人　工　知　能 93－5
（1994．3．8）

人工生命とリアルコンピュテーション

トーマス・レイ
進化システム研究室
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

概要

人工生命に向けた進化的アプローチは、自然淘汰による進化のプロセスを人工的なメディアで実現することから成る。それによって、進化はその人工的メディアの中での生物の自然の形を発見することができる。それらは生命のモデルではなく、独立した生命の実例である。ここで紹介する例は経験的な比較進化生物学でのひとつの実験であるとともに、計算メディアに具現化された人工生命と言える。また、人工生命が、複雑な計算機ソフトウェアの生成に進化的プロセスを利用する手段となりうることを示す。

Artificial Life and Real Computation

Thomas S. Ray
ATR Human Information Processing Research Laboratories
2-2 Hikaridai, Seika-cho, Soraku-gun, Kyoto, 619-02, Japan
ray@hip.atr.co.jp, ray@udel.edu

Abstract

An evolutionary approach to artificial life consists of inoculating the process of evolution by natural selection into an artificial medium. Evolution is then allowed to find the natural forms of living organisms in the artificial medium. These are not models of life, but independent instances of life. Examples are cited of artificial life embedded in the computational medium, where in addition to being an exercise in experimental comparative evolutionary biology, it is also a possible means of harnessing the evolutionary process for the production of complex computer software.

1 Introduction

The process of evolution by natural selection is able to create complex and beautiful information processing systems (such as primate nervous systems) without the guidance of an intelligent supervisor. Yet intelligent programmers have not been able to produce software systems that match even the capabilities of primitive organisms such as insects. Recent experiments demonstrate that evolution by natural selection is able to operate effectively in genetic languages based on the machine codes of digital computers [13, 16, 21]. This opens up the possibility of using evolution to generate complex software. The hope is that artificial life will generate real computing applications.

Evolution by natural selection is a process that enters into a physical medium. Through iterated replication-with-selection of large populations through many generations, it searches out the possibilities inherent in the "physics and chemistry" of the medium in which it is embedded. It exploits any inherent self-organizing properties of the medium, and flows into natural attractors realizing and fleshing out their structure.

Evolution never escapes from its ultimate imperative: self-replication. However, the mechanisms that evolution discovers for achieving this ultimate goal gradually become so convoluted and complex that the underlying drive can seem to become superfluous. Evolution is both a defining characteristic and the creative process of life itself. The living condition is a state that complex physical systems naturally flow into under certain conditions. It is a self-organizing, self-perpetuating state of autocatalytically increasing complexity. The living component of the physical system quickly becomes the most complex part of the system, such that it re-shapes the medium, in its own image as it were. Life then evolves adaptations predominantly in relation to the living components of the system, rather than the non-living components. Life evolves adaptations to itself.

Until recently, life has been known as a state of matter, particularly combinations of the elements carbon, hydrogen, oxygen, nitrogen and smaller quantities of many others. However, recent work in the field of Artificial Life has shown that the natural evolutionary process can proceed with great efficacy in other media, such as the informational medium of the digital computer [6, 13–16, 18–21]

These new natural evolutions, in artificial media, are beginning to explore the possibilities inherent in the "physics and chemistry" of those media. They are organizing themselves and constructing self-generating complex systems. While these new living systems are still so young that they remain in their primordial state, it appears that they have embarked on the same kind of journey taken by life on earth, and presumably have the potential to evolve levels of complexity that could lead to sentient and eventually intelligent beings.

If natural evolution in artificial media leads to sentient or intelligent beings, they will likely be so alien that they will be difficult to recognize. The sentient properties of plants are so radically different from those of animals, that they are generally unrecognized or denied by humans, and plants are merely in another kingdom of the one great tree of organic life on earth [12, 17, 22]. Synthetic organisms evolving in other media such as the digital computer, are not only not a part of the same phylogeny, but they are not even of the same physics. Organic life is based on conventional material physics, whereas digital life exists in a logical, not material, informational universe. Digital intelligence will likely be vastly different from human intelligence; forget the Turing test.

2 The Approach

The objective of the approach discussed here, is to create an instantiation of evolution by natural selection in the computational medium. This creates a conceptual problem that requires considerable art to solve: ideas and techniques must be learned by studying organic evolution, and then applied to the generation of evolution in a digital medium, without forcing the digital medium into an "unnatural" simulation of the organic world.

