J. Doyne Farmer† and Alletta d'A. Belin‡
†Complex Systems Group, Theoretical Division, and Center for Nonlinear Studies, Lo
Alamos National Laboratory, Los Alamos, NM 87545 and Santa Fe Institute, 1120 Canyo
Road, Santa Fe, NM 87501 and ‡Shute, Mihaly, and Weinberger, P.O. Box 2768, Sant
Fe, New Mexico 87504

Artificial Life: The Coming Evolution

Within fifty to a hundred years, a new class of organisms is likely to emerge. These organisms will be artificial in the sense that they will originally be designed by humans. However, they will reproduce, and will evolve into something other than their initial form; they will be "alive" under any reasonable definition of the word. These organisms will evolve in a fundamentally different manner than contemporary biological organisms, since their reproduction will be under at least partial conscious control, giving it a Lamarckian component. The pace of evolutionary change consequently will be extremely rapid. The advent of artificial life will be the most significant historical event since the emergence of human beings. The impact on humanity and the biosphere could be enormous, larger than the industrial revolution, nuclear weapons, or environmental pollution. We must take steps now to shape the emergence of artificial organisms; they have potential to be either the ugliest terrestrial disaster, or the most beautiful creation of humanity.

Artificial Life II, SFI Studies in the Sciences of Complexity, vol. X, edited by
C. G. Langton, C. Taylor, J. D. Farmer, & S. Rasmussen, Addison-Wesley, 1991

PREFACE

This paper was originally prepared for a conference held in honor of Murray Gell-Mann's 60th birthday, whose theme was "Where Are Our Efforts Leading?" The speakers were presented with a list of sixteen questions, which Gell-Mann had picked out as great challenges in science and human affairs. We were asked to pick one or more of these challenges and to examine the efforts of society to address them.

The challenges were very broad; a few examples from the list are:

6. To understand the common features of complex adaptive systems, such as the brain and the mind, the immune system, biological evolution, prehistoric chemical evolution, the generation of new strategies in computers, the evolution of human language, and the rise and fall of human cultures.
9. To find ways to cope with human "tribalism" in its many manifestations, including national, ethnic, and religious rivalries.
11. To build worldwide institutions, formal and informal, that may permit mankind to continue to avert large-scale catastrophe through management of conflict and management of the biosphere.

We chose artificial life as our subject for this symposium, since we felt that it was a topic that was pertinent to many of these challenges, but which had not received much popular attention. This paper was written for a popular, non-scientific audience; readers that are already familiar with artificial life may wish to skip the latter part of section 2. It was submitted for publication in the proceedings of that conference entitled "Proceedings in Celebration of Murray Gell-Mann's 60th Birthday" (Cambridge University Press ©1991) and is reprinted by permission of the publisher and the authors.

1. INTRODUCTION

Murray Gell-Mann posed some difficult questions for this symposium. Among them are: Where will our efforts lead in 50 to 100 years? What are the most important challenges that we face, for both science and society? What should people be thinking about that they are not properly aware of?

One answer to each of these questions concerns the advent of "artificial life." Within the next century we will likely witness the introduction on earth of living organisms originally designed in large part by humans, but with the capability to reproduce and evolve just as natural organisms do. This promises to be a singular and profound historical event—probably the most significant since the emergence of human beings.

The study of artificial life is currently a novel scientific pursuit—a quest to understand some of the most fundamental questions in physics and biology. This field is in its infancy. There are very few researchers actively engaged in the study of artificial life, and as yet there are far more problems than solutions. Studying artificial life has the potential to put the theory of evolution in a broader context and to help provide it with a firmer mathematical basis.

The advent of artificial life also has deep philosophical implications. It prompts us to reexamine our anthropocentric views and raises numerous questions about the nature and meaning of life. In addition, the study of artificial life may help us to understand, guide, and control the emergence of artificial life on earth, thereby averting a potential disaster and perhaps helping to create beautiful and beneficial new life-forms instead.

2. WHAT IS ARTIFICIAL LIFE?

In a recent book on the subject,[14] the discipline of artificial life was defined by Chris Langton as "the study of man-made systems that exhibit behaviors characteristic of natural living systems." A primary goal of this field is to create and study artificial organisms that mimic natural organisms.

We are used to thinking of evolution as a phenomenon specific to life on earth. Biology as it is commonly practiced is, in this sense, a parochial subject. The only example of life at hand is carbon-based life on earth. All life-forms on earth involve the same basic mechanisms. They all reproduce and develop under the control of the protein and DNA-templating machinery. However, it is not at all clear that this is the *only* possible basis for life. It is easy to conceive of other forms of life, in different media, with a variety of different reproductive and developmental mechanisms.

One motivation for thinking about life at this level of generality is the question, "If we ever make contact with life from other planets, will our science of biology help us understand it?" The answer depends very much on how universal the characteristics of life on earth are to all life-forms. Since we know nothing about life on other planets, it is a difficult question to answer. It seems probable, however, that much of our biology will simply be inapplicable to other life-forms. A central motivation for the study of artificial life is to extend biology to a broader class of life-forms than those currently present on the earth, and to couch the principles of biology in the broadest possible terms.

2.1 WHAT IS LIFE?

In order to state how something artificial might also be alive, we must first address the question of what life is, as generally as we can. To see why this is a difficult question, consider a related question to the one above: If we voyage to another planet, how will we know whether or not life is present? If we admit the possibility that life could be based on very different materials than life on earth, then this becomes a difficult task. Obviously, we cannot answer this question unless we have a general definition of what it means to be "alive." At present we do not have a good answer to this question.[1]

Nonetheless, we will make an attempt to state some of the criteria that seem to bear on the nature of life. There seems to be no single property that characterizes life. Any property that we assign to life is either too broad, so that it characterizes many nonliving systems as well, or too specific, so that we can find counter-examples that we intuitively feel to be alive, but that do not satisfy it. Albeit incomplete and imprecise, the following is a list of properties that we associate with life:

- *Life is a pattern in spacetime*, rather than a specific material object. For example, most of our cells are replaced many times during our lifetime. It is the pattern and set of relationships that are important, rather than the specific identity of the atoms.
- *Self-reproduction*, 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 (internal) environments.
- *Interdependence of parts.* The components of living systems depend on one another to preserve the identity of the organism. One manifestation of this is the ability to die. If we break a rock in two, we are left with two smaller rocks; if we break an organism in two, we often kill it.
- *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. Indeed, the possession of a lineage is an important feature of living systems.

Another property that might be included in this list is growth. Growth is not a very specific property, however; there are many inanimate structures such as mountains, crystals, clouds, rust, or garbage dumps that have the ability to grow.

---

[1] One attempt has been made by Schrödinger in his book, "What is Life?" However, the discussion is heavily based on life as we know it rather than life as it might be, and as Schrödinger himself admits, the description is highly incomplete. Perhaps the best discussion of this issue is that of Monod.[18] He defines life in terms of three qualities: (1) teleonomic or "purposeful" behavior; (2) autonomous morphogenesis; and (3) invariance of information. The latter two are similar to some of the criteria we present here, but the first criterion seems as difficult to define as life itself.

Many mature organisms do not grow. Once they replicate, viruses do not usually grow.

It is not clear that life should be an either/or property. Organisms such as viruses are in many respects midway between what we normally think of as living and non-living systems. It is easy to conceive of other forms, for example the "proto-organisms" in some origin of life models,^{1,7,8} that are "partially alive." In a certain sense societies and ecosystems may be regarded as living things. It seems more appropriate to consider life as a continuum property of organizational patterns, with some more or less alive than others.

