions showed the way toward general purpose computers. A remarkable point of development of computers was to turn attention away from the *mechanism* of life to the *logic* of life. Research on the logic of life led to current computer technology and to on understanding of intelligence as logical operations.

Research into the mechanisms of life also led to the development of *cybernetics*. Credit as a founder of cybernetics is usually given to Norbert Wiener. His work led to the research field of systems theory, although some earlier history of systems theory could be given as well. This field is now experiencing a renaissance as it has became obvious that logic will not be able to explain all behavior observable in Nature. Particularly, the non-linear and chaotic behavior of systems is attracting more attention.

Todays general purpose computers provide us a powerful tool to explore the same phenomena that inspired them to be created, namely the *mechanism of life*. The result of this interest studying the mechanism of life will most likely give birth to a new generation of computation technology, or, in general, artifacts resembling the natural organism and their adaptive behavior.

3 Modern Artificial Life

Artificial Life as a word came to common use through the workshop held in Los Alamos, September 1987 [18]. A second ALife workshop was held in Santa Fe, February 1990 [20], as well as a third in June 1992. The fourth workshop will be at MIT, Boston, May 1994. The "key" organizer of the workshops thus far has been Dr. Christopher G. Langton, who defines Artificial Life as follows.

> Artificial Life (ALife) 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, ALife can contribute to theoretical biology by locating *life-as-we-know-it* within the larger picture of *life-as-it-could-be*. [18, page 1]

Thus ALife includes biological life, but extends it to cover some mechanisms other than carbon-chain based chemical forms. This creates a paradox to recognize ALife, because our understanding of life, and its mechanisms, are so bound to a biological context.

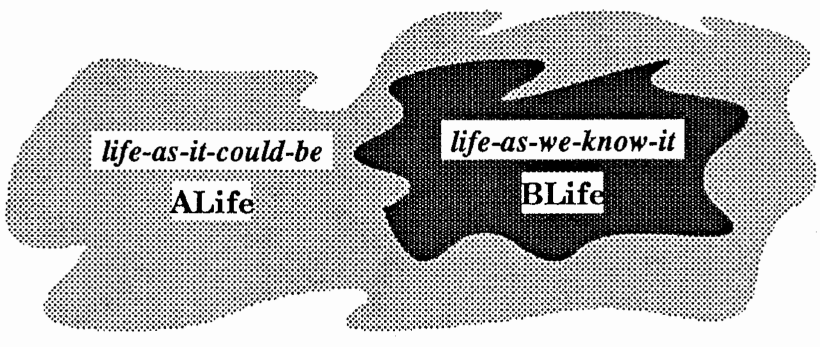

Figure 2: Artificial Life vs. Biological Life

Before discussing the concept of ALife further we need to understand how the life is defined. In fact, a general definition for life is very difficult to give. Instead we list some properties of it [12, page 818].

- *Life is a pattern in spacetime*, rather than a specific material object.
- *Self-production*, if not in the organism itself, at least in some related organisms. (Mules are alive, but cannot reproduce.)
- *Information storage of a self-representation.* For example, contemporary natural organisms store a description of themselves in DNA molecules, which is interpreted in the context of the protein/RNA machinery.
- *A metabolism* which converts matter and energy from the environment into the pattern and activities of the organism. Note that some organisms, such as viruses, do not have a metabolism of their own, but make use of the metabolisms of other organisms.
- *Functional interactions with the environment.* A living organism can respond to or anticipate changes in its environment. Organisms create and control their own local environments.
- *Interdependence of parts.* The components of living systems depend on one another to preserve the identity of the organism.
- *Stability under perturbations* and insensitivity to small changes, allowing the organism to preserve its form and continue to function in a noisy environment.
- *The ability to evolve.* This is not a property of an individual organism, but rather of its lineage.

Thus ALife is a study of artificial systems that possess at least some of the above properties. Artificial Life and Artificial Intelligence have some common properties in that both are ill-defined subjects. Both are merely used as an umbrella to cover various research fields. In the following I shall present my selection of the most important concepts and research directions belonging more or less under the ALife umbrella.

4 Emergence

One of the key concepts in ALife is *emergence* [13]. Emergence may be divided into four parts according to the emergent mechanism. First, the organism as it emerges out of the local interactions of a large number of molecules, without a global control (*cf.* morphogenesis: development of an individual) (see Fig. 3).

Second, an organism with a neural network (or some other system) can modify its behavior based on earlier experiences etc. (*cf.* ontogenesis: individual learning of organisms) (see Fig. 4).

Third, the emerging of a behavioral repertoire through the evolution process (*cf.* phylogenesis: development of species) (see Fig. 5).

Fourth, in societies of systems, a common means of communication emerges as well as other social patterns, that transfer from one generation to the next (*cf.* culture) (see Fig. 6).

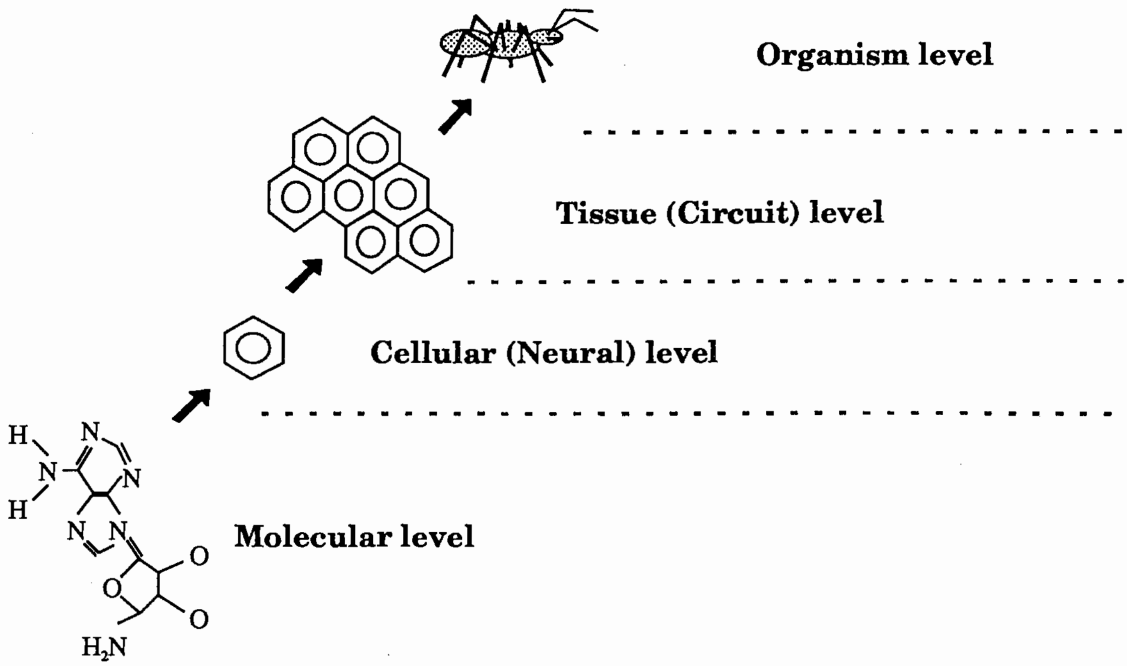

Figure 3: Morphogenesis: The synthesis of organisms.

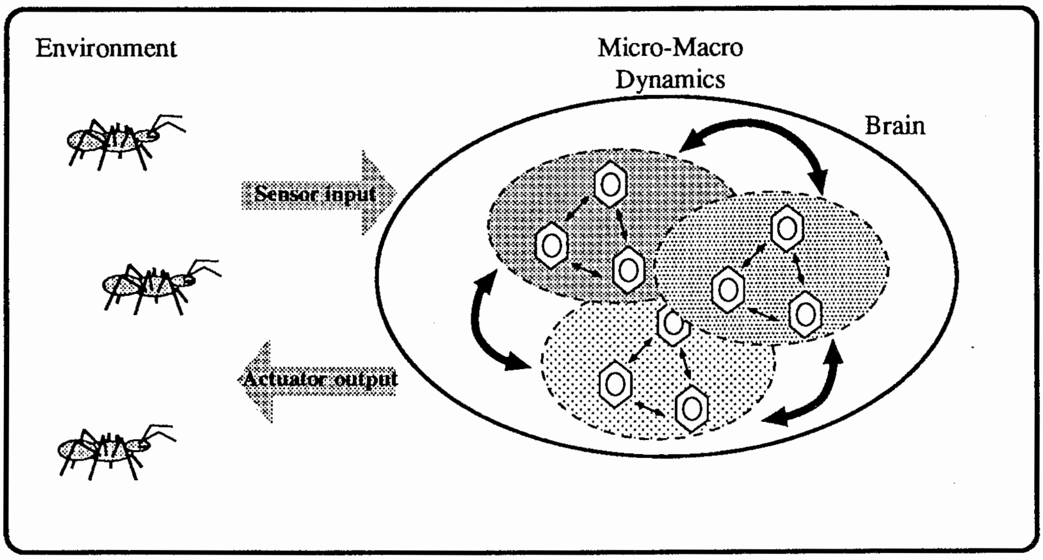

Figure 4: Ontogenesis: Learning from experience.

