

ARTIFICIAL LIFE:

From Life-As-We-Know-It to Life-As-It-Could-Be

MOSHE SIPPER

*Department of Computer Science,
Ben-Gurion University, POB 653,
Beer-Sheva 84105, Israel*

1. Artificial Life

“He wanted to dream a man: he wanted to dream him with minute integrity and insert him into reality.” This was the goal of the silent man who came from the South, in Jorge Luis Borges’s short story *The Circular Ruins*. From Pygmalion, Frankenstein, and the Golem to Star Trek’s Lieutenant Commander Data, the dream of administering the breath of life has fascinated mankind since antiquity (Sipper, 2002). This question moved from the realm of science fiction to that of science with the advent of the field known as *artificial life (ALife)*. The term was coined by Christopher G. Langton, organizer of the first artificial life conference, which took place in Los Alamos in 1987 (Langton, 1989).

“Artificial Life,” wrote Langton (in the proceedings of the second conference), “is a field of study devoted to understanding life by attempting to abstract the fundamental dynamical principles underlying biological phenomena, and recreating these dynamics in other physical media—such as computers—making them accessible to new kinds of experimental manipulation and testing.” (Langton et al., 1992) While biological research is essentially *analytic*, trying to break down complex phenomena into their basic components, artificial life is *synthetic*, attempting to construct phenomena from their elemental units, as such adding powerful new tools to the scientific toolkit. This is, however, only part of the field’s mission. As put forward by Langton, “In addition to providing new ways to study the biological phenomena associated with life here on Earth, *life-as-we-know-it*, Artificial Life allows us to extend our studies to the larger domain of the ‘bio-logic’ of possible life, *life-as-it-could-be*, whatever it might be made of and wherever it might be found in the universe.”

2. Key Challenges

One way to understand the workings and goals of a scientific endeavor is to consider the key challenges its practitioners are facing. Beadau et al. (2000) have recently published such a structured list of fourteen key open problems in artificial life. The

challenges are classified under three broad categories: the transition to life, the evolutionary potential of life, and the relation between life and mind and culture.

- A. How does life arise from the nonliving?
1. Generate a molecular proto-organism *in vitro*.
 2. Achieve the transition to life in an artificial chemistry *in silico*.
 3. Determine whether fundamentally novel living organizations can exist.
 4. Simulate a unicellular organism over its entire lifecycle.
 5. Explain how rules and symbols are generated from physical dynamics in living systems.
- B. What are the potentials and limits of living systems?
6. Determine what is inevitable in the open-ended evolution of life.
 7. Determine minimal conditions for evolutionary transitions from specific to generic response systems.
 8. Create a formal framework for synthesizing dynamical hierarchies at all scales.
 9. Determine the predictability of evolutionary consequences of manipulating organisms and ecosystems.
 10. Develop a theory of information processing, information flow, and information generation for evolving systems.
- C. How is life related to mind, machines, and culture?
11. Demonstrate the emergence of intelligence and mind in an artificial living system.
 12. Evaluate the influence of machines on the next major evolutionary transition of life.
 13. Provide a quantitative model of the interplay between cultural and biological evolution.
 14. Establish ethical principles for artificial life.

As can be seen, the field of artificial life encompasses a wide range of problems, dealing with issues pertaining to the very basics of life, and all the way up to intelligence and societies, along the way passing through basic questions of organization, emergence, and structuring.

Since I cannot elaborate upon all of the above issues within the limited scope of this chapter, I shall try to focus on a small number of key issues, which underlie many of the above challenges. ALife pioneer Thomas Ray wrote, "I would consider a system to be living if it is self-replicating, and capable of open-ended evolution." (Ray, 1992) Indeed, self-replication and evolution are both crucial to ALife research.

