The Garden in the Machine, Chapter One, by Clauss Emmeche, Princeton University Press (1994)

THE GAME OF LIFE

"IN FACT, it's quite easy to create life."

The little man with the dark eyes smiles. He has just summarized his lecture. A weak buzz from his computer confirms his conclusion. We can create life - and not only by reproducing our own flesh and blood; rather, we can create completely new life. The man is Thomas Ray from the University of Delaware. He does not have any supernatural or religious capacity. As a trained ecologist, Ray knows his biology and has a rather unsentimental, down-to-earth relation to living things. The kind of life Ray is talking about creating is biological life, but of a very special kind.

We are at the Technical University of Denmark, just outside Copenhagen, where Ray is visiting a research group that studies chaos theory and artificial life. A few years ago, Ray had renounced the study of tropical ecosystems, which had been his specialty. Ray was not bored, nor had he cynically realized that his object of study was soon going to disappear and that he had to find a new niche in the academic world. No, no matter how much Ray loved his rain forests, there was something else that had attracted his attention. He had gotten a remarkable idea: that with the help of computers he could piece together fragments of computer programs (i.e., instructions) and turn them into artificial organisms that did not just resemble life, but that theoretically speaking were just as alive as real animals and plants.

The idea seems a bit insane, or rather grandiose. Today, however, Ray can conclude that it is in fact "quite easy" to reproduce life. Just think of the computer viruses that plague our computer systems. The viruses are often created by hackers, usually teenagers with an intimate knowledge of their own and others' computers. They hack themselvesinto company or government computer networks and then release their homemade computer programs or viruses into these giant systems. These virulent pieces of computer programming reproduce themselves and then spread to all the computers in the network. Here they often make complete chaos of the most sophisticated systems and erase important data [2] Viruses are, in this sense, a form of life.

Ray was not the first to arrive at the idea of computer life. In 1989, when he began to work on creating organisms, scientists from many different disciplines were already testing the idea of creating life in artificial universes. A common thread began to emerge between them. This new life would be more fun, and it would be a more creative life than the destructive computer viruses that had wreaked havoc on several scientific centers. The project itself was to be interdisciplinary and scientifically ambitious, in the same way as in the 1950s when logicians, computer scientists, and psychologists came together in the attempt to construct "artificial intelligence."

A modest man, Ray entitled his project "An Approach to the Synthesis of Life," implying that there exist other possible methods as well. Life is a process, a complex, rhythmic pattern of matter and energy. What is important is not what kind of matter or what kind of energy we find, but rather the pattern, the process, the form. The computer's powers of information processing can imitate other forms and processing so well that the result is not simply an imitation or a theoretical image of life. The pulsating patterns on the computer screen are themselves new examples of how the process of life itself can take form, says Ray. Such self-developing patterns of processes are life.

There are an infinite number of processes in nature and in society where minute changes have profound, wide-ranging effects. If one removes the lowest can in a supermarket's tower of canned tuna fish, a large number of events occur quickly, one after the other; not only do the cans fall, but the tower itself vanishes. Physically speaking, it is easy enough to understand that a tower of canned tuna, which represents an organized system in an unstable balance, may quickly evolve into a more stable state: a pile of cans on the floor. This happens when the small change that takes place at the bottom of the tower ruptures the fragile symmetry and thereby converts the cans' potential energy into movement, heat, and a considerable amount of disorder. The cans come alive nearly all by themselves. But the cans' development in a supermarket does not have much to do with life in a biological sense. When we say that life is self-organizing, we refer to that which typifies a living organism (and the species' evolution), the organization that is built up and maintained rather than destroyed. Order does not disappear like the tower of cans in the supermarket; something is broken down along the way, to be sure, but order is continually recreated.

Ray discovered that he did not need to do much himself before all kinds of things began to show up on his computer. The events on the screen were more reminiscent of life's self-organization than of the spontaneous fall of a tower of cans. When Ray designed just one simple stem-organism, in the form of a program that he set off in his computer, it evolved into a veritable zoo of various species of descendants. Some were parasites (programs that utilize the computational resources of other programs). Others were hyperparasites (which exploited the parasites); and there also emerged more social organisms. All these organisms evolved at the cost of the machine's memory, but without Ray himself having to control their direction or their consumption of memory. Ray was startled to see the various new life-forms that appeared on the screen after he had released his homemade ancestral organism. I will later go into more detail about similar types of computer life.