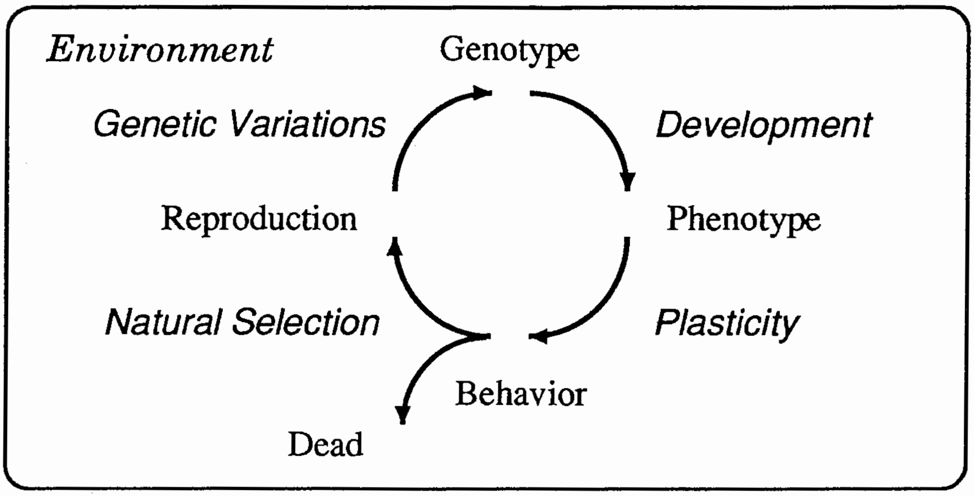

Figure 5: Phylogenesis: The creation of new species.

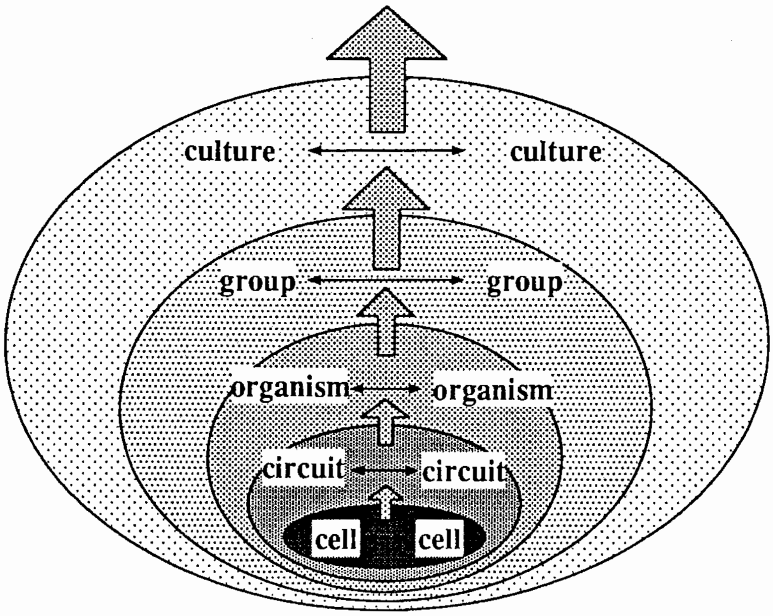

Figure 6: Culture: The development of cultures and societies.

Thus the emergence principle is a bottom-up approach, that is defined as follows by C. Langton.

> Artificial Life 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. [19, page 2]

Emergent Computation

When this bottom-up approach is applied to computation we will have a new research field of emergent computation, that is defined as follows.

> An alternative approach exploits the interactions among simultaneous computation to improve efficiency, increase flexibility, or provide a more natural presentation. Researchers in several fields have begun to explore computational models in which the behavior of the entire system is in some sense more than the sum of its parts. [...] In these systems interesting global behavior *emerges* from many local interactions. When the emergent behavior is also a computation, we refer to the system as an *emergent computation.* [13, page 1].

The point is to avoid any global control that could dictate the result of the computation process. Emergence is a process where any explicit definition of the result is avoided as far as possible. Instead of having only a single model, a population of independent models is created and the global control is replaced by a mechanism of *selection*.

Creativity: Selection vs. Elimination

The essence of the emergence principle is *creativity*. The emergence process should be able to create new forms. Usually Darwinian natural selection is understood as selection of the best individuals for reproduction. In this approach, selection of individuals is based on their *fitness* values, which are defined explicitly by the creator of the model. This approach has been applied successfully to the engineering field and it is referred to as *Genetic Algorithms* [7].

However, this approach creates a paradox of predetermining the goal, which is not the case in Nature. Instead we should concentrate on the elimination of the individuals based on the criteria of being not preferable. The definition of the elimination rules is not complementary from the definition of the selection rules. In the latter case it is necessary to define explicitly what you want, and thus to close the space of the possible products of evolution. The former case only defines what you absolutely would *not* like to have leaving the possible results open, and we can expect some creativity from the model.

5 A Dynamic Embodiment with the Environment

In Nature the behavior of systems is dynamic. There is a continuous interaction between the environment and the systems. Every change in the environment causes changes in the behavior of the system. Also every behavior of the systems causes changes in the environment. This dynamic and continuous process is modeled in ALife.

In traditional computation the focus has been on *final results*. In ALife the focus is on the *ongoing dynamic behavior*. In biological systems the behavior is a collection of smaller behaviors resulting in a dynamic behavior. Similar dynamic behavior should be achieved in artificial systems (see Fig. 7).

> [...], Artificial Life studies natural life by attempting to capture the behavioral essence of the constituent components of a living system, and endowing a collection of artificial components with similiar behavioral repertoires. If organized correctly, the aggregate of artificial parts should exhibit the same dynamic behavior as the natural system. [19, page 3]

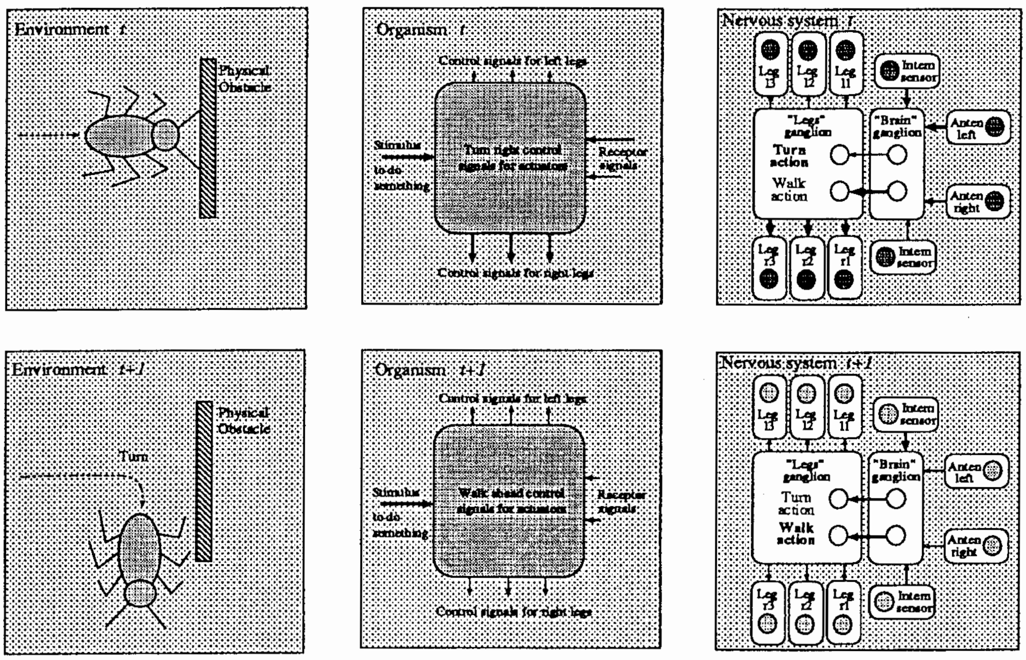

Figure 7: The behavior at the environment level emerges from the lower level interactions. At the organism level stimulus signals from the environment are converted to input for the neural structure. At the circuit level the input signals are propagated to cause response activations, that are further converted into physical movement.