We must derive inspiration from observations of organic life, but we must never lose sight of the fact that the new instantiation is not organic, and may differ in many fundamental ways. For example, organic life inhabits a Euclidean space, however computer memory is not a Euclidean space. Inter-cellular communication in the organic world is chemical in nature, and therefore a single message generally can pass no more information than on or off. By contrast, communication in digital computers generally involves the passing of bit patterns, which can carry much more information.

The fundamental principal of the approach being advocated here is *to understand and respect the natural form of the digital computer, to facilitate the process of evolution in generating forms that are adapted to the computational medium, and to let evolution find forms and processes that naturally exploit the possibilities inherent in the medium.*

Situations arise where it is necessary to make significant changes from the standard computer architecture. But such changes should be made with caution, and only when there is some feature of standard computer architectures which clearly inhibits the desired processes.

3 The Computational Medium

The computational medium of the digital computer is an informational universe of boolean logic, not a material one. Digital organisms live in the memory of the computer, and are powered by the activity of the central processing unit (CPU). Whether the hardware of the CPU and memory is built of silicon chips, vacuum tubes, magnetic cores, or mechanical switches is irrelevant to the digital organism. Digital organisms should be able to take on the same form in any computational hardware implementing the same logic.

Digital organisms might as well live in a different universe from us, as they are not subject to the same laws of physics and chemistry. They are subject to the "physics and chemistry" of the rules governing the manipulation of bits and bytes within the computer's memory and CPU. They never "see" the actual material from which the computer is constructed, they see only the logic and rules of the CPU and the operating system. These rules are the only "natural laws" that govern their behavior. They are not influenced by the natural laws that govern the material universe (e.g., the laws of thermodynamics).

A typical instantiation of this type involves the introduction of a self-replicating machine language program into the RAM memory of a computer subject to random errors such as bit flips in the memory or occasionally inaccurate calculations [1,2,5,9,13]. This generates the basic conditions for evolution by natural selection as outlined by Darwin [3]: self-replication in a finite environment with heritable genetic variation.

In this instantiation, the self-replicating machine language program is thought of as the individual "digital organism" or "creature". The RAM memory provides the physical space that the creatures occupy. The CPU provides the source of energy. The memory consists of a large array of bits, generally grouped into eight bit bytes and sixteen or thirty-two bit words. Information is stored in these arrays as voltage patterns which we usually symbolize as patterns of ones and zeros.

The "body" of a digital organism is the information pattern in memory that constitutes its machine language program. This information pattern is data, but when it is passed to the CPU, it is interpreted as a series of executable instructions. These instructions are arranged in such a way that the data of the body will be copied to another location of memory. The informational patterns stored in the memory are altered only through the activity of the CPU. It is for this reason that the CPU is thought of as the analog of the energy source. Without the activity of the CPU, the memory would be static, with no changes in the informational patterns stored there.

The instruction set of the CPU, the memory, and the operating system together define the complete "physics and chemistry" of the universe inhabited by the digital organism. They constitute the physical environment within which digital organisms will evolve. Evolving digital organisms will compete for access to the limited resources of memory space and CPU time, and evolution will generate adaptations for the more agile access to and the more efficient use of these resources.

4 The Living Environment

Some rain forests in the Amazon region occur on white sand soils. In these locations, the physical environment consists of clean white sand, air, falling water, and sunlight. Embedded within this relatively simple physical context we find one of the most complex ecosystems on earth, containing hundreds of thousands of species. These species do not represent hundreds of thousands of adaptations to the physical environment. Most of the adaptations of these species are to the other living organism. The forest creates its own environment.

Life is an auto-catalytic process that builds on itself. Ecological communities are complex webs of species, each living off of others, and being lived off of by others. The system is self-constructing, self-perpetuating, and feeds on itself. Living organisms interface with the non-living physical environment, exchanging materials with it, such as oxygen, carbon-dioxide, nitrogen, and various minerals. However, in the richest ecosystems, the living components of the environment predominate over the physical components.