That computers only do what they have been programmed to do has become a truism with modifications. All Ray did was to create a little "universe" with his machine, give it a few rules to work with, and then allow the system to evolve by itself. Just like Isaac Newton, who imagined that God, after having created the universe, the elements, and the eternal laws of motion, sat back and rested, without any further participation. Newton, however, did not regard the universe as being completely stable, so his God had to intervene from time to time in order to restore balance. Ray does not need to. Ray's microuniverse organizes itself, to a certain degree. Here it is not stability but evolution that really counts.

THE CASE OF THE FLICKERING GAME

In order to provide an initial, tangible notion of how artificial life behaves in a computer, it is easiest to examine a game called "Life," invented by the mathematician John Horton Conway in 1970. "Life" is so simple that it can be played on a chess board with a few pieces, pawns for instance. It is easier than chess, and the kinds of pieces used are in fact unimportant.

Each square on the board, regardless of color, is a "cell." A cell can have one of two states: "on" (when there is a piece on the square), or "off." Each cell has a total of eight neighboring cells (surrounding squares). The game is a solitary one; there is no opponent. It begins as follows:

Choose a beginning situation. For example, allow three cells in a row to be on and the remainder off. This is the beginning position. The game now proceeds step by step such that for each individual step, one calculates its next state (next generation) using two simple rules:

1. A cell is turned on if three of its neighbors are turned on
2. A cell remains on if two or three of its neighbors are also on; otherwise it is turned off.'

Conway formulated the rules slightly more organically and with more words:

3. Survival: each piece with two or three neighbors survives in the next generation
4. Death: each piece with four or more neighbors dies from overpopulation, and each piece with one or no neighbors dies of isolation
5. Birth: each empty cell with precisely three inhabited neighbor cells is a birth cell, where a new piece is born in the next generation.

The choice of words, however, is not important as long as the logic of the game remains the same.

For each cell, one counts the number of neighbors and uses the rules in order to determine its state in the next generation. When this has been done for all cells and the pieces are placed correctly, one has in reality "updated the cellular automaton" (for a moment, we can think of the game as a mechanical device, an automaton). In other words, one has computed the state of the cells on the entire playing board for the next time-step. One can thereafter continue to calculate the next generation, and the next, and so on until one decides to stop.

All cells in a generation change their state (i.e., are born, survive, or die) simultaneously. This means that one must not remove the pieces in a field that neighbors a cell if the states of the next generation have not yet been calculated. Conway recommended the following technique:

6. Start with a pattern of black pieces
7. Identify those who must die and lay a black piece on top of them
8. Find the birth cells and lay a white piece in each
9. Checking and double-checking, remove the dead pieces and replace the newborn pieces with black ones.

One soon discovers that depending on the initial state chosen (i.e., how many cells are on and how dispersed they are from each other) quite different things begin to happen. Sometimes the game dies out rapidly, and all the cells are turned off. Groups of cells can be frozen solid in a crystalline state, a stiff pattern that cannot be changed (unless "pushed" by new patterns). However, there may also appear "wave" sequences that are impossible to maintain within the boundaries of the chess board because the patterns grow; in this case more and more neighboring cells are born, and it takes increasingly more time to calculate the state of the entire next generation, the number of pieces is exhausted, and so on.

Conway quickly encountered problems because instead of a chess board he tried the game out in his home, on the checked tiles in the large entrance hallway, and he used dishes to mark the different cells. He quickly became tired of this when he literally began to step on the pieces. Luckily, the game can be played more effectively on a computer, and the computer can also calculate the generation change much more rapidly as well. Here the game really begins to resemble something living: small vibrating patterns go back and forth in wave-like rhythms, and the new structures emerge from simple seedlings. One of these new structures is a "glider," which crawls its way diagonally across the chess board until it crashes into other structures and allows itself to be swallowed up. A glider looks like this:

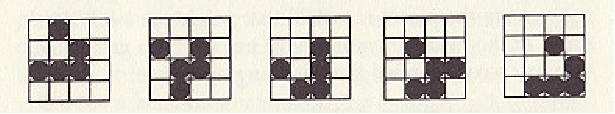

FIGURE 1.1. Microscopic section of John Conway's game "Life," in five consecutive generations, where a "glider" moves diagonally down over the network of cells. Along the way its form changes periodically as shown. White means "off" (empty field), black means "on" (a piece on the field). For each generation, a cell's state is determined according to the rules cited above. Only sixteen cells are shown but the edge cells also have eight neighbors, which cannot all be seen, and which are considered off. The active cell in the highest row in the first generation (to the left) has only one neighbor on; therefore it is off in the second generation, etc.