The dynamic models give a new perspective for artificial intelligence as well. Artificial life concepts can be used to create artificial intelligence as adaptive behavior.

From Artificial Intelligence to Artificial Life

Traditional AI understands intelligence to be *explicitly* definable. This implies that AI systems tend to be intelligent *solutions* rather than intelligent *behaviors*. AI also started from natural intelligence, but because of lack of understanding of natural intelligence, it developed more in the artificial direction. Thus at present AI is understood as a symbolic representation of knowledge with techniques to manipulate the knowledge. These are applied to solve problems, that are thought to need intelligence. This approach has created the AI-paradox: "Whenever you explain how AI solves the problem, AI looses its intelligence".

However, there are two new ideas for AI. First, intelligence is not explicitly defined, but built up from primitive operations using the above defined emergent methods. This gives us *Emergent Intelligence*. Second, intelligence is understood to be closely embedded into the environment, and to develop together and within it. This gives us *Embedded Intelligence*.

Both ideas are directly related to ALife. In fact, it could be said that understanding intelligence is understanding life. A good description of this is given by Valentino Braitenberg [3]. In his 'mind-game' he develops vehicles (see Fig. 8) as an experiment on synthetic psychology. What we now have in ALife is a methodology to create them on computers.

In ALife we could develop a system possessing some neural-like computational elements and a set of sensors and effectors. All these can evolve during the evolution process (see Fig. 9). The result is a gradual increase in the behavior of repertoire of artifacts (see Fig. 10).

6 Implementation of ALife

So far we have discussed general concepts of ALife and two its distinguishing features: emergence and dynamic embodiment with the environment. Now we change the focus to actual work done in the field of ALife. ALife, as AI, covers interdisciplinary research areas of wetware synthesis, software synthesis, hardware synthesis and philosophical aspects of life. In the following we will give a short introduction to each of these in the light of a few examples.

6.1 Wetware synthesis

Wetware synthesis uses real chemical compounds as building elements. In the broadest sense this covers all genetic engineering and other enforcement methods to redirect biological life. However, in a more narrow sense this could be restricted to cover the molecular biology of evolving new artificial molecules.

The exploration of possibilities has just begun with very promising results. However, in this context we will only refer to the work of Gerald Joyce [2, 15] as a good starting point in this area.

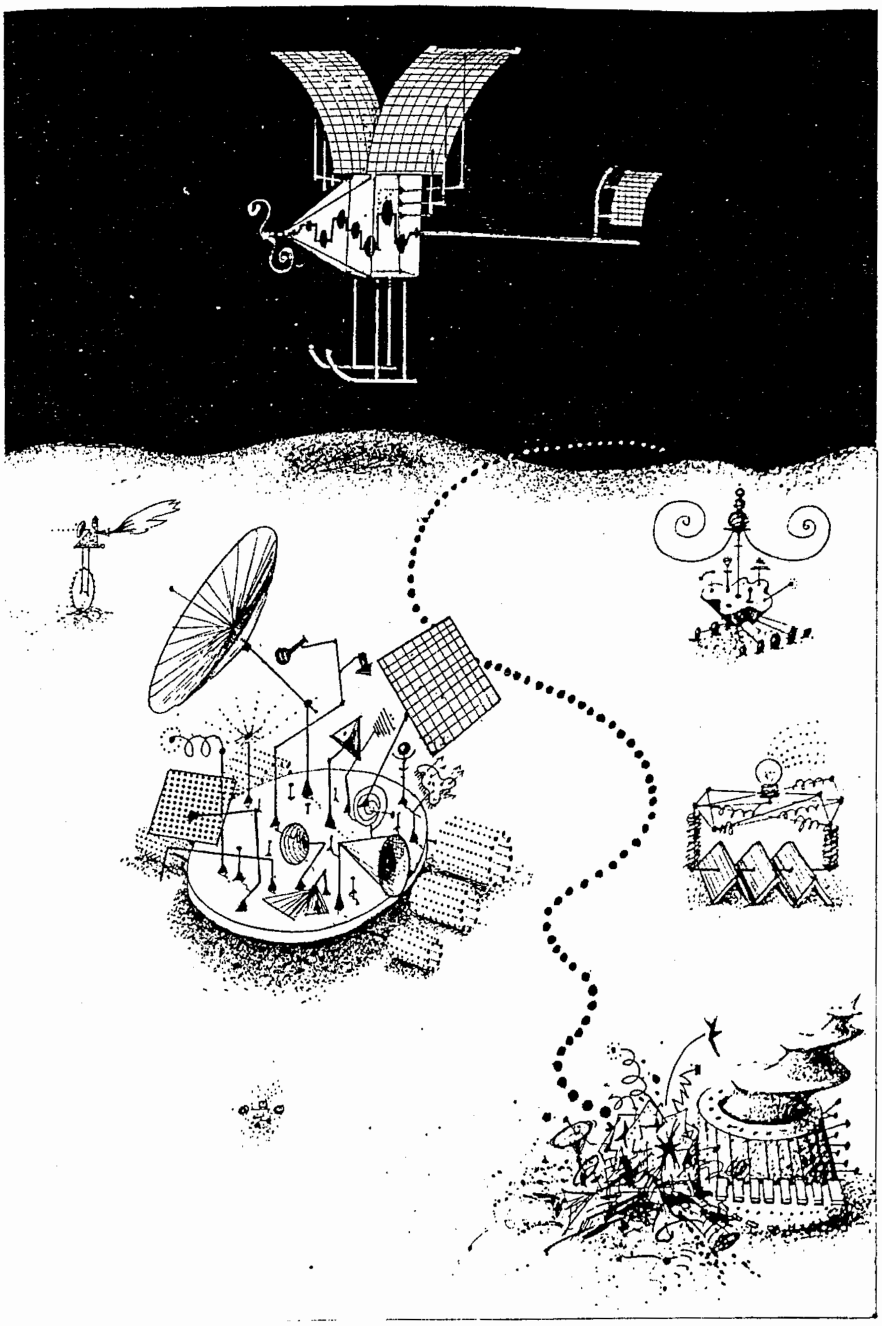

Figure 8: An artificial world as conceived by Maciek Albrecht[3, page 91]

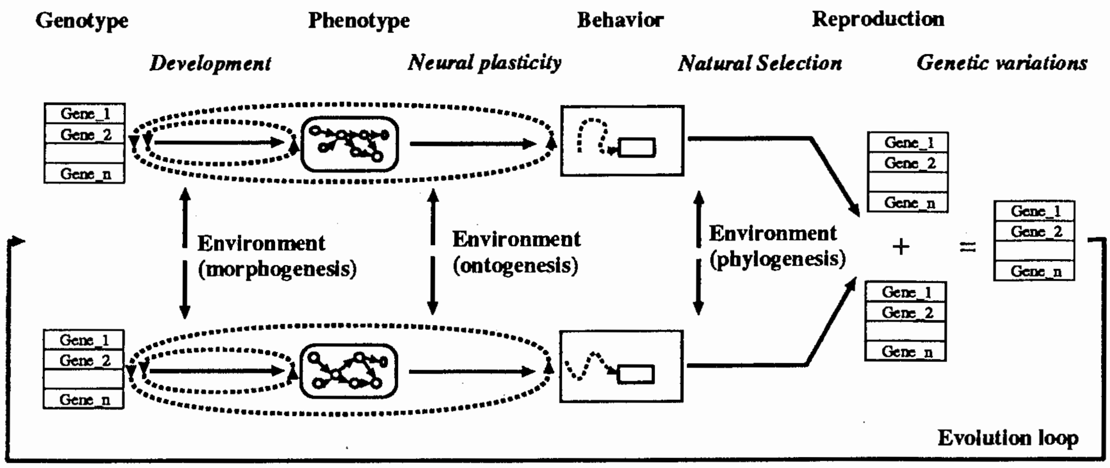

Figure 9: The evolution loop of artifacts.