With living organisms constituting the predominant features of the environment, the evolutionary process is primarily concerned with adaptation to the living environment. Thus ecological interactions are an important driving force for evolution. Species evolve adaptations to exploit other species (to eat them, to parasitize them, to climb on them, to nest on them, to catch a ride on them, etc.) and to defend against such exploitation where it creates a burden.

This situation creates an interesting dynamic. Evolution is predominantly concerned with creating and maintaining adaptations to living organisms which are themselves evolving. This generates evolutionary races among groups of species that interact ecologically. These races can catalyze the evolution of upwardly spiraling complexity as each species evolves to overcome the adaptations of the others. Imagine for example, a predator and prey, each evolving to increase its speed and agility, in capturing prey, or in evading capture. This coupled evolutionary race can lead to increasingly complex nervous systems in the evolving predator and prey species.

What this discussion points to is the importance of embedding evolving synthetic organisms into a context in which they may interact with other evolving organisms. A counter example is the standard implementations of genetic algorithms in which the evolving entities interact only with the fitness function, and never "see" the other entities in the population. Many interesting behavioral, ecological and evolutionary phenomena can only emerge from interactions among the evolving entities.

5 Multi-cellularity

Multi-celled digital organisms are parallel processes. By attempting to synthesize multi-celled digital organisms we can simultaneously explore the biological issues surrounding the evolutionary transition from single-celled to multi-celled life, and the computational issues surrounding the design of complex parallel software.

5.1 Biological Perspective — Cambrian Explosion

Life appeared on earth somewhere between three and four billion years ago. While the origin of life is generally recognized as an event of the first order, there is another event in the history of life that is less well known but of comparable significance. The origin of biological diversity and at the same time of complex macroscopic multi-cellular life, occurred abruptly in the Cambrian explosion 600 million years ago. This event involved a riotous diversification of life forms. Dozens of phyla appeared suddenly, many existing only fleetingly, as diverse and sometimes bizarre ways of life were explored in a relative ecological void [7,11].

The Cambrian explosion was a time of phenomenal and spontaneous increase in the complexity of living systems. It was the process initiated at this time that led to the evolution of immune systems, nervous systems, physiological systems, developmental systems, complex morphology, and complex ecosystems. To understand the Cambrian explosion is to understand the evolution of complexity. If the history of organic life can be used as a guide, the transition from single celled to multi-celled organisms should be critical in achieving a rich diversity and complexity of synthetic life forms.

5.2 Computational Perspective — Parallel Processes

It has become apparent that the future of high performance computing lies with massively parallel architectures. There already exist a variety of parallel hardware platforms, but our ability to fully utilize the potential of these machines is constrained by our inability to write software of a sufficient complexity.

There are two fairly distinctive kinds of parallel architecture in use today: SIMD (single instruction multiple data) and MIMD (multiple instruction multiple data). In the SIMD architecture, the machine may have thousands of processors, but in each CPU cycle, all of the processors must execute the same instruction, although they may operate on different data. It is relatively easy to write software for this kind of machine, since what is essentially a normal sequential program will be broadcast to all the processors.

In the MIMD architecture, there exists the capability for each of the hundreds or thousands of processors to be executing different code, but to have all of that activity coordinated on a common task. However, there does not exist an art for writing this kind of software, at least not on a scale involving more than a few parallel processes. In fact it seems unlikely that human programmers will ever be capable of actually writing software of such complexity.

5.3 Evolution as a Proven Route

It is generally recognized that evolution is the only process with a proven ability to generate intelligence. It is less well recognized that evolution also has a proven ability to generate parallel software of great complexity. In making life a metaphor for computation we will think of the genome, the DNA, as the program, and we will think of each cell in the organism as a processor (CPU). A large multi-celled organism like a human contains trillions of cells/processors. The genetic program contains billions of nucleotides/instructions.

In a multi-celled organism, cells are differentiated into many cell types such as brain cells, muscle cells, liver cells, kidney cells, etc. The cell types just named are actually general classes of cell types within which there are many sub-types. However, when we specify the ultimate indivisible types, what characterizes a type is the set of genes it expresses. Different cell types express different combinations of genes. In a large organism, there will be a very large number of cells of most types. All cells of the same type express the same genes.