3. Self Replication

The study of self-replicating structures in man-made systems began in the late 1940s, when John von Neumann—one of the twentieth century's most eminent mathematicians and physicists—posed the question of whether a machine can self-replicate (i.e., produce copies of itself). He wrote:

Living organisms are very complicated aggregations of elementary parts, and by any reasonable theory of probability or thermodynamics highly improbable. That they should occur in the world at all is a miracle of the first magnitude; the only thing which removes, or mitigates, this miracle is that they reproduce themselves. Therefore, if by any peculiar accident there should ever be one of them, from there on the rules of probability do not apply, and there will be many of them, at least if the milieu is reasonable. (von Neumann, 1966).

Von Neumann was not interested in building an actual machine, but rather in studying the theoretical feasibility of self-replication from a mathematical point of view. He succeeded in proving (mathematically) that machines can self-replicate, laying down along the way a number of fundamental principles involved in this process.

To conduct a formal mathematical investigation of the issue, von Neumann used a model conceived by his colleague, mathematician Stanislaw Ulam. The model, known as *cellular automata (CA)*, consists of a large grid of cells (similar to a checkerboard), each possessing a certain state at a given moment. The number of possible states per cell is finite and is usually small (in figures, each state is usually represented by a different color). All cells change state simultaneously such that the state of a cell at the next time step depends only on its state at the current time step and the states of its neighboring cells. The principle that guides state transformations is applied identically to all cells and is referred to as the *transition rule*. Note that this rule is entirely *local*, thus any global phenomena occurring throughout the grid are due to interactions of numerous cells.

A *machine* in the CA model is a collection of cells that can be regarded as operating in unison. Thus, one can observe simple “creatures” that are able to move within this austere universe, as demonstrated in Figure 1 for a well-known CA rule—the “game of life.”

One of von Neumann’s main conclusions was that the replicative process uses the assembly instructions in two distinct manners: as interpreted code (during actual assembly), and as uninterpreted data (copying of assembly instructions to offspring). During the following decade, when the basic genetic mechanisms began to unfold, it became clear that nature had “adopted” von Neumann’s conclusions. The process, by which assembly instructions (DNA) are used to create a working machine (proteins), makes dual use of information: as interpreted code and as uninterpreted data, the former referred to in biology as *translation*, the latter as *transcription* (Figure 2).

A major problem with von Neumann’s approach (as well as a number of his successors) is the complexity of the constructor, which requires hundreds of thousands of cells. In addition, each cell can be in one of 29 states rather than just two. In 1984 Christopher Langton observed that although the capacity for universal construction is a *sufficient* condition for self-replication, it is not a *necessary* one; furthermore, natural systems are not capable of universal construction. He devised a small CA “loop,” capable solely of self-replicating, but which was much simpler than von Neumann’s design, with but a few dozen cells (Langton, 1984; for a demonstration of its workings see <http://necsi.org/postdocs/sayama/sdsr/java/>).




Figure 1. The “game of life” CA was defined by John H. Conway in 1968: “Life occurs on a virtual checkerboard. The squares are called cells. They are in one of two states: alive or dead. Each cell has eight possible neighbors, the cells which touch its sides or its corners. If a cell on the checkerboard is alive, it will survive in the next time step if there are either two or three neighbors also alive. It will die of overcrowding if there are more than three live neighbors, and it will die of exposure if there are fewer than two. If a cell on the checkerboard is dead, it will remain dead in the next time step unless exactly three of its eight neighbors are alive. In that case, the cell will be ‘born’ in the next time step” (Berlekamp et al., 1982). Shown above at time step 0 are two patterns, a stationary square, known as a “block,” and a moving creature, known as a “glider.” The latter displaces itself one square diagonally every four time steps. Upon meeting other creatures, such as the block, both are subject to this simple universe’s basic rule, or “laws of physics,” which cause their mutual annihilation in this case. Note that actions such as “movement” are purely in the eyes of the beholder—the most basic level consists of simple state transformations, with no movement at all. However, by considering higher-level views of the system, one can introduce new terms to describe the perceived phenomena. Such phenomena have often been termed “emergent,” although this appellation is highly problematic (Ronald et al., 1999).