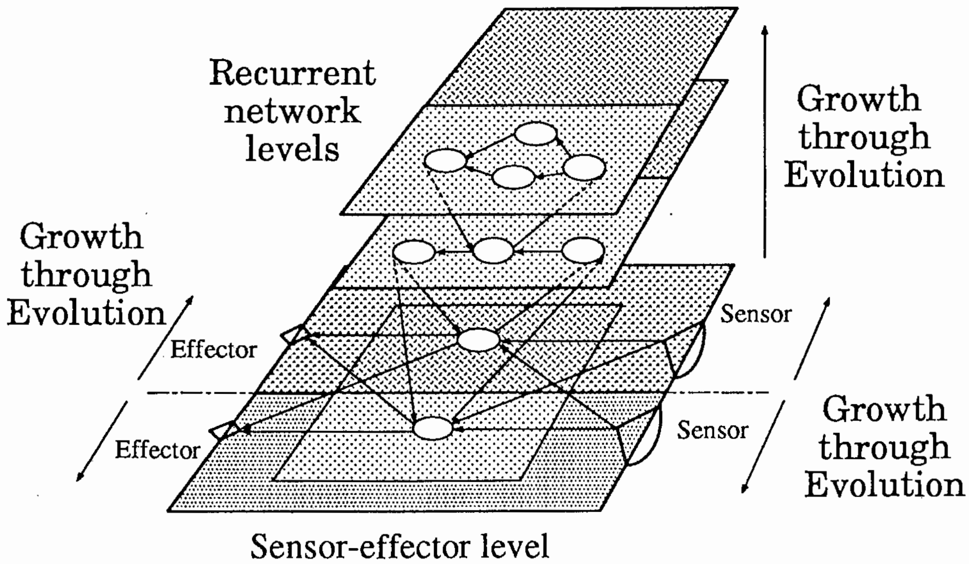

Figure 10: The growth of complexity of artifact.

6.2 Software synthesis

By software synthesis we understand software models, that try to capture the phenomena of life, particularly the developmental process, and the evolution process.

The implementation of ALife based software models can be summarized by the words of Langton as follows.

> A new approach to computation is required, one that focus on *ongoing dynamic behavior* rather than on any *final result.*
>
> The essential features of computer-based Artificial Life models are:
>
> - They consist of populations of simple programs or specifications.
> - There is no single program that directs all of the other programs.
> - Each program details the way in which a simple entry reacts to local situations in its environment, including encounters with other entities.
> - There are *no* rules in the system that dictate global behavior.
> - Any behavior at levels higher than the individual programs is therefore emergent.
>
> [19, page 3]

6.2.1 Developmental Models

The developmental process, starting from one cell, and through repeated cell divisions, cell movements, and cellular differentiation, gradually forming a mature organism is still much a mystery. This process is called *morphogenesis* and there are some computer models that try to capture the process.

Cellular Automata

Self-production has long been a dream of human technology. The first computational approach to self reproduction was due to John von Neumann.

> ... [John von Neuman] was not trying to simulate the self-production of a natural system at the level of genetics and biochemistry. He wished to abstract from the natural self-production problem its logical form.
>
> *Christopher Langton [19, page 13]*

von Neuman's ideas on cellular automata (CAs) developed from a suggestion of Stan Ulman. The basic idea in CAs is to have a set of local rules that change the state of a cell depending on its own state and on the state of the neighborhood cells.

According to Wolfram [28], CAs can be characterized by the following points.

- *Discreteness in space:* They consist of a discrete grid of spatial cells.

- *Discreteness in time:* The value of each cell is updated in a sequence of discrete time steps.
- *Discrete states:* Each cell has a finite number of possible states.
- *Homogeneous:* All cells obey the same set of state transmition rules, and are arranged in a regular array.
- *Synchronous updating:* All cell states are updated synchronously, each depending on its own state and on the states of neighboring cells.
- *Deterministic rule:* Each cell state is updated according to a fixed, deterministic, rule.
- *Spatially local rule:* The rule at each cell depends only on the states of a local neighborhood of cells around it.
- *Temporally local rule:* The rule for the new value of a site depends only on the states for a fixed number of preceding steps (usually just one).

Fig. 11 demonstrates a simple cellular automata model of self-production. Based on cellular automata studies C. Langton has proposed a theory of "Life at edge of chaos" [20], where life is defined as a complex system whose behavior lies between periodic and chaos.

The same techniques can be used to develop neural networks with growing connections. This is illustrated by the work of de Garis [10] in Fig.12.

Lindenmayer Systems

L-systems can be categorized as a rewriting system such as the *snowflake curve* proposed in 1905 by von Koch. A description of rewriting systems is given by Lindenmayer in [23, Page 1]:

> Rewriting is a technique for defining complex objects by successively replacing parts of a simple initial object using a set of *rewriting rules* or *productions.*

These rewriting systems operate best for character strings. The Chomsky's work on formal grammars spawned a great interest in string rewriting in the late 1950s. They were very widely applied in computer science.

In 1968 a biologist, Aristid Lindenmayer, introduced a new type of string-rewriting mechanism, called Lindenmayer systems (L-systems). The following differences between Chomsky grammars and L-systems are given as [23, pp. 2-3]:

> The essential difference between Chomsky grammars and L-systems lies in the method of applying productions. In Chomsky grammars productions are applied sequentially, whereas in L-systems they are applied in parallel and simultaneously replace all letters in a given word. This

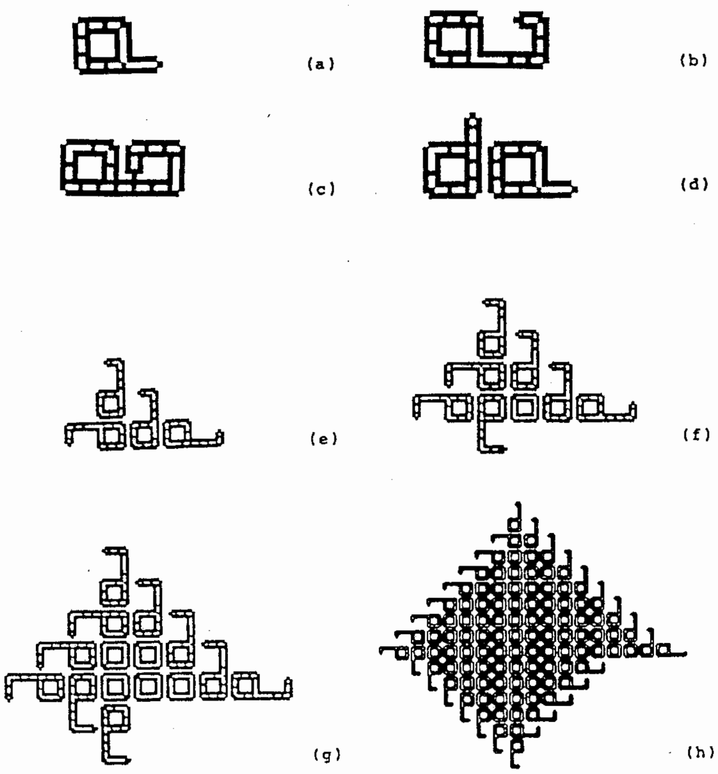

Figure 11: A cellular automata model of self-production, due to C. Langton. Signals propagating around the "Adam" loop (a) cause the short arm to grow and curl back on itself (b,c,d), producing an offspring loop (e). Each loop then goes on to produce further offspring, which also reproduce (f). This process continues indefinitely, resulting in an expanding colony of loops (g,h), consisting of a "living" reproductive fringe surrounding a growing "dead" core, as in the growth of a coral.[12, page 824]