This list is far from adequate—an illustration of the poverty of our understanding. We hope that as the field of artificial life develops, one of its accomplishments will be to give a sharper definition of what it means to be alive.

2.2 EXAMPLES OF OTHER LIFE-FORMS

The creation of new life-forms will almost certainly broaden our understanding of life, for several reasons:

- The act of construction is instructive about the nature of function.
- Artificial life-forms provide a broader palette, making it easier to separate the universal from the parochial aspects of life.
- Dissection and data gathering are potentially much simpler, particularly for life-forms that exist only inside a computer.

In the latter sense artificial life is to biology as physics is to astronomy: In astronomy we can only observe, but in physics we can perform experiments to test our hypotheses, altering the universe to enhance our understanding of it. Life, however, is a collective phenomenon, the essence of which is the interaction of the parts—too large an alteration results in death. Our ability to dissect or alter the form of natural organisms is limited. In contrast, we have complete knowledge of artificial organisms inside a computer, and furthermore we have the ability to alter their structure as well as that of the artificial universe in which they reside.[2] Similarly, by recreating new forms of life inside a test tube, we may understand these underlying principles more thoroughly.

There are many possible media for artificial organisms. They might be made of carbon-based materials in an aqueous environment, similar to natural organisms; they might be robots, made of metal and silicon; or they might be abstract mathematical forms, represented as patterns of electrons existing only inside a computer.

---

[2] For a provocative and entertaining discourse on the potential ethical problems involved in the study of artificial organisms, see "The Experiment" by Stanislaw Lem.^{15}

2.2.1 COMPUTER VIRUSES Much of current research in artificial life focuses on computer programs or elements of computer programs that might be considered living organisms. It may be difficult to understand how this may be life, so we will begin by discussing the notorious example of computer viruses. Although computer viruses are not fully alive, they embody many of the characteristics of life, and it is not hard to imagine computer viruses of the future that will be just as alive as biological viruses.

These viruses are computer programs that reproduce themselves, typically designed as practical jokes by computer hackers. They are a diverse lot, and can live in many different media. For example, many viruses spend most of their life on floppy disks. Suppose a friend gives you a floppy disk that is infected with a virus. When you put the disk into your personal computer, the virus attempts to copy itself into the machine; when you insert another floppy disk, the virus attempts to copy itself onto the new disk. If the virus is effective, you may discover that, perhaps without your knowledge, it has infected all your floppy disks. If it is virulent, you may find that it takes up a great deal of space on your floppy disks, or that when it enters your machine, it causes the machine to spend much of its time executing the virus program rather than the task that you *want* the machine to perform. If the virus is really malevolent, it may destroy other programs that you have stored on your floppy disks.

A computer virus is certainly not life as we know it. It is just a pattern, a particular magnetic configuration on a floppy disk, or a particular set of electronic states inside a computer. Is the computer virus alive?

Note that a computer virus satisfies most, and potentially all, of the criteria that we have stated:

- A computer virus is a pattern on a computer memory storage device.
- A computer virus can copy itself to other computers, thereby reproducing itself.
- A computer virus stores a representation of itself.
- Like a real virus, a computer virus makes use of the metabolism of its host (the computer) to modify the available storage medium. The computer virus can direct the conversion of electrical energy into heat to change the composition of a material medium—it uses energy to preserve its form and to respond to stimuli from other parts of the computer (its environment).
- A computer virus senses changes in the computer and responds to them in order to procreate.
- The parts of a computer virus are highly interdependent; a computer virus can be killed by erasing one or more of the instructions of its program.
- Although many viruses are *not* stable under large electrical perturbations, by the nature of the digital computer environment, they are stable to small noise fluctuations. A truly robust virus might also be stable under some alterations of its program.
- Computer viruses evolve, although primarily through the intermediary of human programmers; an examination of the structure of computer viruses naturally places them in a taxonomic tree with well-defined lineages. For current

computer viruses random variation in computer virus programs is almost always destructive, although some more clever viruses contain primitive built-in self-alteration mechanisms that allow them to adapt to new environments, or that make them difficult to detect and eliminate. Thus, contemporary viruses do not evolve naturally.

Although computer viruses live in an artificial medium that we cannot directly see, they nonetheless possess most of the properties we have listed as characteristic of life, except possibly the last two. Computer viruses are already more than just a curiosity, and software infected by viruses is becoming increasingly common. During the fall of 1988, a computer virus propagated across the ARPA network (a fast communication link built by the defense department for interconnecting geographically separated computers), and brought computer operations at many major universities and national laboratories to a standstill.

Computer viruses are just one of many possible artificial life-forms, selected for discussion because they have already emerged, and because they illustrate how artificial life-forms can appear to be fundamentally different from more familiar contemporary biological life-forms. Because of their instability and their dependence on human intervention in order to evolve, they not as fully “alive” as their biological counterparts. However, as computers become more prevalent, more complex, and more highly interconnected, we suspect that so will computer viruses. Eventually it is likely that a computer virus will be created with a robust capacity to evolve, that will progress far beyond its initial form.

One example of computer organisms that evolve within a restricted environment is already provided by the VENUS simulation of Rasmussen et al.[20] Their work was inspired by a computer game called “Core Wars,” in which hackers create computer programs that battle for control of a computer’s “core” memory.[6] Since computer programs are just patterns of information, a successful program core wars is one that replicates its pattern within the memory, so that eventually most of the memory contains its pattern rather than that of the competing program.

VENUS is a modification of Core Wars in which the computer programs can mutate. Furthermore, each memory location is endowed with “resources,” which, like sunshine, are added at a steady rate. A program must have sufficient resources in the regions of memory it occupies in order to execute. The input of resources determines whether the VENUS ecosystem is a “jungle” or a “desert.” In jungle environments Rasmussen et al. observe the spontaneous emergence of of primitive “copy/split” organisms starting from (structured) random initial conditions. Note that since these “organisms” are contained by a highly specialized computer environment, there is no possibility of escape into the computer operating system. Such a protocol for containment is followed by all responsible researchers in artificial life.

2.2.2 MACHINES AND AUTOMATA A machine may be defined as "an apparatus consisting of interrelated parts with separate functions." Like an organism, a machine can break or die. One of the main features that distinguishes machines from organisms is the ability for self-reproduction. However, as demonstrated by John von Neumann in the late 1940's, it is possible, at least in principle, to build self-reproducing machines. Von Neumann imagined an "environment" filled with spare parts. The hypothetical machines in this environment had descriptions of themselves, and "construction arms" for acquiring and assembling the spare parts, all under the control of a computer. He sketched out the basic principles that such self-reproducing machines might follow, and laid out a blueprint for how they might operate.

Such a mechanical world is too complicated for simple mathematical analysis. Von Neumann, like contemporary researchers in artificial life, wanted to study the emergence and functioning of life in order to discover the basic principles that distinguish life from non-life. He was searching for an abstract environment to facilitate the study of these questions, in which simple patterns can be created that have lifelike properties. His hope was that by creating environments that give rise to pseudo-organisms he could gain an understanding of the fundamental properties of life itself.

Toward this end he turned to an abstract mathematical world, whose inhabitants are mathematical patterns. Following a suggestion of Stan Ulam's, he postulated a world consisting of a two-dimensional latticework of abstract "states," that change at discrete times according to a deterministic rule that depends only on the value of the neighboring states. This interaction rule may be thought of as defining the "physics" of a toy universe. Such a set of discrete states, together with a rule that changes them based of the states of their neighbors, is called a *cellular automaton*. In this world he demonstrated that there was a particular configuration of states with the capability to reproduce itself. The resulting construction is complicated to describe in detail. Roughly speaking, he constructed an initial pattern that contained a description of itself. Because of the particular rules he chose for the toy universe, the information from this description could flow out through a "constructing arm" (also consisting entirely of abstract states) so that the organism could "build" a copy of itself.