Conway and his colleagues gradually assembled a virtual zoo of forms that could persist, send signals (like a "glider gun"), keep blinking, or crawl around (more on this in chapter 3). Copies of this simple program soon began to circulate among computer enthusiasts, who themselves could explore this world and discover within it new life-form.

Conway's game is simple, but the structures that it generates can be quite complex. For followers of artificial life like Thomas Ray, the game of Life is a fully valid example of real life because it embodies the same principles of computation that are achieved in the more advanced forms of artificial life, including his own model.

SELF-ORGANIZATION: DARWINISM'S UNRESOLVED QUESTION

If life in a computer really is life and it organizes itself, it can perhaps fulfill an old dream of biology. Some of the researchers who in the infant days of biology observed the growth of plants and the fetal development of animals dreamed of one day being able to explain the riddle of the generation of form and regulation of growth. They dreamed of a "Newton of the grassblade" who could do for biology what Sir Isaac Newton had done for physics.' Next to nothing was understood about the mechanics of life and metabolism; there were no formulas for a simple blade of grass. The ingenious functionality of organisms and their ability to live in ecologically marginal areas and adapt to harsh conditions could be explained only by invoking some kind of divine design. The sublime intricacy of life-forms seemed to be sufficient proof of the Creator's magnificence. Behold His masterful work!

After Darwin, in the middle of the last century, created his theory of evolution based on natural selection and finally published it, it was not long before Ernst Haeckel, one of his supporters, declared Darwin to be the Newton of biology. According to Haeckel, Darwin had achieved the fundamental physical explanation for the multiplicity of life. The wealth of life-forms had evolved via a long, gradual selection process, carried forward via the individual organism's struggle for survival. The divine design explanation was now unnecessary.

Even though several decades went by before Darwin's theory of natural selection was accepted, it later became established as a fundamental biological law, at least from the neo-Darwinian point of view. In principle, the theory is still valid today as the natural explanation. The problem, however, is that we have realized that it is only half the explanation.

Darwin's evolutionary theory can be characterized as a theory about survival and gradual genetic change (toward increased fitness) among types of organisms. It is a theory about organisms that are *already* partially well adapted, well-constructed, and highly organized biophysical systems. What Darwinian evolutionary theory fails to explain is how animals and plants were created originally. Natural selection did not just leave the gigantic riddle about the origin of life unsolved; Darwin's theory also offered no insight about the creation of the individual multicellular organism in each new generation. Darwin knew nothing about self-organization. His theories were decisive for the subsequent development of biology, but the lack of understanding about the dynamics of the creation of a living structure left theoretical biology in a quandary. This became a major stumbling block for twentieth-century research. It is only within the last two decades that we have really begun to discover some of the principles for the creation of form, the creation of the complex from the simple, that must supplement Darwinian theory.

There are many paths to this new understanding. One of them is embryology itself, which took on the task of minutely sketching out the phases in the development of the vertebrate embryo. Another path is that of molecular biology, which studies the mutations that disturb the normal developmental sequences. The sequence that a fetus follows has been called the "epigenetic landscape," i.e., the total possible paths of cell development, growth, and anatomical creation of form. New insights have also come from physical theories of self-organization, such as the creation of ordered patterns in matter driven by energy flowing through the system (such as, e.g., hexagonal convection cells in a liquid heated from below to produce coherent motion of ensembles of molecules, the so-called "Bernard instability" phenomenon that organizes the system spatially). Finally, there is the contribution of computer science, both as a tool to simulate some of the complex morphogenetic processes and as a source of various metaphors for programmed development and coding. Ray and his colleagues, who claim to have created artificial life, are participants in a much larger process of intellectual history, a process with profound consequences for our way of viewing biology and life.