Figure 2. A schematic diagram of von Neumann’s self-replicating CA. The machine is a universal constructor (UC) capable of constructing, through the use of a “constructing arm,” any configuration whose description (genome) can be stored on its input tape. Thus, given its own description, the machine is capable of constructing a copy of itself, i.e., self-replicate. The parent UC first *interprets* the genetic information stored on the tape as building instructions to create an offspring UC, and then copies the genomic tape *uninterpreted* and attaches it to the offspring, so the latter can replicate in its turn.

The study of self-replication has been taking place now for more than half a century (Sipper, 1998; also <http://www.cs.bgu.ac.il/~sipper/selfrep>). This research might better our understanding of self-replication in nature, as well as find many technological applications. There is much talk today of nanotechnology, where self-replication is of vital import: we would like to be able to build one miniature machine, which would then sally forth and multiply.

As noted above self-replication is but one of two major characteristic of life, the second being evolution.

4. Evolutionary Algorithms

The only process currently known to have produced an ecosystem of living creatures, and in particular, of intelligent beings, is that of natural evolution. Darwin laid out the core of the currently accepted theory of evolution, its major elements being (Ray, 1994):

- Individuals vary in their viability in the environments that they occupy.
- This variation is heritable.
- Self-replicating individuals tend to produce more offspring than can survive on the limited resources available in the environment.
- In the ensuing struggle for survival, the individuals best adapted to the environment are the ones that will survive to reproduce.

The continual workings of this process over the millennia causes populations of organisms to change, generally becoming better adapted to their environment.

Having witnessed the study of self-replication in an artificial setting, we now do the same for evolution. The idea of applying the biological principle of natural evolution to artificial systems was introduced in the 1950s and the 1960s, when several researchers studied evolutionary systems with the idea that evolution could be used as an optimization tool for engineering problems. Central to all the different methodologies is the notion of solving problems by evolving an initially random population of candidate solutions, through the application of operators inspired by natural genetics and natural selection, such that in time “fitter” (i.e., better) solutions emerge (Sipper, 2002). Nowadays, these so-called evolutionary algorithms are ubiquitous, having been successfully applied to numerous problems from different domains, including optimization, automatic programming, machine learning, economics, operations research, immune systems, ecology, and population genetics. In particular, they are central to ALife research.

Based on (and inspired by) the workings of evolution by natural selection, the basic meta-algorithm is seductively (and, I might add, deceptively) simple, and can be expressed in a mere 8 lines of pseudocode:

1. produce an initial **population**. of individuals, these latter being candidate solutions to the problem at hand
2. **evaluate** the fitness of each individual in accordance with the problem whose solution is sought

3. *while* termination condition not met *do*
4. **select** fitter individuals for reproduction
5. **recombine (crossover)** individuals
6. **mutate** individuals
7. **evaluate** fitness of modified individuals
8. *end while*

Over the past two decades evolutionary algorithms have proven their worth beyond a doubt, at times not only solving hard problems but indeed competing with their human designers (Koza et al., 2003).

A major reason for the success of this seemingly blind search is its not being blind at all (as in nature): though randomness and probability do play a part, selection is a function of fitness, and the recombination of good sub-pieces—or building-blocks—from two (or more) individuals enables the creation of offspring that outperform their parents (Figure 3).

ALife researchers often use some form of evolutionary algorithm to study the issues delineated in Section 2. For example, where the transition from the nonliving to the living is concerned, one can explore the evolution of self-replicating entities, as done, e.g., by Lohn and Reggia (1997). They used an evolutionary algorithm to automatically discover transition rules that govern emergent self-replicating processes in cellular automata. Given dynamically evolving automata, one of the most difficult tasks is that of identifying effective fitness functions for self-replicating structures. Simply ascribing a higher fitness to better replicators is of no use, since at the outset (generation 0) there are no replicators at all, i.e., all individual structures have a fitness of zero. And with no variability evolution grinds to a halt. Lohn and Reggia were able to ingeniously solve the fitness problem by assigning values to “partial” replicators, thereby demonstrating that self-replicating structures can be evolved.