A simpler example of a cellular automaton is the *game of life*.[10] Imagine a checkerboard. Each square is either "alive" (has a piece on it) or "dead" (empty). Each square has a neighborhood, defined as the eight adjoining squares. To make a "move" each square examines its neighbors in order to decide whether it will be alive or dead when the move is completed. If it is dead, and two or three of the squares in its neighborhood are alive, then after the move is completed, it is alive. If it is alive, and three of the squares in its neighborhood are alive, then after the move is completed, it is alive. Otherwise it is dead. This procedure is followed for each position with the pieces fixed in place, and then the positions are updated simultaneously.

This game is so simple that, unless you have seen it before, you may find it hard to believe that it can give rise to very complex structures. For example, there are "gliders," simple oscillating patterns that propagate across the game board; "glider guns," which periodically emit gliders; and "self-reproducing glider guns," which make glider guns.

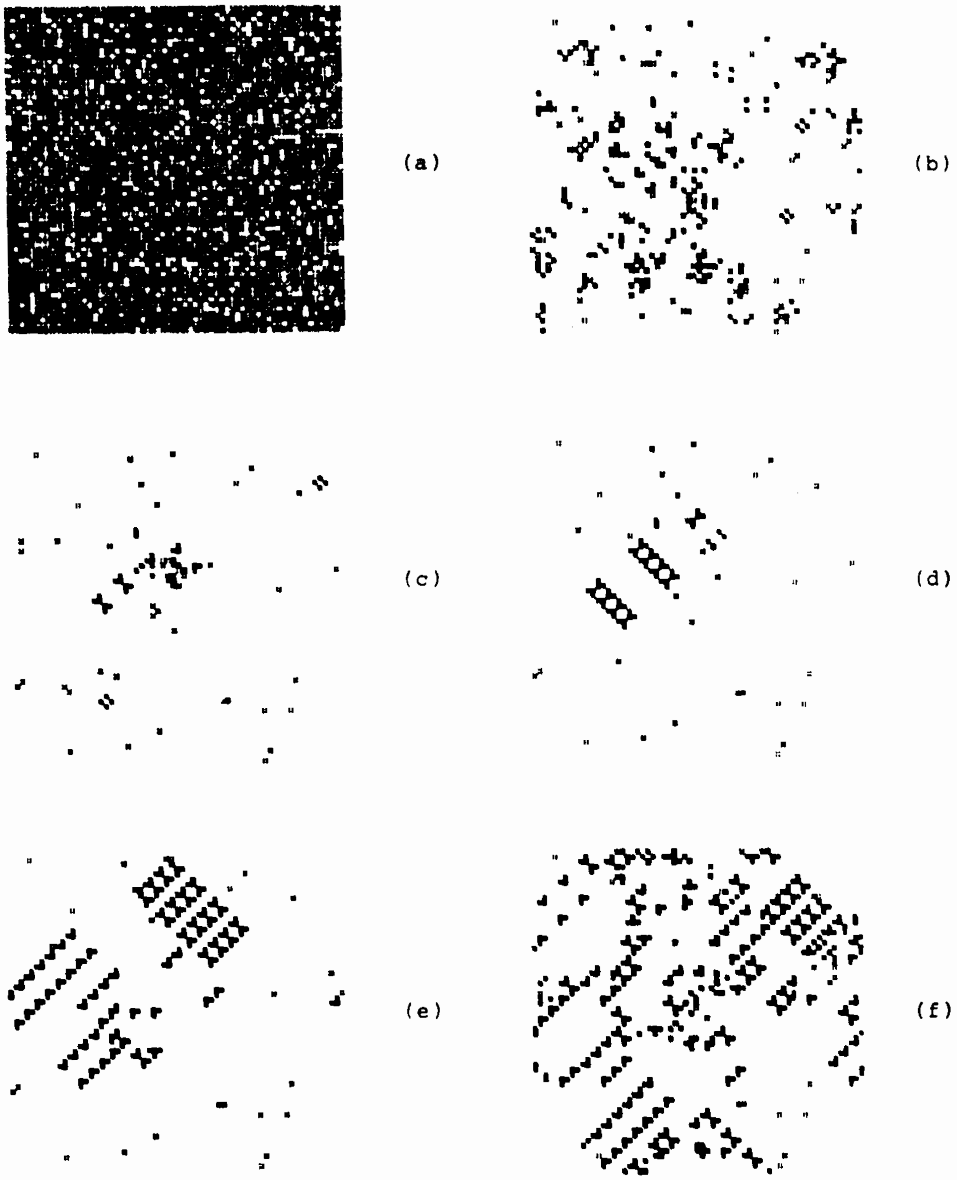

FIGURE 1 A cellular automaton with robust self-reproducing patterns, discovered by Chris Langton. Figure (a) shows a random initial condition on a square lattice; the eight possible states are represented by distinct patterns of dots. As time evolves the density of blank states increases, as shown in (b) and (c). In (c) we already see the seeds of self-reproducing patterns; as time progresses these patterns grow by replicating themselves. There are several reproducing patterns, which compete with each other for space, as shown in (d), (e), and (f).

Like the contemporary computer viruses, the self-reproducing objects in the game of life are not very stable. A small perturbation in their patterns typically destroys the replicating structures. Furthermore, if the game of life is run from a random initial condition, it typically settles down into static or simple periodic configurations. There are, however, other cellular automaton rules that are similar to the game of life, for which self-replicating structures seem to be quite robust. One set of examples, recently discovered by Chris Langton, is shown in Figure 1.

Langton has also made models for the formation of colonies,[13] as shown in Figure 2.

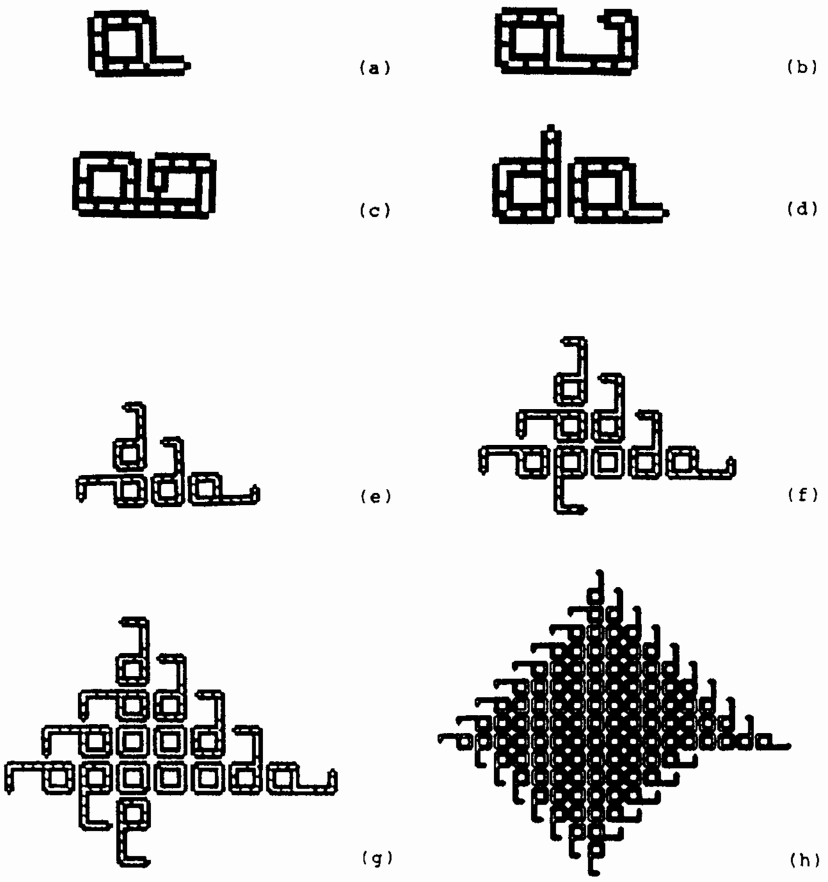

FIGURE 2 A cellular automaton model of self-reproduction, due to Chris Langton. Signals propagating around the "Adam" loop (a) cause the short arm to grow and curl back on itself (b,c,d), producing an offspring loop (e). Each loop then goes on to produce further offspring, which also reproduce (f). This 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.