BETWEEN MYSTICISM AND BANALITY: FRACTAL KNOWLEDGE

The news that we can create artificial life sounds nearly magical. People program the material, and miraculously viruses, worms, and ants appear. We put form into the formless, pump sense into the senseless, let dust form into a system - do we now control the miracle of life? The idea is indeed dizzying, but the dizziness is due more to the challenge to our conventional ideas than to artificial life itself. Is it the conceptual gap between living and dead matter (and between controlled behavior and spontaneity) that is artificial, that gives us goose bumps when we suddenly look into a new space full of living forms that begin to fill the gap?

The creature on the computer screen, however, develops from a quite simple point of departure. There is no reason to mystify either artificiality or system-imposed spontaneity. What is mystical is not life presented as a biological phenomenon, nor life presented as a computer-based imitation. The mystical, rather, is the life we ourselves live - our own fascinations, aspirations, and conceptions, as well as the myriad of feelings that are aroused in us when we encounter the unknown, or when we find ourselves confronting the collapse of outdated concepts.

For years biochemists, physiologists, and geneticists have investigated the processes that are linked to the wet insides of living cells. Since its advent we have used the computer as an aid for the storage and manipulation of data. Yet until very recently, the idea that real life could be synthesized, let alone using the computer as a substratum, was a totally foreign notion. Whereas we had previously viewed organisms as machines, we were now not only faced with the strange idea that synthetic life could in principle be created, but that it could be created from the computer's silicon chips, entirely without carbon compounds, the basic elements found in all life that we know of on earth.

The mechanistic tradition persists in small alcoves of biology and periodically appears in certain researchers' extemporaneous philosophizing. The mechanistic biologist believes that an organism is in reality nothing more than a collection of atoms, a simple machine made of organic molecules. Mainstream biology, however, asserts that physics, despite the undeniable physical and material properties of organisms, is nevertheless inadequately equipped to explain them. An organism's structure and function are phenomena that exist in their own right, and must be described with biology's own conceptual apparatus. (This rather mundane biological idea has been given the imposing name "organicism.")'

Nevertheless, biologists have often employed a range of metaphors to describe the real nature of organisms, and the metaphors have typically been borrowed from the technology that happened to be most fashionable at the moment. An ant, for example, can be viewed as a mechanical piece of clockwork, with precise, finely tuned parts, each with its distinct function. From a subsequent perspective, the ant can be viewed as a piece of energy technology: a thermodynamic design that - in analogy to a steam engine - consumes chemically bound energy by combustion and performs work while developing heat. Today we might view the ant as a little computer with associated sensory and motor organs: it processes a mass of information about the external world and reacts by feeding back various responses.

Regardless of how one views living beings - in terms of mechanics, thermodynamics, as information-processing machines, as entirely unique types of systems in their own right, seeing that none of the existing metaphors as entirely satisfying - the emergence of artificial life nevertheless arouses within us a fundamental curiosity. It generates equally disturbing questions and among some people spawns an equally deep-seated intellectual and emotional resistance, as did research in artificial intelligence in its attempt to understand the human mind. If artificial life was simply a collective designation for a few computer models or a new method brought about by the cooperation of several natural-science disciplines, it would hardly have aroused the same virulent reactions. Why the undeniably curious sentiments about artificial life?

It has been said that science demystifies the world. It is closer to the truth to say that science, when it is at its best, opens the world up for us, bringing daily realities under a kind of magic spell and providing the means to see the limitations of what we think we know, and the scope of what we do not at all understand. Even though physics, chemistry, biology, and other sciences have helped to draw a grand map of reality, the blank spots on the map are never completely filled in. The more one draws, the more details must be investigated, and the more new horizons expand. Thus, knowledge has a kind of fractal structure: for every newly collected fragment of knowledge, we reach "farther down"; that is, we can with greater accuracy say more about nature and extend the universe of conversation about it and with it. However, we simultaneously extend our limits to the unknown. And if we restrict ourselves to the fragments, without trying to consolidate them into larger wholes, we succeed only in digging ourselves deeper and deeper into the same hole, and we see only a steadily decreasing part of the horizon. This is the geometry of discovery.