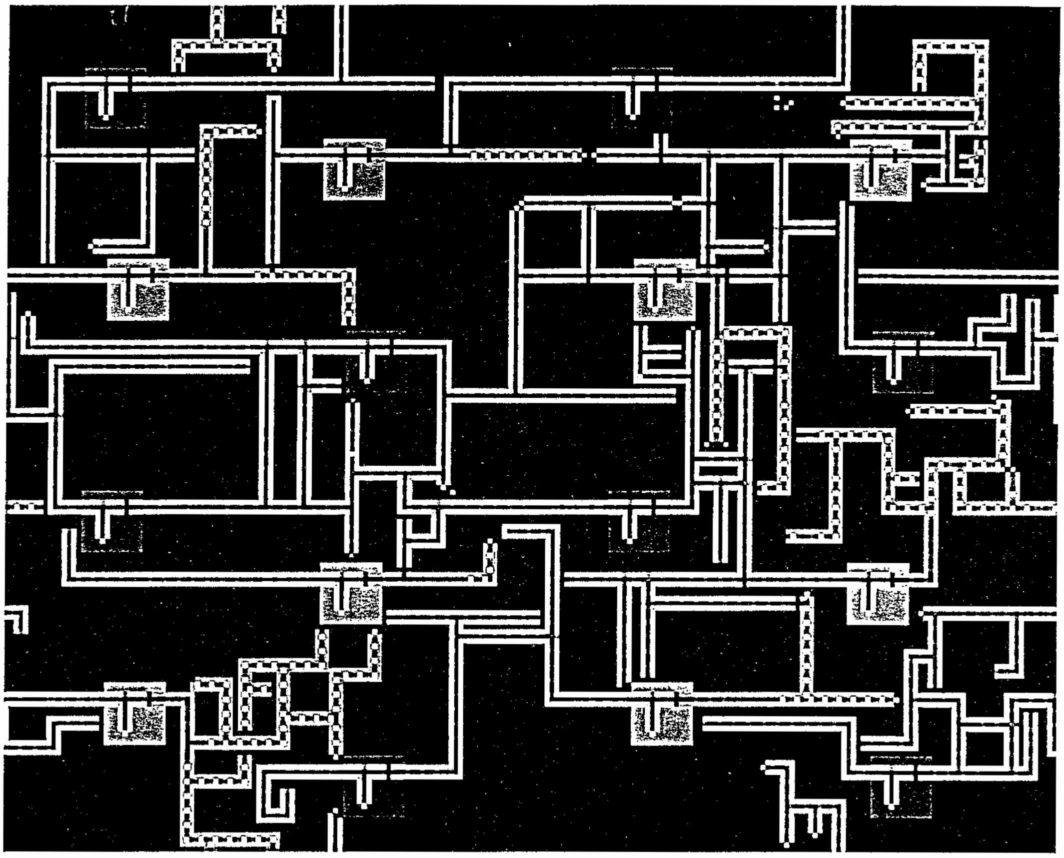

Figure 12: An evolvable cellular automata based neural network with axons and dendrites growing and forming connections with other neurons. Courtesy of Hugo de Garis: Evolved Cellular Automata Based Neural Network

difference reflects the biological motivation of L-systems. Productions are intended to capture cell divisions in multi-cellar organisms, where many divisions may occur at the same time.

The Lindenmayer systems (L-systems) have been used mainly for Computer Graphics to model plants [23]. The basic idea, however, has much wider applicability. L-systems have been successfully applied in biological modeling to capture *biological recurrent behavior.*

The basic elements of the L-system are *the (initial) word* and *production rules.* The word represents the model (data) and the production rules represent the modification rules of the model (instructions).

The word (symbol string) is a one-dimensional array of letters (symbols). In each *derivation step* all letters of the word are changed (rewritten) according to the production rules.

An example of L-systems is given in Fig. 13.

6.2.2 Models of Evolution

The simulation of evolution has been one of the most successful areas of ALife. In the following we present some of the most interesting models, although the selection is not by any means complete.

Morphology

Morphology is an interesting and important part of life. Nature creates a great variation in the details of morphology, observable in bats which are much alike except their faces (see Fig. 14). To explore the capabilities of the computer to produce different structures, Richard Dawkins developed a program called Biomorphs [9]. It is a simple system which demonstrates how small changes in the genotype cause a large change in the phenotype (see Fig. 15).

Digital Organisms — Tierra

Tom Ray's Tierra [24] is an example of how digital organisms can be created and evolved based on very simple rules. In the system the computer is thought of as an energy source, and the memory as the living space. The simple self-modifying assembler programs compete for available memory space. The programs try to copy themselves on the available memory area. The free memory is made available by eliminating old and badly behaving programs. The programs begin to show new behavior patterns while the execution is progressing. They could change their length, execute others code, and otherwise show behaviors observable in biological systems, such as parasitism, resistance to parasitism, *etc.*

Genetic Programming

John Koza [17] has applied genetic algorithm technique to LISP programming. The idea is that the programs could be created by cross-over of LISP subtrees and se-

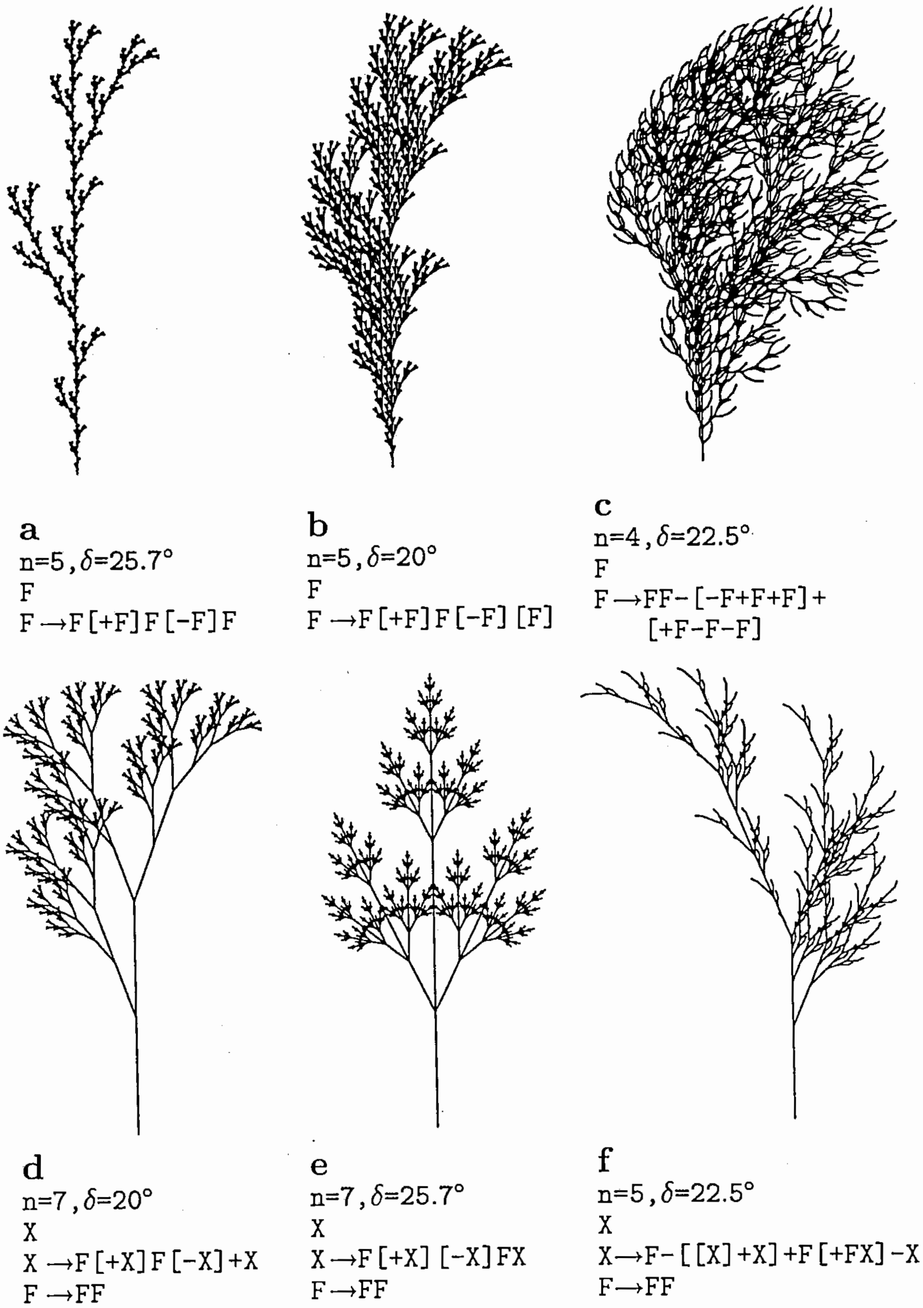

Figure 13: Examples of plant-like structures generated by OL-systems.[23, page 25] The letter "F" has a simple graphical representation of a segment. "+" and "-" represent turning to the left and right by an angle "δ" respectively. "[" and "]" represents pushing and popping the drawing position.