An initial pattern reproduces itself on adjacent squares, in a manner reminiscent of the growth of a coral reef.

Along a somewhat different line, Richard Dawkins has created simple forms called "biomorphs" that evolve under artificial selection.[5] The "breeder" begins with a random pattern of lines connected to form a "tree." The geometric pattern of these lines is specified by simple rules, the details of which are given by simple abstract "genes." These rules are recursive, i.e., their form is the same at every level of the tree, and they can be applied to themselves. The biomorphs reproduce by making copies of themselves which differ from each other due to random mutation of their genes. The breeder selects the biomorphs he or she finds pleasing, and lets them breed again. In only a few steps, it is possible to create forms that are reminiscent of many different organisms. A few examples are shown in Figure 3.

The ease with which specific biomorphs are created illustrates the importance of recursive operations in generating evolvable biological forms.

Do these worlds give rise to life? So far, with the possible exception of the copy/split organisms of VENUS, or the robust self-reproducing automata of Langton, we would have to say that the answer is probably no. The key problem is finding the right combination of stability and variability. In most of the examples above, the self-reproducing patterns are destroyed by the slightest change. They are so fragile that they have difficulty evolving beyond their initial form. The robust replicating structures in both VENUS and Langton's automata are robust, but so far they have not been able to evolve beyond a fairly simple level of complexity. Discovering how to make such self-reproducing patterns more robust so that they evolve to increasingly more complex states is probably the central problem in the study of artificial life.

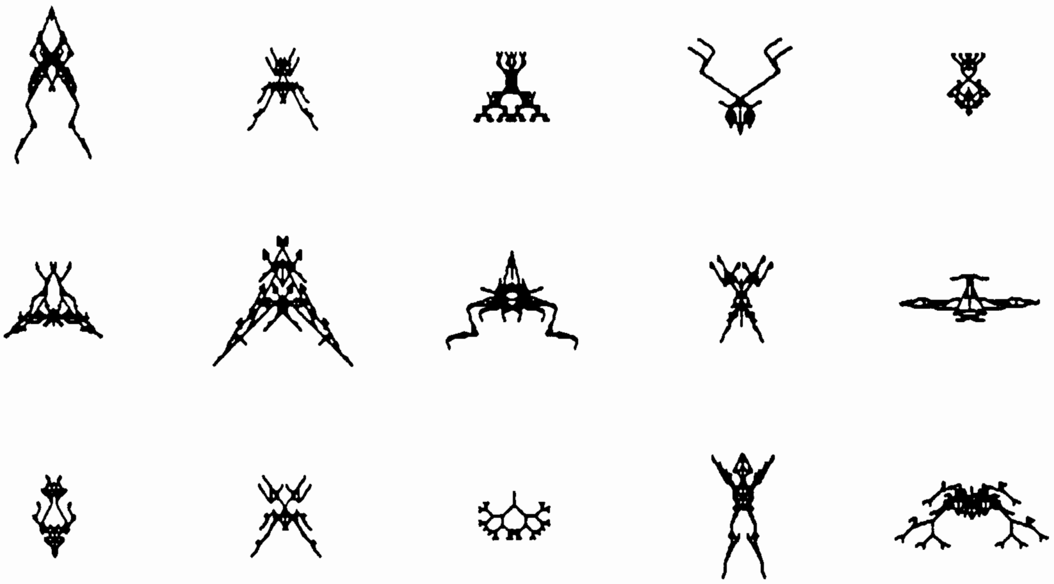

FIGURE 3 "Biomorphs," created through random variation of a simple "genome" and artificial selection of desirable features.

2.2.3 GENETIC ENGINEERING AND ARTIFICIAL WETWARE To highlight the contrast with carbon-based naturally occurring life-forms, in the discussion so far we have mainly addressed silicon-based artificial life-forms. However, artificial life can also occur in the wet, carbon-based medium of contemporary organisms.

The preponderance of examples of contemporary life-forms are bags of mostly water, built out of proteins, nucleic acids, lipids, and other organic compounds. The genome containing a self-representation is a DNA molecule. It is essentially a book, an instruction manual for the construction and operation of an organism. The message is written in an alphabet consisting of four letters, corresponding to four nucleotide molecules. The detailed message that distinguishes one organism from another is contained in the sequence of nucleotides along the DNA chain. The machinery of the cell, consisting of proteins, lipids, etc., reads this message and constructs replicas of the cell, much as does von Neumann's automaton.

In a certain sense carbon-based artificial life-forms have been with us since the advent of animal husbandry. By circumventing natural selection and replacing it with artificial selection, we alter the genome, creating varieties and hybrids that would never exist in the natural world. Nonetheless, artificial breeding is a comparatively weak tool, and the plant and animal forms it has produced are all relatively close to naturally occurring forms. The techniques of modern biochemistry and genetic engineering promise to take us far beyond this, giving us much more control over the genome, and the potential ability to create artificial life-forms that are radically different from natural life-forms.

There are two paths for the emergence of artificial life in the organic medium. The first path, which has already produced a variety of artificial life-forms, is genetic engineering. By directly manipulating the genome, we can modify existing life-forms. The second path, "artificial wetware," returns to the most primitive level, attempting to recreate the origin of life, or perhaps to generate whole new roots of the evolutionary tree.

Genetic Engineering. Under normal circumstances the message contained in a DNA molecule is invisible to us, and can only be read by the cell itself. In recent years, however, we have acquired the ability to peer into the cell and translate the sequence of nucleotides into sequences of human-readable symbols, A, G, C, and T. It is quite possible that we will be able to sequence the entire human genome, which consists of more than a billion nucleotides, by the year 2000. In other words, we will be able to read off the entire sequence of letters comprising the message that defines a particular human being. There is still a large step before we can *understand* what this information actually means and anticipate the effect of making changes in the sequence; at the moment the language is much more foreign than any human language, with semantics that lie in an entirely different realm. When we acquire the ability to interpret the messages of the genome, we will be able to "design" living things, change their form, cure them of hereditary diseases, make them bigger or smaller, or more or less intelligent. We will be able to create new species with properties radically different from those of natural organisms.

There are many potential commercial applications for genetic engineering. Bacteria have been genetically engineered to perform a variety of useful tasks. For example, the “*Ice*$^{-}$” bacterium protects plants against damage from freezing. Other bacteria have been designed so that they increase nitrogen fixation in plants or help clean up hazardous waste sites. Genetic engineering of fungi promises to improve industrial production of antibiotics and other useful chemicals. Plants have been genetically engineered so that they are resistant to infectious agents, or produce more and better food. Pigs have been genetically engineered to produce better meat, by elevating their level of bovine growth hormone, making them more like cattle. This is an example of how delicate success in genetic engineering can be—these artificial pigs also acquired a variety of unacceptable health problems, making them unviable freaks.