There is more between DNA and the organism than James D. Watson ever dreamed of when he told the world that he and Francis Crick had solved the riddle of life. It was in 1953 that they realized that DNA had to be built up as a double helix. But there is more - not because life is a mystical, vital essence, as previously believed - but because the research within molecular biology that followed Watson and Crick's discovery, has revealed a highly organized system of processes, from the individual cell's microcosm to the total ecosystem. This points toward new questions and new limits for what we can conclude from the analysis of individual chemical compounds. We can understand how DNA specifies the sequence of building blocks that make up the proteins, yet the DNA molecule cannot tell us how the proteins, metabolism, cells, and 5organ systems in an animal function as a coherent whole.

IN THE BEGINNING WERE SOME BITS

How does a cell emerge? What is it that controls metabolism? How is a hand created? What are the mechanisms of evolution? In order to answer these kinds of questions, it is no longer sufficient to continue with the methodological reduction of complex systems to their individual components, which has been the normal procedure for biology. It is necessary to find methods to describe the *total* system, the cell organization of the organism being analyzed. The structure of the body and of the cell, its dynamics in time and space, and the sudden appearance of new traits at higher levels than those of the component parts (often called "emergence") are quite important themes in the new biology.

It is here that the computer's abilities to manage data and perform calculations on many bits of data also come into play. Artificial life is interesting, not only as a proof that life is easy to create, but also as a broad analogy, a class of models of complex calculated systems that share ecological and evolutionary conditions with many of the real organisms found in nature.

If life is characterized by processing the special form of biological information found in the genetic material, then life might be viewed as having begun as a kind of natural information-processing system. Life started in the primordial soup, from which originated the large molecules that could store - "remember" - information. This is one of the general theories of life's emergence.[6] In the beginning there was information.

In this sense, it seems plausible that our information-processing machines ought to be able to imitate Nature's trick. But what is information anyway? Is information simply a sequence of bits that allow themselves to be counted and manipulated by a machine or by an organism?

One can define information in several ways: one is as anything that can function as an answer to a question: In the beginning was ... the answer? Can the answer indeed lie at the beginning? If life is information, and information is an answer, what was the question? Does information not assume a cell, a unit in space, that can interpret and utilize this information in its own interest - a cell that (anthropomorphically speaking) can ask the question, "How do I organize my survival?" and that seeks to interpret the signs it encounters in its environment? The artificial-life approach asks these questions in a radically new way.

A-LIFE AT Los ALAMOS

In September 1987, 160 computer scientists, physicists, philosophers, biologists, anthropologists, and several other kinds of academics gathered for the first international conference on artificial life, or a-life as it came to be called. The conference took place at the Los Alamos National Laboratory, the sprawling, largely military research center in New Mexico where the atom bomb was built. Today, the mountains of New Mexico are also home to more peaceful pursuits. It is here that the data base for the human genome is being constructed (the complete micromap of human DNA). The a-life conference was sponsored by three organizations: Los Alamos's own "Center for Non-Linear Studies"; the small, but important Santa Fe Institute, long one of the leaders in research on complex systems; and finally, Apple Computers, which in fact delivered the hardware for a large number of artificial life-forms.

During the five intense days of the conference, a myriad of interesting systems was presented. Models for the emergence of life and its evolution were presented as well, including models for selection and adaptation of cooperating "organisms" in simulated ecosystems; self-reproducing automata; flocks of birds and schools of fish; computer plants that grew large and beautiful with the help of simple algorithms (the recipes from which programs are written, in this case formulas for tree-branching and length of growth); and much more. Everything that the heart could desire was here, from the more serious attempts to illustrate and test existing theories of growth and development to pure fooling around. The latter had more in common with arcade computer games and screen patterns in science-fiction films than any serious natural-science research. (The conference's nonexclusionary attitude toward what might be considered the field of artificial life has led some traditional and empirically oriented scientists to deny a-life any kind of scientific relevance or legitimacy. In this way they also avoid the difficulty of having to concern themselves with what it is really all about.)