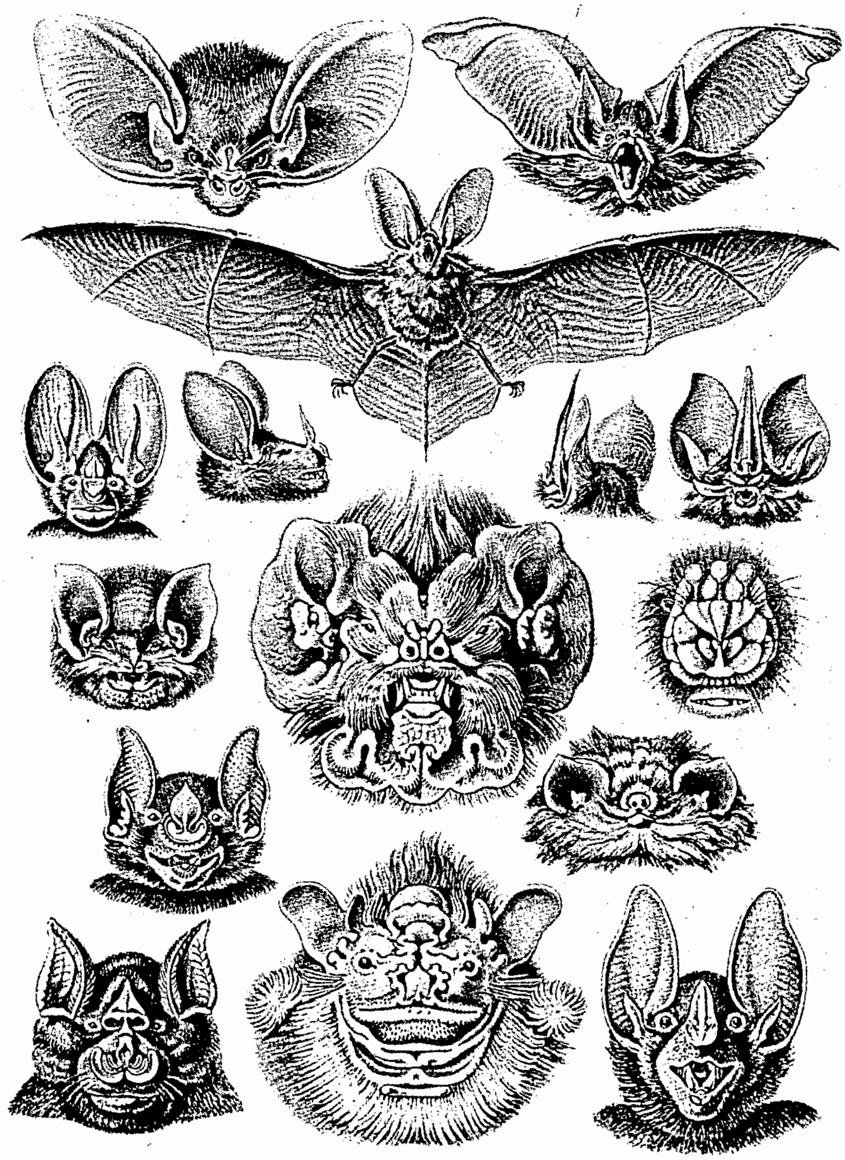

Figure 14: The faces of otherwise very similar bats.[14]

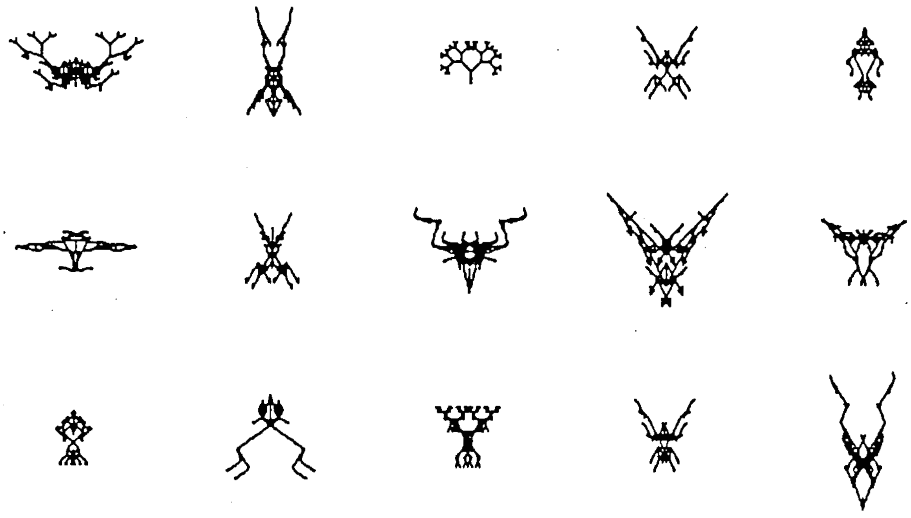

Figure 15: Computer created morphologies based on genetic variations.[9, page 209]

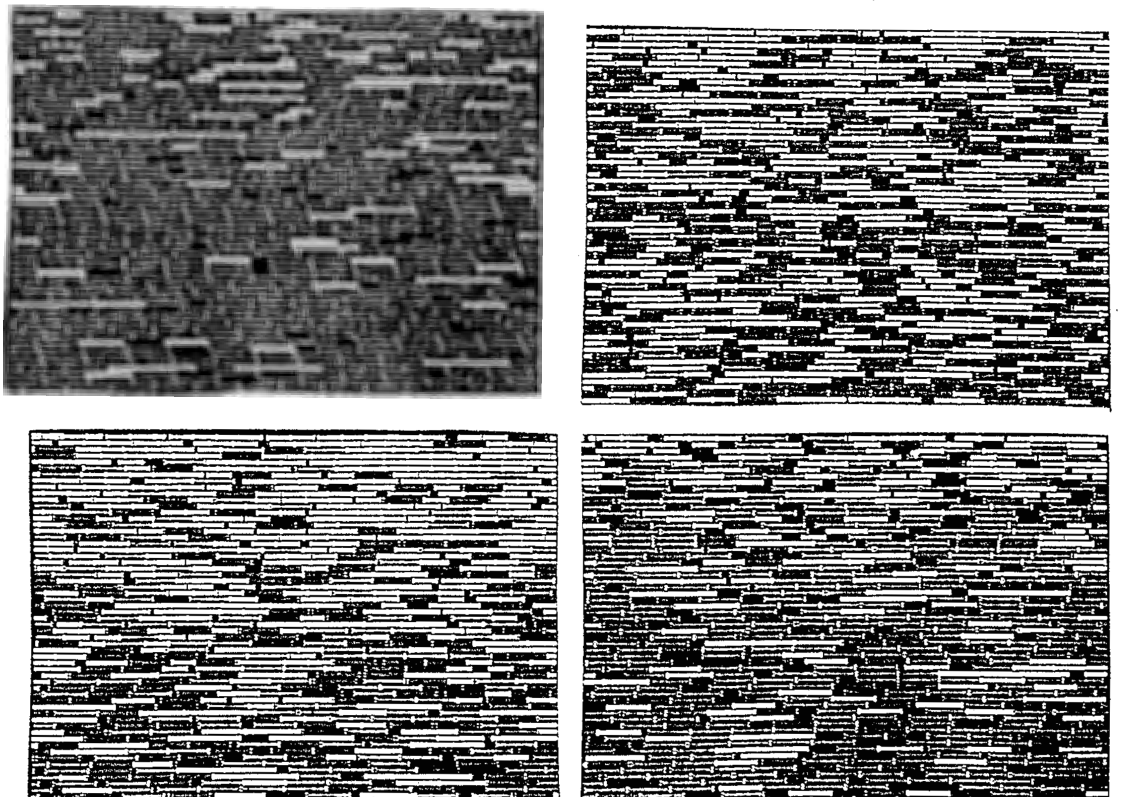

Figure 16: The evolved digital organisms are shown in the memory space. The colors represent different organisms. Top left, almost all organisms are of the same type. Top right, after a short time new types of organisms emerge as the original type loses living space to the new types. Below, the types of organisms keep changing as in Nature, while the evolution progresses.

lecting the best ones for reproduction. An example of this is presented in Fig. 17, where a robot program is evolving to perform a particular task.

Differente from the previous example of Tierra, here we have an explicitly defined task, where the Genetic Programmer plays an important role in preselecting the functions and terminal conditions used in the evolution. However, this is a good example of how the genetic algorithms can be used for extending existing programs.