5. Getting “Wetter”: Lipidia

The presentation up until now may have seemed somewhat theoretical, and indeed much research in artificial life is, ipso facto, of an abstract nature. But in recent years ALife researchers have been attempting to get “closer” to biology, by building systems and models that attempt to answer biological questions through ALife methods. As an example I will briefly describe below the work my colleagues and I have undertaken in the area of Origin of Life: Lipidia.

Lipidia is a novel simulation system related to the “Lipid World” scenario for the origin of life. Lipidia allows conducting experiments with a population of assemblies containing lipid-like molecules on a two-dimensional grid (Naveh et al., 2004). The dynamics of the assemblies is modeled using the Graded Autocatalysis Replication Domain (GARD) model (Segre et al., 1998).

The “RNA World” is possibly today’s most popular theory for the origins of life. Because RNA molecules can act as catalysts in addition to acting as templates, it is hypothesized they might have been able to do both: to store alphabet-based genetic information *and* to catalyze their own creation. Life, according to this theory, began

when certain RNA molecules achieved the capability to replicate themselves. This scenario, despite its elegance, suffers from difficulties.

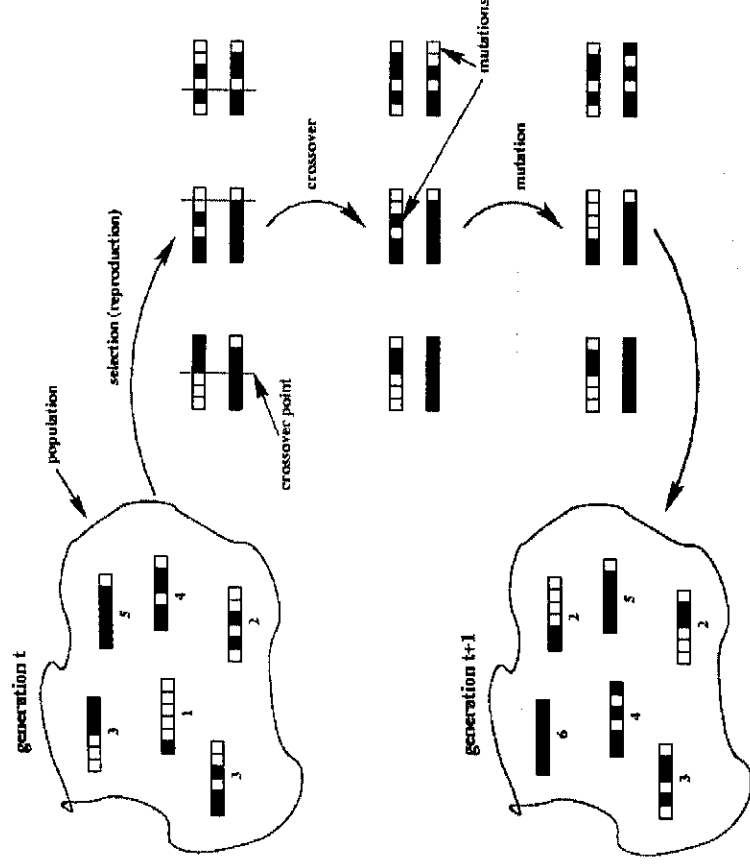


Figure 3. One generation (evaluation-selection-crossover-mutation cycle) of a simple evolutionary algorithm. Individuals are represented as bit strings (the so-called *genomes*). Fitness in this toy example is the number of 1 (black) bits. Selection is performed proportionately to fitness, so that high-fitness individuals are more likely to be selected. Recombination—or crossover—is performed on two individuals by selecting a crossover point at random and exchanging the chunks beyond this point. Mutation is performed by flipping a small number of bits (with low probability). As can be seen, this simple evolutionary scenario has produced a perfect solution by the next generation, in this case due to the crossover operation, which has glued a good sub-piece from one individual (3 1s on the right) with another good sub-piece from the second individual (3 1s on the left).