The conference was organized and led by Chris Langton, who for years has worked with a type of mathematical structure called cellular automata, which are ideal to implement and simulate on computers. (Conway's Life game is in fact an example of such a cellular automaton.) Langton has said that he had long felt frustrated about the fragmentary nature of the field of biological modeling and simulation. For years he plowed his way through the literature in various libraries, data bases, and bookstores in order to obtain an overview of the field, although the field did not exist at all as an independent discipline. Theoretical biology had always been relatively ignored compared with theoretical physics. Biologists are experimentalists; only physicists can with good conscience and full sympathy from their colleagues allow themselves the luxury of being theoreticians. There are, of course, good reasons why this is so - witness biology's slower and subsequent development as a discipline and as a science - but people like Langton had difficulties in finding an environment for their interests.

Ultimately, however, Langton's initiative bore fruit. Not only did the Los Alamos a-life meeting testify to an enormous excitement at the new possibilities of understanding biological-development processes via modeling, of which many of the participants proudly exhibited examples. There also appeared something new: people who had been working in isolation, each tinkering with their own model of, say, cell growth, now met others who had worked with other organisms or at an entirely different level of organization, but had used much the same methods. Langton relates that there gradually emerged a collective feeling about the "essence" of artificial life. It was a vague feeling that did not, in the first round at least, lead to any programmatic declarations or any explicit research program, but that Langton later came to formulate in his book *Artificial Life.'* Langton's vision of a new research program, which is likely to become an important part of theoretical biology, contains several elements.

THE SEVEN COMMANDMENTS OF ARTIFICIAL LIFE

Seven central points, when taken together, comprise the vision of artificial life in its strongest, most ambitious form.[8] These apply both to the creation of a new research area and to the technical-scientific realization of new types of systems. The significance of these points will be reiterated in the chapters that follow. In its ambitious version, the concept of artificial life encompasses the following ideas.

1. The biology of the possible. Artificial life does not concern the special wet and carbon-based life as we know it here on earth, which is the subject of experimental biology. Artificial life deals with *life as it could be.* Since biology is only based on one example, life on earth, it is too empirically limited to help create truly general theories. Here artificial life is a clear and quite decisive supplement. It is not certain that we

appear as we do simply because of previously existing earthly materials and the accidental evolutionary sequence. Evolution could rest on much more general organizational laws, but these are laws that we simply do not know yet. Biology today is only the biology of actual life. It must become a biology of any possible life-forms.

2. Synthetic method. Where traditional biological research has placed emphasis on *analyzing* living beings and explaining them in terms of their smallest parts, the artificial-life perspective attempts to *synthesize* life-resembling processes or behavior in computers or other media.

3. Real (artificial) life. Artificial life is the study of humanly created systems that exhibit behavior characteristic of natural, living systems. However, what is it in artificial life that is artificial, in the sense of false, unnatural, or humanly created? That which is "artificial" about life *in silico - all* gadgets and information structures in the form of machines, models, and constructed "organisms" - is not the behavior as such. The behavior, the generalized process, is just as genuine as the behavior exhibited by real-life organisms. No, the "artificial" of artificial life rests solely in the *components* (like the silicon chips, formulas, computational rules, and the like) of which it consists. These are designed by us. The behavior, however, is produced by the artificial life itself.

4. All life is form. Neither actual nor possible life is determined by the matter of which it is constructed. Life is a process, and it is the *form* of this process, not the *matter,* that is the essence of life. One can therefore ignore the material and instead abstract from it the *logic* that governs the process, taking it out of the concrete material form of the life we know. Hence, one can thus achieve the same logic in another material "clothing" or substratum. Life is fundamentally independent of the medium.

These four theses are related. It is precisely because life is a form that the biology of the possible can be studied by a synthetic method in which the form or formal mathematical descriptions (via computer calculations, for example) can be made to achieve a behavioral sequence that is just as authentic as earthly life itself. From here also follow three additional commandments about the way in which artificial life must be constructed:

5. Bottom-up construction. The synthesis of artificial life takes place best via a principle of computer-based information processing called "bottom-up programming": at the bottom many small units and a few simple rules for their internal, purely local interaction are defined. (This is the real programming.) From this interaction arises the coherent "global" behavior at the general level, behavior not previously programmed according to specific rules. Bottom-up programming corresponds to the fact that our proteins are "programmed" relatively explicitly by DNA, but there is no gene that directly specifies the form of the face or the number of fingers. This kind of programming contrasts with the dominant programming principle within artificial intelligence (AI). Here one attempts to construct intelligent machines by means of programs made from the top down: the total behavior is programmed a priori by dividing it into strictly defined subsequences of behavior, which are in turn divided into precise subroutines, smaller subsubroutines, etc., all the way down to the program's own machine code. The bottom-up method in artificial life imitates or simulates processes in nature that organize themselves. We might also call these processes "simulated self-organization."