The cells of a single cell type can be thought of as exhibiting parallelism of the SIMD kind, as they are all running the same "program" by expressing the same genes. Cells of different cell types exhibit MIMD parallelism as they run different code by expressing different genes. Thus large multi-cellular organisms display parallelism on an astronomical scale, combining both SIMD and MIMD parallelism into a beautifully integrated whole. From these considerations it is evident that evolution has a proven ability to generate massively parallel software embedded in wetware. The computational goal of evolving multi-cellular digital organisms is to produce such software embedded in hardware.

5.4 Digital "Neural Networks" — Natural Artificial Intelligence

One of the greatest challenges in the field of computer science is to produce computer systems that are "intelligent" in some way. This might involve for example, the creation of a system for the guidance of a robot which is capable of moving freely in a complex environment, seeking, recognizing and manipulating a variety of objects. It might involve the creation of a system capable of communicating with humans in natural spoken human language, or of translating between human languages.

It has been observed that natural systems with these capabilities are controlled by nervous systems consisting of large numbers of neurons interconnected by axons and dendrites. Borrowing from nature, a great deal of work has gone into setting up "neural networks" in computers [4,8]. In these systems, a collection of simulated "neurons" are created, and connected so that they can pass messages. The learning that takes place is accomplished by adjusting the "weights" of the connections.

Organic neurons are essentially analog devices, thus when neural networks are implemented on computers, they are digital emulations of analog devices. There is a certain inefficiency involved in emulating an analog device on a digital computer. For this reason, specialized analog hardware has been developed for the more efficient implementation of artificial neural nets [10].

Neural networks, as implemented in computers, either digital or analog, are intentional mimics of organic nervous systems. They are designed to function like natural neural networks in many details. However, natural neural networks represent the solution found by evolution to the problem of creating a control system based on organic chemistry. Evolution works with the physics and chemistry of the medium in which it is embedded.

The solution that evolution found to the problem of communication between organic cells is chemical. Cells communicate by releasing chemicals that bind to and activate receptor molecules on target cells. Working within this medium, evolution created neural nets. Inter-cellular chemical communication in neural nets is "digital" in the sense that chemical messages are either present or not present (on or off). In this sense, a single chemical message carries only a single bit of information. More detailed information can be derived from the temporal pattern of the messages, and also the context of the message. The context can include where on the target cell body the message is applied (which influences its "weight"), and what other messages are arriving at the same time, with which the message in question will be integrated.

It is hoped that evolving multi-cellular digital organisms will become very complex, and will contain some kind of control system that fills the functional role of the nervous system. While it seems likely that the digital nervous system would consist of a network of communicating "cells", it seem unlikely that this would bear much resemblance to conventional neural networks.

Compare the mechanism of inter-cellular communication in organic cells (described above), to the mechanisms of inter-process communication in computers. Processes transmit messages in the form of bit patterns, which may be of any length, and so which may contain any amount of information. Information need not be encoded into the temporal pattern of impulse trains. This fundamental difference in communication mechanisms between the digital and the organic mediums must influence the course that evolution will take as it creates information processing systems in the two mediums.

It seems highly unlikely that evolution in the digital context would produce information processing systems that would use the same forms and mechanisms as natural neural nets (e.g., weighted connections, integration of incoming messages, threshold triggered all or nothing output, thousands of connections per unit). The organic medium is a physical/chemical medium, whereas the digital medium is a logical/informational medium. That observation alone would suggest that the digital medium is better suited to the construction of information processing systems.

If this is true, then it may be possible to produce digitally based systems that have functionality equivalent to natural neural networks, but which have a much greater simplicity of structure and process. Given evolution's ability to discover the possibilities inherent in a medium, and it's complete lack of preconceptions, it would be very interesting to observe what kind of information processing systems evolution would construct in the digital medium. If evolution is capable of creating network based information processing systems, it may provide us with a new paradigm for digital "connectionism", that would be more natural to the digital medium than simulations of natural neural networks.

6 Acknowledgements

This work was supported by grants CCR-9204339 and BIR-9300800 from the United States National Science Foundation, a grant from the Digital Equipment Corporation, and by the Santa Fe Institute, Thinking Machines Corp., IBM, and Hughes Aircraft. This work was conducted while at: School of Life & Health Sciences, University of Delaware, Newark, Delaware, 19716, USA, ray@udel.edu; and Santa Fe Institute, 1660 Old Pecos Trail, Suite A, Santa Fe, New Mexico, 87501, USA, ray@santafe.edu.