In an attempt to come up with a probable scenario, having observed that no known bio-molecule is capable of self replication in its naked form, it has been suggested that self replication might not have been achieved by a single molecule, but rather by a molecular ensemble (Kauffman, 1995). This work is based on “The Lipid World” scenario (Segre et al., 2001), which follows that line of thought. The scenario assumes that self-replication was initially achieved by non-covalent assemblies of lipid-like molecules that contained mutually catalytic sets. RNA according to this scenario, while possibly playing an important role, came later.

We used the GARD model to quantitatively model and simulate the developmental process of non-covalent assemblies of lipid-like molecules. Previous studies using the GARD model have mostly examined such assemblies in a one-at-a-time fashion. The behavior of assembly populations has been largely unexplored. In this work we expanded the model to a *population* of assemblies and obtained quantitative and qualitative results regarding its behavior. Also, previous studies assumed idealization of an infinite environment where the assembly's effect on the environment is negligible and "food" molecules are in infinite supply. In our work we introduced a finite environment to the model, which allows cross-interactions between assemblies via the environment. We also compared the effect of finite environments vs. infinite environments (Figure 4).




Figure 4. A LipidA screen shot, typifying an A Life scenario. LipidA is based on a two-dimensional interaction *grid*, as with cellular automata. Each square on the *grid* is called a *grid location*, or *location* for short. For each location there is a defined *environment* containing a variety of *molecules*. Each location may contain zero or more *assemblies* of molecules. The location's environment is common for all assemblies contained within it. Molecules from the environment may *join* an assembly, and molecules from the assemblies may *leave* their assembly back to the environment. "Matter" on the grid is therefore preserved—no matter is ever lost or created. The finite environment model can be replaced by an infinite environment model of fixed concentrations. For full details see (Navet et al., 2004).

Our results showed that a finite environment produces more attractors (species), and faster, than an infinite environment. A finite environment allows more assemblies to occur in more attractors and in greater numbers (Naveh et al., 2004). Thus, diversity increases.

The results might be considered surprising. One might think that having an infinite supply of resources, in the form of “food” molecules, might help to “do more.” According to our results, it only helps to “do more of the same.” Diversity seems to spring when resources are limited. It is when resources for the “best solutions” run out that the race towards alternative solutions begins.

While the details are beyond the scope of this chapter, this work exemplifies an ALife research: the formation of digitally (but not carbonaceously) testable questions in a biological domain (in this case, Origin of Life), followed by the conception of a synthetic model, its implementation and the running of simulations, and the drawing of conclusions.

6. Strong Artificial Life?

The research described heretofore is often referred to as *weak* artificial life, namely, the application of synthetic—often digital—tools to the study of biological phenomena. But can we create bona fide life? This latter is the goal of *strong* artificial life, which for now remains mostly in the realm of theory and philosophy.

One way to create life, currently being explored by a number of groups, is to start at the very bottom—with bacteria. These have small genomes—on the order of a few hundred nucleotides—and we may be able to fabricate such small living organisms. A related research subjects simple bacteria to knock-out experiments, wherein genes are systematically taken out, in order to arrive at the very minimal organism which is still viable.

At the other end, we might employ *open-ended* evolution. As opposed to evolutionary algorithms described above, where a fitness function is imposed by the human designer, open-ended evolution—such as occurs in nature—has no such external imposition, but rather seeks out a trajectory on its own. Creating a set-up for such an open-ended process is a major challenge, probably still far in the future (a lengthy discussion of this challenge can be found in (Sipper, 2002)). Yet one cannot help but wonder, what if: What if we came up with the prodigious resources (computer power, “wet” power) necessary to mount an open-ended evolutionary process? What if we were able to *design* such an experiment—itself a fundamental problem? Would we then see—after a day, a month, or a decade—the emergence of living creatures? If so, would they really be artificial? After all, artificial life might in fact be an oxymoron—how can life be artificial? If something is truly alive—assuming we can somehow agree upon this fact, as no ironclad definition exists—then what is artificial about it?