6. Parallel processing. While information processing in a classical computer takes place sequentially - similar "one-logical-step-at-a-time" thinking is also found in classical AI - the principle for information processing in artificial life is based on a massive parallelism that occurs in real life. In real life the brain's nerve cells work alongside each other without waiting for their neighbor to "finish his work"; in a flock of birds it is the simultaneity of the many birds' individual small changes in the direction of flight that gives the flock its dynamic character. Artificial neural networks are a typical example of parallel information processing and, hence, a kind of artificial life. (That the parallelism of neural network models can be simulated on sequential computers is simply a stroke of luck and says nothing about the computational principle itself.)

7. Allowance for emergence. The essential feature of artificial life is that it is not predesigned in the same trivial sense as one designs a car or a robot. The most interesting examples of artificial life exhibit "emergent behavior." The word "emergence" is used to designate the fascinating whole that is created when many semisimple units interact with each other in a complex (nonlinear) fashion.' In computational terms, it is the bottom-up method that allows for the emergence of new, unforeseen phenomena on the superordinate level, a phenomenon that is crucial for living systems.

An ecological balance in a small lake with plants, plankton, invertebrate animals, and fish is a fine example of dynamic emergent behavior that cannot be explained without the entire process of which it consists. The balance may well contain chaotic aspects (the system will perhaps never completely repeat its own movements). The important aspect, however, is that this dynamism forms an integrated behavioral whole for the entire system; that is, it has a property that does not characterize its individual component ele-ments.10 It is similar to partly genetically determined embryo development, where one cannot predict the entire system's phenotype (the individual's actual appearance) from the genotype (the specific set of genes), even if the genotype could be completely known. (Often it is the reverse: one deduces the genotype of a particular trait from a knowledge of its phenotype.) We can generalize as follows: in artificial life, the system's P-type cannot be predicted from its G-type, in any case not for any random program.11 The G-type comprises the simple rules under which the system operates; for instance, the two rules in Conway's "Life" game. The P-type is the model's overall emergent behavior, such as the glider's diagonal wriggling down through the Life grid.

Artificial life is emergent life, but it is not the mystical emergence that the old vitalists dreamed of: artificial life is in a way a mechanist holism.

With artificial life we have a new, creative way of dealing with the contradiction in biology between reductionism and holism. J. Doyne Farmer, who has worked with a model for the emergence of life, believes that artificial life makes it possible to be a mechanist and a vitalist at the same time.[12] Creativity within the field is also expressed linguistically: new expressions are arising almost daily. The researchers who regard themselves as belonging to the community of artificial-life researchers have begun to call themselves "a-lifers." Not every alifer would wholeheartedly adhere to all seven of the basic ideas listed above. Among alifers there is remarkable interest in the philosophical and ethical questions that the entire initiative raises. This was seen especially at the second international artificial-life meeting, held in Santa Fe, New Mexico, in February 1990, where a central topic of discussion was how life could be defined at all.

Do we really know what life is? We know it intuitively, of course, but this is the same as asking, "What is thinking?" There are psychologists and philosophers who do not regard research in artificial intelligence as real psychology; that is, as at all relevant for understanding the specifically human mode of reasoning. People often think in analogies and images, for example. The AI-based expert systems do not. Artificial intelligence may be seen, rather, as a kind of nonempirical experimental logic, mathematics, or semiotics.[13] In the same way, research in artificial life can be seen as an experimental (and biologically inspired) philosophy that investigates the conditions under which anything can be considered to be alive, to be biological life, b-life. Perhaps it is here that a-life research will have its greatest impact.

Biologists find it difficult to describe the mechanisms of organisms molecule by molecule. It is not much easier for the alifer, who wants to simulate the logic of self-reproduction or other aspects of the game of life. Yet life itself, before being inserted into the scientific framework and systems, seems to contain an imponderable, inexplicable, unutterable, unbelievable, and unapproachable lightness.