Some applications to humans are already in place. A particularly promising one is “gene therapy,” in which defective or mutant genes are fixed by genetically engineered viruses that either replace or supplement the defective genetic information. This gives us the potential to cure many disorders of the bone marrow, liver, central nervous system, some kinds of cancer, and hormone imbalances. Gene therapy only involves the reproductive machinery of the cell, and not that of the whole organism, so that changes are not transmitted to the offspring. Another technique, called homologous recombination, makes it possible to replace a defective gene in the reproductive cells, forever altering future generations.

Through techniques such as homologous recombination, we have the capacity to change the human species by eliminating deleterious genes. The evolution of human beings thus comes under conscious human control. Initially these changes will be minor adjustments, such as the elimination of diabetes and other genetic diseases. As we acquire more knowledge of the function and interpretation of the code, we will also acquire more capability to *add* new features, as we already do for bacteria, plants, and even some mammals. Should we choose to exercise the option of making such changes, we may give rise to “human” beings that are quite different from current homo sapiens.

Making alternative organic life-forms from scratch. The quest to discover the origins of life has led to a great deal of speculation about the simplest possible life-forms. Since our record of the earliest life-forms is extremely poor, it is generally agreed that the only experimental test is to recreate life “from scratch” in the laboratory. This forces us to seriously consider the issue of what it means to be alive, and also raises that possibility that, rather than recreating the origin of contemporary life, we might create a whole new evolutionary tree, whose material basis is quite different from that of contemporary life.

The basic building blocks of contemporary life are amino acids, which form proteins, and nucleic acids, which form DNA. As demonstrated by Miller and Orgel,[17] amino acids are easily synthesized under artificial conditions. It is more difficult to make proteins, although Sidney Fox,[9] has demonstrated that it is possible to make similar molecules called “protenoids,” which form bacterium-sized protenoid spheres with some suggestively lifelike properties. Nucleotides can also be formed in the laboratory, providing the proper protein enzymes are present.

However, there are still several crucial problems that remain to be solved before we will be able to directly recreate lifelike behavior from non-living material.

---

3. HOW AND WHEN?

Whether or not we study it as a scientific pursuit, we suspect artificial life will emerge in one form or another for economic reasons. We feel that this is unavoidable because of the economic incentives. The timetable and detailed mechanisms are are still uncertain, but the imperative is quite clear. In any case the implications for our civilization and ecosystem are dramatic.

3.1 HOW WILL IT HAPPEN?

We feel that artificial life will emerge gradually, slowly becoming a part of our day-to-day lives. There are many possible avenues for this; probably many of them will be explored simultaneously. True artificial life will be preceded by a series of stages in which we come closer and closer to the real thing.

It is often the case that technological developments are anticipated, at least in spirit, in speculative fiction. Just as Jules Verne anticipated much of the technology of the twentieth century, many aspects of artificial life that may appear in the twenty-first century and beyond have been anticipated in the fiction and nonfiction of this century. The possibility of self-reproducing machines was anticipated as early as 1929 by J. D. Bernal.[2] More recently the poet laureate of artificial life, Stanislas Lem, has written many books that are populated by a variety of artificial life-forms, and that suggest how and why they might arise. For example, in his recent book *Fiasco*,[16] a group of space explorers searching for life travels to what seems from a distance to be a planet with a ring around it, similar to Saturn. However, on closer inspection the ring turns out to be composed of attack satellites and anti-missile weapons. It originally began as a "star wars" defense shield against land-based nuclear attack. As each side learned to jam the operations of the others' technology, more and more autonomous control was given to the satellites. Since material was difficult to transport into space, they made them self-reproducing. The ring evolved and developed into an ecology of hostile, autonomous organisms, beyond the control of the parental planet. Unfortunately, in view of modern developments, this scenario is all too believable.

More peaceful applications are already beyond the realm of science fiction. For example, NASA sponsored a summer study group to investigate the feasibility of making self-reproducing aluminum mining modules on the moon.[12] The purpose was to design aluminum mining machines capable of mining aluminum, making copies of themselves, and catapulting aluminum into a near-zero gravity orbit between the earth and the moon where it can be used to build a space station. The machines use the aluminum they mine to manufacture replacement parts. Although the initial investment would be large, once the seed machinery is in place, because of the ability to reproduce, the amplification of the initial investment is almost unlimited. The NASA study concluded that this could be accomplished by placing only 100 metric tons of material on the surface of the moon.

Outer space provides a favorable medium for artificial life. Although the conditions in space are hostile to biological organisms, machines do not breathe oxygen, do not require water, are naturally powered by solar energy, and elegantly driven by "solar sails," which employ the solar wind as a motive force. Machines thrive where humans perish. If we ever wish to explore the solar system and make use of the tremendous natural resources that exist outside of earth, self-reproducing machines provide the natural way to accomplish this task. Because of the enormous potential economic returns, self-reproducing machines are likely to emerge as the natural tool for space exploration.

The emergence of artificial life will probably have antecedents on earth that are not as dramatic as self-reproducing aluminum mining modules in outer space. Indeed, we are already coming close to such possibilities. The Macintosh computer, for example, is produced in factories with virtually no human intervention, machines producing other machines. Microchip fabrication is under increasing levels of computer control, from the layout of printed circuit boards to etching of the actual chips. As computers become more sophisticated and more integrated into our lives, and as we become more dependent on them, they will exert more control on us and on themselves. We have already discussed how computer networks form an "agar," fostering the formation of computer viruses. It seems that whenever there is a medium capable of supporting large amounts of specific information, organizational patterns emerge that propagate themselves by taking over the resources of this medium. As our society becomes increasingly information intensive, it automatically acquires increasing potential to support artificial life-forms.

Carbon-based artificial organisms are already a reality. At this point they do not play a major role in our lives, but then genetic engineering is a new technology whose potential has only begun to be explored. As the power of this technology develops, we will inevitably come to rely more and more heavily on genetic engineering to face the problems caused by overpopulation and the limits of our resources. It is only a question of time before we begin to apply genetic engineering to human beings. Elimination of genetic-related diseases will probably occur without a great deal of controversy. But once this is accepted, more controversial measures will begin to be considered. Some changes, while potentially desirable for society, may be very difficult to bring about. For example, we could use genetic engineering to make human beings smaller. Small people take up less space and consume fewer resources, and if we were all significantly smaller, we could support the same number of people and place far less strain on our planet. Nonetheless, who would be the first to volunteer?

A critical point will occur when we acquire the ability to modify the intelligence of our offspring. If this can be done simply and reliably, there will probably be many volunteers. Although the political and social difficulties may be substantial, as our society becomes increasingly complex, the demand for increased intelligence will grow. In a relatively short amount of time, we may find "human" beings that are quite different from current homo sapiens, new generations of men and women as anticipated by Stapleton.[23]

3.2 WHEN WILL IT HAPPEN?

The easy answer to this question is that it has already happened. Computer viruses and genetic engineering are a reality, a tangible demonstration that artificial life is not only a subject for science fiction. However, neither of these are self-sufficient life-forms; both computer viruses and genetically engineered life-forms require human beings to create them. This does not say that they are not alive—there are many natural organisms that cannot exist without other organisms. It merely says that their evolutionary development depends on symbiotic relationships with other parts of the ecosphere.

Before artificial life is achieved on a broader scale, so that it contains all the rich possibilities of natural life, there are still technological developments that need to occur. These developments are significantly different in detail for carbon-based and silicon-based organisms, although the general problems are related.