Evolutionary Art

ALife concepts are already in use in art [27, 26]. This is maybe the most successful area of the evolutionary approach. The idea is to create a series of images with slight variations. Some of the generated images are selected as an intial value (genetic information) for the next generation cycle. Gradually the images become more interesting and attractive based on the judgment of the artist.

An example of this is Todd and Latham's work on evolutionary art (see Fig. 18).

6.3 Hardware synthesis

Hardware synthesis is focusing on behavior in the real environment. The field is very close to the autonomous systems research field. The approach focuses on the incremental increase of functionality resembling the evolution process of human engineering.

Subsumption Architecture

This is the approach that Brooks uses in his construction of "animal-behaving robots" where different behavior-levels are used to control the actions [4, 5] (see Fig. 19).

6.4 Philosophical Aspects

From the philosophical point of view ALife generates a lot of questions. The main question is, of course, what is life, and how to recognize it when and if we succeed to create it. There is a lot philosophical discussions concerning whether life could be explained as an emergence of physical phenomena. Thus far the discussions are based on self-modifying systems and how they can create life patterns [16, 25]. The basic idea is that all that is needed is a self-modifying system that by interpreting the genetic information and the environment, modifies itself (and thus its behavior). However, this theory is still quite far from proven inductively, or by synthesizing such a system capable to life-like behaviors. A good candinate for such a general theory is *autopoiesis*, by Maturana and Varela [21, 29].

Maybe the most famous philosophical view for life is presented by Richard Dawkins in his book "The Selfish Gene" [8]. His idea is simply that genes build a machine (organism) around them in order to increase the survival probability. Based on this simple principle he has nicely explained the birth and evolution of life.

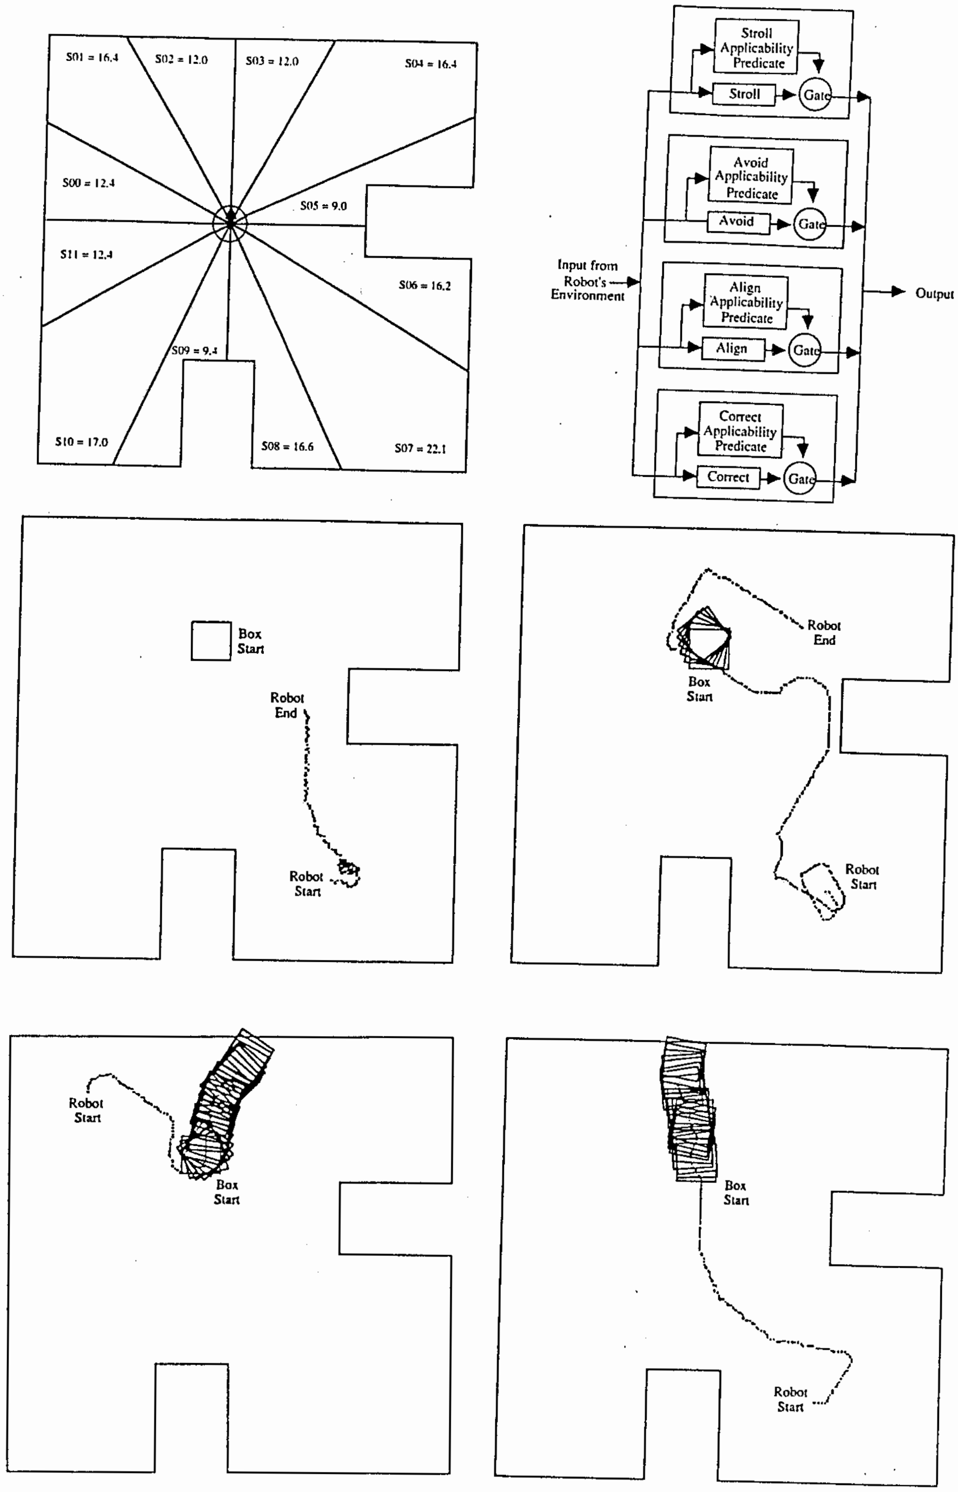

Figure 17: First row, a robot consists of 12 ultrasonic sensors and a subsumption based architecture which is genetically programmable for particular tasks: e.g. stroll, avoid, align, and correct. The object of the experiment is to evolve a program to control the robot to find a box and to move it to a wall. Second row, the first generation of programs are capable of moving the robot in the environment, but not to accomplish the task. Last row, after 45 generations the robot is capable of accomplishing the task starting from different positions in the environment. [17, page 381–387]