References

[1] Barton-Davis, Paul. Unpublished. Independent implementation of the Tierra system, contact: pauld@cs.washington.edu.

[2] Brooks, Rodney. Unpublished. Brooks has created his own Tierra-like system, which he calls Sierra. In his implementation, each machine instruction consists of an opcode and an operand. Successive instructions overlap, such that the operand of one instruction is interpreted as the opcode of the next instruction. Contact: brooks@ai.mit.edu

[3] Darwin, Charles. 1859. On the origin of species by means of natural selection or the preservation of favored races in the struggle for life. London: Murray.

[4] Dayhoff, Judith. 1990. Neural network architectures. New York: Van Nostrand Reinhold. Pp. 259.

[5] de Groot, Marc. Unpublished. Primordial soup, a Tierra-like system that has the additional ability to spawn self-reproducing organisms from a sterile soup. Contact: marc@kg6kf.ampr.org, marc@toad.com, marc@remarque.berkeley.edu

[6] Feferman, Linda. 1992. Simple rules... complex behavior [video]. Santa Fe, NM: Santa Fe Institute. Contact: fef@santafe.edu, 0005851689@mcimail.com

[7] Gould, Steven J. 1989. Wonderful life. W. W. Norton & Company, Inc. Pp. 347.

[8] Hertz, John, Anders Krogh, and Richard G. Palmer. 1991. Introduction to the theory of neural computation. Addison-Wesley Publishing Co. Pp. 327.

[9] Maley, Carlo C. 1993. A model of early evolution in two dimensions. Masters of Science thesis, Zoology, New College, Oxford University. Contact: cmaley@oxford.ac.uk

[10] Mead, Carver. 1993. Analog VLSI and neural systems. Addison-Wesley Publishing Co. Pp. 371.

[11] Morris, S. Conway. 1989. Burgess shale faunas and the Cambrian explosion. Science 246: 339–346.

[12] Ray, T. S. 1979. Slow-motion world of plant 'behavior' visible in rainforest. Smithsonian 9(12): 121–30.

[13] ———. 1991. An approach to the synthesis of life. *In*: Langton, C., C. Taylor, J. D. Farmer, & S. Rasmussen [eds], Artificial Life II, Santa Fe Institute Studies in the Sciences of Complexity, vol. X, 371–408. Redwood City, CA: Addison-Wesley.

[14] ———. 1991. Population dynamics of digital organisms. *In*: Langton, C. G. [ed.], Artificial Life II Video Proceedings. Redwood City, CA: Addison Wesley.

[15] ———. 1991. Is it alive, or is it GA? *In*: Belew, R. K., and L. B. Booker [eds.], Proceedings of the 1991 International Conference on Genetic Algorithms, 527–534. San Mateo, CA: Morgan Kaufmann.

[16] ———. 1991. Evolution and optimization of digital organisms. *In*: Billingsley K. R., E. Derohanes, H. Brown, III [eds.], Scientific Excellence in Supercomputing: The IBM 1990 Contest Prize Papers, Athens, GA, 30602: The Baldwin Press, The University of Georgia.

[17] ———. 1992. Foraging behaviour in tropical herbaceous climbers (Araceae). Journal of Ecology. 80: 189–203.

[18] ———. 1992. Tierra.doc. Documentation for the Tierra Simulator V4.0, 9–9–92. Newark, DE: Virtual Life. The full source code and documentation for the Tierra program is available by anonymous ftp at: tierra.slhs.udel.edu [128.175.41.34] and life.slhs.udel.edu [128.175.41.33], or by contacting the author.

[19] ———. 1994. An evolutionary approach to synthetic biology: Zen and the art of creating life. Artificial Life 1(1/2): 195–226

[20] ———. In press A. Evolution and complexity. *In*: Cowan, George A., David Pines and David Metzger [eds.], Complexity: Metaphor and Reality. Addison-Wesley Publishing Co.

[21] ———. In press B. Evolution, complexity, entropy, and artificial reality. Physica D.

[22] Strong, D. R. and T. S. Ray. 1975. Host tree location behavior of a tropical vine (*Monstera gigantea*) by skototropism. Science, 190: 804–06.