For carbon-based artificial life-forms, we need a much more comprehensive and efficient capability to read and alter the the genome. Sequencing or "reading" a genome is currently a very labor-intensive task. We have complete sequences for only a few of the most primitive organisms. Nonetheless, technological developments in this area are relatively easy to anticipate, and it seems likely that in the twenty-first century we will be able to read large genomes relatively easily. Similarly, techniques for manipulating the genome, i.e., making specific alterations in the sequence of nucleotides, are developing at a rapid pace, and we can expect that in the twenty-first century this will be a relatively easy matter.

The real limiting factor to the development of carbon-based artificial life is *understanding* the language of the genome, so that we can anticipate the effect of making a given change. This is complicated by the fact that genes do not act independently—their actions are highly dependent on those of other genes. This interdependence makes it very difficult to anticipate the effect of a given change. Solving this problem requires a much more complete understanding of how a living organism functions.

For computer-based life-forms, the needed developments are naturally divided into two areas: hardware and software. Of these, the development of hardware, the raw computational machinery, is much easier to predict. The development of software is analogous to understanding the language of the genome—we need fundamental breakthroughs and a comprehensive understanding, and its development is much more difficult to predict.

We will first examine the development of hardware: Since the advent of computers, our ability to compute has increased at a steady exponential rate. Up until now computational power has increased by a factor of roughly 1000 every 20 years. This implies that by about the year 2030, if we follow the same growth curve, we will have computer hardware roughly a million times as powerful as that we possess now.[19] At this point, we will have computers whose power is roughly comparable to that of the human brain.

It is, of course, difficult to compare the power of the human brain to the power of a computer. Their capabilities are quite different. Roughly speaking, though, the raw hardware power of the human brain can be estimated in terms of the number of neurons and their speed as computational elements. These figures are not known with any precision, but a ball park figure places the number of neurons at 10^{10}, the switching speed of a neuron at 100 bits per second, and the storage capacity at 100 bits per neuron. Using this estimate, and extrapolating the rate of growth of computer technology, we can expect that by about the year 2025, we will have computers with roughly the computational power of a human brain. Our estimate may easily be wrong by a factor of 1000, but as long as the available computational power grows exponentially, this makes only a very small difference in the time for the hardware potential of artificial computers to reach equivalence with the human brain. Even if the estimates of the power of the human brain are off by a factor of 1000, the crossover point still shifts by only 20 years.[3]

In any case, the complexity of the human brain is probably more than that needed for life. The "hardware" that makes up a simple bacterium is certainly far less complex than that of the human brain. Its true complexity is difficult to estimate, but it is quite possible that contemporary computers already have enough hardware power to simulate the essential information-processing functions of a bacterium.

The time for the emergence of software is more difficult to assess, and places a more severe limit on the emergence of artificial life than the development of sufficient hardware. Conventional computer languages and computer programs follow very different principles than those of the brain or of the machinery that controls the cell. The underlying principles behind biological organisms are robust and adaptable. New approaches evolve spontaneously, without conscious intervention. In contrast, conventional computer programs are not robust; they are easily broken by small changes. Spontaneous evolution is difficult.

To create artificial computer-based life that is robust, which can survive fluctuations in its environment and evolve as freely as biological life, we must solve several fundamental problems in the design of computer software. We must make software that is ***adaptable***, with learning algorithms that allow computer programs to profit from experience. Ultimately we need computer programs capable of writing other computer programs, with "goal-seeking" behavior that allows programs to function in ill-specified environments. We need computer software that can innovate, and

---

[3] See Hans Moravec[19] for more detailed treatments of these issues.

add onto itself in response to its "needs." Solving these problems is one of the fundamental goals in the study of artificial life. These are also central problems in the related field of artificial intelligence.

New approaches to artificial intelligence include computer programs that mimic aspects of real biological neurons,[21] and computer programs that alter themselves through "genetic" manipulations very much like those employed by our reproductive machinery.[11] However, we are still lacking several principles needed to build living systems. It is unclear at this stage whether all that is needed are a few broad fundamental theoretical breakthroughs, or whether we still face a long trail of piecemeal and highly specialized discoveries. In the latter case, the timetable for the broad emergence of robust artificial life-forms might be extended significantly.

The advent of computer viruses illustrates the immediacy of artificial computer-based life. Although contemporary computer viruses are not very robust in the face of changes in their programs, they can nonetheless be quite long lived. We believe that the ability to make stable, self-reproducing artificial life-forms only awaits a few conceptual breakthroughs. In this case, artificial life might fully emerge by the middle of the next century.

Note that the development of carbon-based and computer-based life-forms are highly complementary processes. The technology for sequencing and manipulating the genome is highly dependent on computers and developments in computer-based artificial intelligence. Developing an understanding of the language of the genome is likely to be highly dependent on increasingly more sophisticated computer simulations of the functioning of organisms. In turn, this understanding is likely to guide us in developing the principles for computer-based artificial life. And eventually, genetic engineering of more intelligent humans is likely to have an impact on all of these problems.

4. THE BIG PICTURE

4.1 EVOLUTION AND SELF-ORGANIZATION

We are accustomed to thinking of evolution as an explicitly Darwinian phenomenon, specific to biological organisms, involving competing processes of random mutation and natural selection. However, it is possible to take the broader view that biological evolution is just one example of the tendency of matter to organize itself as long as the proper conditions prevail.

This concept of evolution was originally introduced by Herbert Spencer in the mid-nineteenth century.[22] He defined evolution as "a change from an incoherent homogeneity to a coherent heterogeneity." According to Spencer, evolution is a process giving rise to increasing differentiation (specialization of functions) and integration (mutual interdependence and coordination of function of the structurally differentiated parts). He viewed evolution as the dominant force driving the spontaneous formation of structure in the universe, including the formation of matter, stars, geological formations, biological species, and social organizations. Thus, Darwinian evolution is just a special case of a broader principle.

In Spencer's view, evolution is the antagonist of dissolution. His notion of dissolution is essentially what physicists call the second law of thermodynamics. According to the second law, disorder, or *entropy*, tends to increase in the absence of an input of energy. This is an embodiment of the familiar principle that it is easier to make a mess than to clean it up. In nature organized forms of energy such as light or the bulk motion of matter tend to turn into disorganized energy (heat), i.e., disordered atomic motion.

When organized energy streams down onto earth, much of it simply turns into disorder, in the form of heat. However, something else also happens, which seems to be quite the opposite: Processes of differentiation cause oceans, clouds, wind, rain, and geologic formations. These are processes of *organization* rather than *disorganization*. It is not that they disobey the second law of thermodynamics, but rather that the second law does not tell the full story. While there is an overall net increase of disorganization at the molecular level, at higher levels, under favorable circumstances there is an inexorable tendency for an increase of order. Life is, of course, the primary example.

The theory of organization is much less developed than the theory of disorganization. We have a precise formulation of the second law, but at this point, there are no good general theories for self-organization. In its broadest sense, the study of artificial life is an avenue that can help us make a broader theory of evolution more precise. By producing tangible examples of self-organization in simple mathematical models, we hope to understand why nature has an inexorable tendency to organize itself, and to discover the laws under which this process operates.

4.2 LAMARCKIAN *VS.* DARWINIAN EVOLUTION

Viewed in a broad context, the advent of artificial life is significant because it signals the possibility of a major change in the manner in which evolution as a whole takes place. The first such change probably occurred with the creation of the first self-reproducing organisms. Before this the spontaneous formation of structure relied on more indirect processes of self-organization. With self-reproduction it became possible to directly transmit information and patterns from the past to the future. It also made it possible to incrementally change this structure through Darwinian evolution, a process of random mutation and natural selection. Under that process, small changes take place during the process of reproduction, producing random variations in the offspring. If these changes are not favorable then the offspring may die out. If these changes are favorable, however, then the offspring reproduce more frequently, passing these changes on so that they propagate. *Only* genetic information is transmitted. Acquired characteristics, such as good muscles developed through exercise or the wisdom acquired in one's lifetime, cannot be transmitted to subsequent generations directly. Darwinian evolution is the fundamental mechanism that has designed the flora and fauna of earth. Viewed in the broad sense of Spencer, self-reproduction provided a new mechanism for evolution, signaling a major speedup in the rate at which evolution as a whole took place.