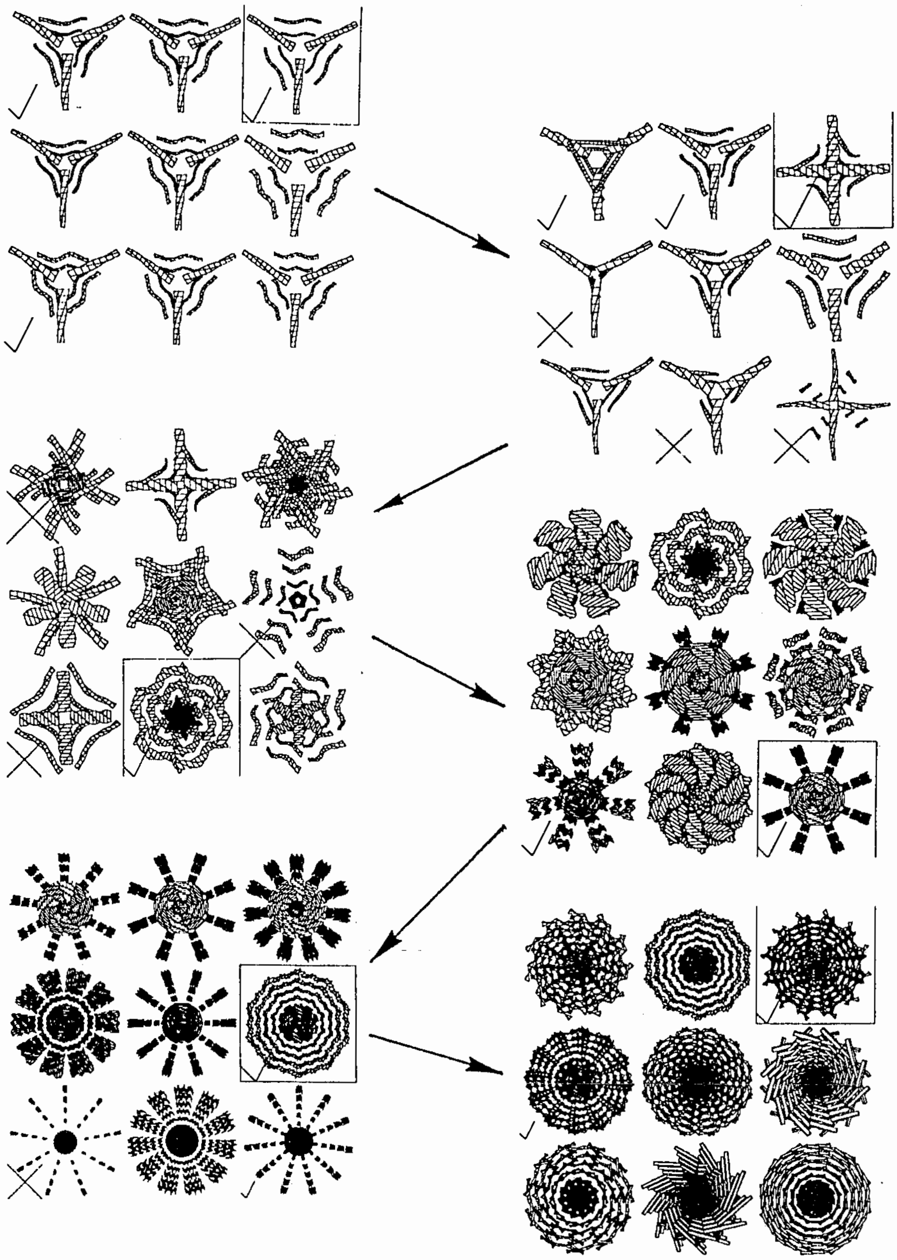

Figure 18: The evolution steps and interactive selection for the next step with artificially generated mutations.[27, page 86]

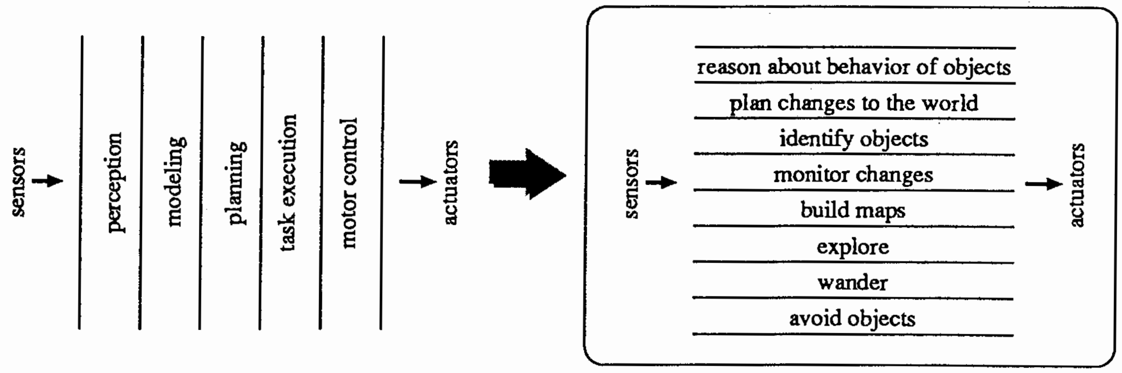

Figure 19: Subsumption based architecture.

Life could be looked at also from the culture point of view. Hans Moravec has a view that our genes are being overtaken by computers as the machine of survival [22]. This "doomsday" view might be a little bit premature, but at least it initiates some thoughts on what we are doing, and what our goals should be.

7 Conclusion

In this short ALife primer we first gave an overview of the most important concepts, namely *emergent processes*, and *a dynamic embodiment with the environment*. Then we overviewed several works done within ALife.

The principle concepts of ALife are widely applicable to the engineering field. ALife attacks the most fundamental problems of engineering: automatic evolution of complexity, self-organization of soft/wet/hardware, and emergence of intelligence without explicit design. If it is able to solve these areas ALife will become the next revolution of engineering. It has all possibilities to succeed.

There are still many problems to be solved. Biological systems are thus far the only known method to create "materialized" life. How to combine computer simulations and hardware realizations? Can Natural Selection (or more generally the evolutionary process) converge towards a desired goal? What are the necessary constraints for creation of an emergent process? There is still a long way to go to the science fiction image of human created life.

Acknowledgment

Thanks are due Tom Ray and Hugo de Garis for comments and corrections during preparation of this manuscript.

References

[1] M. A. Arbib. Simple self-reproducing universal automata. *Information and Control*, 9:177–189, 1966.

[2] A. A. Beaudry and G. F. Joyce. Directed evolution of an RNA enzyme. *Science*, 257:635–641, Jul 1992.

[3] Braitenberg. *Vehicles - Experiments in Synthetic Psychology.* The MIT Press, 1984.

[4] Rodnay A. Brooks and Paul A. Viola. Network based autonomous robot motor control: from hormones to learning. In Rolf Eckmiller, editor, *Advanced Neural Computers*, chapter 6: Motor Control with Neural Networks, pages 341–348. Elsevier Science Publisher B.V. (North-Holland), 1990.

[5] Rodney A. Brooks. Intelligence without representation. *Artificial Intelligence*, (47):139–159, 1991.

[6] Alfred Chapuis and Edmond Droz. *Automata: A Historical and technological Study.* B.A. Batsford Ltd.

[7] Lawrence Davis. *Handbook of Genetic Algorithms.* Van Nostrand Reinhold, 1991.

[8] Richard Dawkins. *The Selfish Gene.* Oxford University Press, 1976.

[9] Richard Dawkins. The evolution of evolvability. In Langton [18], pages 201–220.

[10] Hugo de Garis. Genetig Programming — GenNets, Artificial Nervous Systems, Artificial Embryos. (Wiley manuscript).

[11] Gerald M. Edelman. *Bright Air, Brilliant Fire, On the Matter of the Mind.* Basic Books, New York, 1992.

[12] J. Doyne Farmer and Alletta d'A. Belin. Artificial life: The comming evolution. In Langton et al. [20], pages 815–840.

[13] Stephanie Forrest, editor. *Emergent Computation.* The MIT Press, 1991.

[14] Ernst Haeckel. *Art Forms of Nature.* Dover Publications, Inc., 1974. Originally published by the Verlag des Bibliographischen Institus, Leipzig and Vienna, 1904.

[15] Gerald Joyce. Directed molecular volution. *Scientific American*, pages 90–97, Dec. 1992.

[16] George Kampis. *Self-Modifying Systems in Biology and Cognitive Science.* Pergamon Press, 1991.

[17] John R. Koza. *Genetic Programming: On the programming of computers by means of natural selection.* The MIT Press, 1992.

[18] Christopher G. Langton, editor. *Artificial Life.* Addison-Wesley Publishing Company, 1989.

[19] Christopher G. Langton. Artificial life. In *Artificial Life* [18], pages 1–48.

[20] Christopher G. Langton, Charles Taylor, J. Doyne Farmer, and Steen Rasmussen, editors. *Artificial Life II.* Addison-Wesley Publishing Company, 1992.

[21] H. R. Maturana and F. J. Varela. *Autopoiesis and Cognition: The Realization of the Living.* Reidel, 1980.

[22] Hans Moravec. Human culture: A genetic takeover underway. In Langton [18], pages 167–200.

[23] Premyslaw Prusinkiewicz and Aristid Lindenmayer. *The Algorithmic Beauty of Plants.* Springer-Verlag, 1990.

[24] Thomas S. Ray. An approach to the synthesis of life. In Langton et al. [20], pages 371–408.

[25] Robert Rosen. *Life Itself – A comprehensive Inquire Into the Nature, Origin, and Fabrication of Life.* Columbia University Press, New York, 1991.

[26] Karl Sims. Artificial evolution for computer graphics. *Computer Graphics*, 25(4):319–328, 1991.

[27] Stephen Todd and William Latham. *Evolutionary Art and Computers.* Academic Press, 1992.

[28] S. Wolfram, editor. *Theory and Application of Cellular Automata.* World Scientific, Singapore, 1986.

[29] M. Zeleny, editor. *Autopoiesis: A Theory of Living Organization.* North Holland, 1981.