Such speculative, futuristic science aside, artificial life today is providing a novel outlook on many scientific issues, and enabling the execution of experiments too costly or downright impossible in a “real” laboratory.

7. References

- Bedau, M. A., McCaskill, M. A., Packard, N. H., Rasmussen, S., Adami, C., Green, D. G., Ikegami, T., Kaneko, K. and Ray, T. S. (2000) Open Problems in Artificial Life, *Artificial Life*, **6**: 363–376.
- Berlekamp, E. R., Conway, J. H. and Guy, R. K. (1982) *Winning Ways for your Mathematical Plays*, Academic Press, New York.
- Kauffman, S. A. (1995) *At Home in the Universe: The Search for the Laws of Self-Organization and Complexity*, Oxford University Press, UK.
- Koza, J. R., Keane, M. A., Streeter, M. J., Mydlowec, W., Yu, J. and Lanza, G. (2003) *Genetic Programming IV: Routine Human-Competitive Machine Intelligence*, Kluwer Academic Publishers Boston.
- Langton, C. G. (1984) Self-reproduction in cellular automata, *Physica D*, **10**: 135–144.
- Langton, C. G. (1989) *Artificial Life: Proceedings of an Interdisciplinary Workshop on the Synthesis and Simulation of Living Systems*, Addison-Wesley, Redwood City, CA.
- Langton, C. G., Taylor C., Farmer J. D. and Rasmussen S. (1992) *Artificial Life II: Proceedings of the Workshop on Artificial Life*, Addison-Wesley, Redwood City, CA.
- Lohn, J. D. and Reggia, J. A. (1997) Automatic discovery of self-replicating structures in cellular automata. *IEEE Transactions on Evolutionary Computation*, **1**: 165–178.
- Naveh, B., Sipper, M., Lancet, D. and Shenhav, B. (2004) Lipidia: An Artificial Chemistry of Self-Replicating Assemblies of Lipid-like Molecules, In: J. Pollack, M. Bedau, P. Husbands, T. Ikegami, and R. A. Watson (eds.) *Artificial Life 9: Proceedings of the Ninth International Conference on Artificial Life*, pp. 466–471, The MIT Press, Cambridge, Massachusetts.
- Ray, T. S. (1992) An approach to the synthesis of life, In: C. G. Langton, C. Taylor, J. D. Farmer, and S. Rasmussen (eds.) *Artificial Life II*, Addison-Wesley, Redwood City, CA, pp. 371–408.
- Ray, T. S. (1994) Evolution, complexity, entropy and artificial reality, *Physica D*, **75**: 239–263.
- Ronald, E. M. A., Sipper, M. and Capearrete, M. S. (1999) Design, observation, surprise! A test of emergence, *Artificial Life*, **5**: 225–239.
- Segre, D., Lancet, D., Kedem, O. and Pilpel, Y. (1998) Graded autocatalysis replication domain (GARD): Kinetic analysis of self-replication in mutually catalytic sets, *Origins of Life and Evolution of the Biosphere*, **28**: 501–14.
- Segre, D., Ben-Eli, D., Deamer, D. W. and Lancet, D. (2001) The lipid world, *Origins of Life and Evolution of the Biosphere*, **31**: 119–145.
- Sipper, M. (1998) Fifty years of research on self-replication: An overview, *Artificial Life*, **4**: 237–257.
- Sipper, M. (2002) *Machine Nature: The Coming Age of Bio-Inspired Computing*, McGraw-Hill, New York.
- von Neumann, J. (1966) *Theory of Self-Reproducing Automata*, University of Illinois Press, Illinois, Edited and completed by A. W. Burks.