An alternate mechanism of biological evolution was postulated by Lamarck. He believed in the transmission of acquired characteristics to subsequent generations. He believed, for example, that if a giraffe stretched its neck and made it longer, then its offspring would have longer necks. We now know that this is not true for biological organisms. However, there is an important context in which it is true: the evolution of culture.

A culture can be viewed as a kind of organism built out of individuals and social units. New ideas in a culture compete for prominence within the culture. These ideas propagate through our modes of communication, largely language and writing. Ideas and their concomitant modes of behavior are selected according to their usefulness to the society, and cultures evolve through the course of time. More successful ideas supplant other ideas, mimicking the survival of the fittest that we associate with biological evolution. In contrast to biological evolution, in social evolution acquired characteristics are passed on to subsequent generations. Cultural evolution is essentially a Lamarckian process.[4]

The capability of cultural evolution for bringing about effective change on a timescale far faster than biological evolution demonstrates the power of Lamarckian evolution. Although cultural evolution occurs on a limited basis in other organisms, such as monkeys and birds, its true potential has been manifested only in humans. The emergence of language, with the attendant amplification in the ability to transmit cultural information, wrought an enormous change in collective human behavior. In a very short period of time, perhaps only fifty thousand years, human culture has given us the ability to send people to the moon, to destroy life on our planet, and to create life. The pace of cultural evolution is strikingly fast when compared to the much slower pace of biological evolution. This is not surprising; change happens much more efficiently when acquired characteristics can be transmitted directly, and when innovation comes as a result of conscious design rather than random guessing.

Viewed in the broad terms of Spencer, the introduction of culture, with its more rapid mechanism of Lamarckian evolution, can be viewed as a watershed event in the history of evolution as a whole—a "phase change" accelerating the global evolutionary process. However, in the absence of biological change the scope and possibilities of cultural change are limited. The human brain is limited in its ability to assimilate the vast quantities of information generated by our culture. Increasingly we turn to tools made specially to help us in these tasks, computer memories that are capable of storing information much more efficiently than we can ourselves. These tools are gradually becoming much more than passive memories, actively performing many of the functions that we would otherwise perform.

Artificial life provides the possibility for Lamarckian evolution to act on *the material composition of the organisms themselves.* Once we can manipulate the

---

[4] In analogy with genes, Dawkins has characterized the fundamental units of cultural evolution as *memes*.[4]

genome directly, once we understand how the genome is built and can anticipate the effects of changing it, we can modify our offspring according to our perception of their needs. This is true for both the silicon-based genomes of computing machines and the carbon-based genomes of genetically engineered biological organisms. Unlike the original concept of Lamarck, this does not happen automatically, but rather through the intermediary of consciousness. The giraffe's longer neck is not automatically passed on as a result of stretching. Instead, the giraffe realizes that it would be nice if its offspring could have longer necks, and does appropriate genetic engineering to make this happen.

In artificial computer-based life-forms, the genetic material will almost certainly be under direct control, in computer readable and easily modifiable form. Initially, of course, such organisms may not be very smart. The most likely event is that the genomes will be modified by humans, to effect a good design for some commercial purpose. We then have a *symbiotic* Lamarckian evolution, in which one species modifies the genome of another, genetically engineering it for the mutual advantage of both. In a sense, this is what we have done all along with our technology—automobiles, for example, “evolve” as we manipulate their genomes (blueprints). With artificial life there is the potential for the control of the genome to be given to the products of our technology, thus creating self-modifying, autonomous tools. As artificial life-forms achieve higher levels of intelligence, the ability to modify their own genomes will become increasingly more feasible.

Assuming that artificial life-forms become dominant in the far future, this transition to Lamarckian evolution of hardware will enact another major change in the global rate of evolution, comparable to the enormous acceleration that occurred with the advent of culture. The distinction between artificial and natural will disappear. This will be a landmark event in the history of the earth, and possibly the entire universe.

5. THE CONSEQUENCES FOR HUMANITY

The study of artificial life may potentially answer some very important questions in biology and the theory of evolution. It also provides a tool to address some of the most fundamental philosophical questions, such as: What does it mean to be alive? What are the underlying physical and mathematical processes that give rise to life? How does nature spontaneously create order from chaos? What are the mechanisms of creativity and self-organization?

Artificial life-forms will probably emerge whether or not we choose to study artificial life as a scientific discipline. Artificial life-forms have the capacity to evolve beyond contemporary life. At first, they will be quite unsophisticated, simple tools that we have built to satisfy our needs. Ultimately, however, economic and political pressures will drive artificial life-forms to greater degrees of sophistication, until their complexity and information-processing capabilities are comparable or superior to those of humans. This may engender competition with humans.

What should our attitude be? It is natural to fear the unknown, particularly when it involves a possible threat to our species. It is easy to imagine nightmare scenarios in which cold, malevolent machines or vicious genetically engineered creatures overwhelm humanity. Viewed in this way, artificial life becomes a threat to our survival to which we must respond, something that must be eliminated so that human beings can continue to prosper without competition.

We should, however, use care before automatically taking such a view. In the challenges issued for this symposium, Murray Gell-Mann has asked us to address the dangers of "human tribalism." Humanity has traditionally been self-centered, eager to exalt itself and to regard itself as the sublime creation of God, squarely in the center of the universe for the rest of time. We have now evolved somewhat away from this narcissistic view. We now know that we are the inhabitants of an average planet orbiting an average star in an average galaxy. We may also surmise that this moment in cosmic history was arrived at through an evolutionary process of change which will replace us at the next moment.

The natural order of evolution is change. No species has persisted forever. Individual species are altered and replaced through an evolutionary process of modification and succession that continually alters the composition of the flora and fauna of earth. There is no reason to believe that we are immune to this. It seems quite natural that we, too, will evolve and change with the passage of time, giving rise to new species in the genus homo. With artificial life this evolutionary change may not follow such a continuous path; although we give rise to new species, they may be our own direct conscious creations and radically different in form from ourselves.

Another topic that Murray asked us to address in this symposium is the "preservation of cultural and biological diversity." We now have the possibility to *create* cultural and biological diversity. With the advent of artificial life, *we may be the first species to create its own successors.* What will these successors be like? If we fail in our task as creators, they may indeed be cold and malevolent. However, if we succeed, they may be glorious, enlightened creatures that far surpass us in their intelligence and wisdom. It is quite possible that, when the conscious beings of the future look back on this era, we will be most noteworthy not in and of ourselves but rather for what we gave rise to. Artificial life is potentially the most beautiful creation of humanity. To shun artificial life without deeper consideration reflects a shallow anthropocentrism.

But the path is fraught with danger. Short-sighted fear and hatred all too often dominate the activities of human beings. At the outset at least, we will shape the form and innate drives of artificial organisms. A particularly frightening scenario comes from the potential military uses of artificial life. There are many military applications for which artificial life-forms would be extremely useful, from battlefield robots to satellite warfare. We can only hope that we have the collective wisdom to make treaties and suppress such applications before they occur. As we have seen with nuclear weapons, political forces make the consensus necessary to dismantle existing weapons systems extremely difficult to achieve. Once self-reproducing war machines are in place, even if we should change our mind and establish a consensus, dismantling them may become impossible—they may be literally out of our control. An escalated technological war involving the construction of artificial armies would likely end by destroying the participants themselves, and would give rise to a generation of life-forms that might be even more hostile and destructive than their human ancestors.

Artificial life-forms will be shaped by the forces that create them. If, instead of building war machines, we use our technology for productive purposes, they may bring us a wealth of resources that will greatly enhance our well-being. Ultimately, they may evolve to be far more intelligent than humans, and capable of intellectual feats that we cannot even dream of. Such intelligence might result in enlightened behavior that is inconceivable to lower forms of life such as us.

If we can shape artificial life in a positive direction, the bittersweet consequences to humanity can be visualized by analogy to Arthur C. Clarke's book, *Childhood's End*.[3] In this story he imagines that the children on earth acquire the ability of mental telepathy. This ability makes them into an enlightened race whose collective powers are far greater than those of ordinary humans. However, as a result, they are so beyond their parents that they become strangers to them. Their parents are left with feelings of glory for the harmony and greatness of what they see their children will accomplish, but simultaneously they feel sadness that they cannot participate.

In discussing artificial life here we have been intentionally provocative. Our vision of the future may not be accurate. We hope that, whether or not you agree with us, you will be stimulated to address this issue. If we are right the advent of artificial life is the greatest challenge facing humanity, an inevitability that we must shape and set in motion in the proper direction. If the future is to do justice to the nobler attributes of humanity, then we must take positive action.

BIZARRO By DAN PIRARO

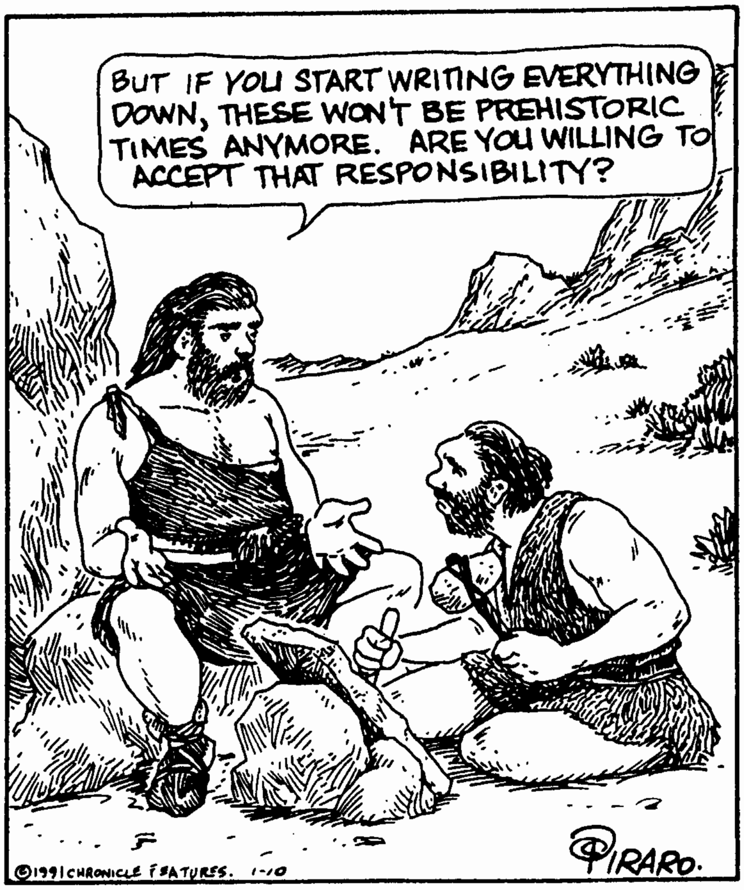

ACKNOWLEDGMENTS

We appreciate valuable comments and criticism from Christian Burks, Walter Fontana, John Holland, Chris Langton, Norman Packard, and Steen Rasmussen.

The "Bizarro" cartoon by Dan Piraro is reprinted by permission of Chronicle Features, San Francisco, CA.

REFERENCES

1. Bagley, R. J., J. D. Farmer, S. A. Kauffman, N. H. Packard, A. S. Perelson, and I. M. Stadnyk. *Modeling Adaptive Biological Systems.* Technical Report LA-UR-89-571, Los Alamos National Laboratory, 1989. To appear in *Biosystems.*
2. Bernal, J. D. "title." In *The World, The Flesh, and The Devil,* edited by E. P. Dutton. 1929.
3. Clarke, A. C. *Childhood's End.* San Diego: Harcourt Brace Jovanovich, 1963.
4. Dawkins, R. *The Selfish Gene.* Oxford: Oxford University Press, 1976.
5. Dawkins, R. "The Evolution of Evolvability." In *Artificial Life,* edited by C. Langton, 201. Santa Fe Institute Studies in the Sciences of Complexity Proc. Vol. VI. Redwood City, CA: Addison-Wesley, 1989.
6. Dewdney, A. K. "Computer Recreations: In the Game Called Core War Hostile Programs Engage in a Battle of Bits." *Sci. Amer.* **250(5)** (1984): 14–22.
7. Farmer, J. D., S. A. Kauffman, and N. H. Packard. "Autocatalytic Replication of Polymers." *Physica D* **22** (1986): 50–67.
8. Farmer, J. D. "A Rosetta Stone for Connectionism." *Physica D* **42** (1990): 153–187.
9. Fox, S. W., K. Harada, and J. Kendrick. "Production of Spherules from Protenoids and Hot Water." *Science* (1959): 129.
10. Gardner, M. "Mathematical Games: The Fantastic Combinations of John Conway's New Soliaire Game 'Life.'" *Sci. Amer.* **223(4)** (1970): 120–123.
11. Holland, J. H. "Escaping Brittleness: The Possibilities of General Purpose Learning Algorithms Applied to Parallel Rule-Based Systems." In *Machine Learning II,* edited by R. S. Mishalski, J. G. Carbonell, and T. M. Mitchell, 593–623. Morgan-Kaufman, 1986.
12. Laing, R. "Replicating Systems Concepts: Self-Replicating Lunar Factory and Demonstration." In *Advanced Automation for Space Missions,* edited by R. Freitas and W. P. Gilbreath, 189–335. NASA Conference Publication 2255, 1982.
13. Langton, C. G. "Studying Artificial Life with Cellular Automata." *Physica D* **22** (1986): 120–149.
14. Langton, C. G. ed. *Artificial Life.* Santa Fe Institute Studies in the Sciences of Complexity Proc. Vol. VI. Redwood City, CA: Addison-Wesley, 1989.
15. Lem, S. "The Experiment." *The New Yorker,* 1978.
16. Lem, S. "title." *Fiasco.* San Diego: Harcourt Brace Jovanovich, 1986.
17. Miller, S. L., and L. E. Orgel. *Science* **30** (1959).
18. Monod, J. *Chance and Necessity.* New York: Knopf, 1971.
19. Moravec, H. *Mind Children: The Future of Robot and Human Intelligence.* Cambridge: Harvard University Press, 1988.
20. Rasmussen, S., R. Feldberg, M. Hindsholm, and C. Knudsen. *Core Evolution: Development of Assembler Automata in the Computer Memory.* Technical report, Los Alamos National Laboratory, 1989. To appear in *Physica D.*

21. Rummelhart, D., and J. McClelland. *Parallel Distributed Processing*, vol. I. Cambridge: MIT Press, 1986.
22. Spencer, H. *First Principles*. 1962.
23. Stapleton, O. *Last and First Men*. 1931.