

The scientific and philosophical scope of artificial life

Mark Bedau

Department of Philosophy, Reed College, Portland OR 97202, USA

mab@reed.edu, <http://www.reed.edu/~mab>

Abstract

The new interdisciplinary science of artificial life has had a connection with the arts from its inception. This paper provides an overview of artificial life, reviews its key scientific challenges, and discusses its philosophical implications. It ends with a few words about the implications of artificial life for the arts.

Artificial life is a young interdisciplinary collection of research activities aimed at understanding the fundamental behavior of life-like systems by synthesizing that behavior in artificial systems. As befits a journal for artists who use science and developing technologies, papers discussing artificial life regularly appear in the pages of *Leonardo*. There is also traffic in the other direction; for example, the biennial International Conference on Artificial Life is the primary vehicle for publishing all the latest scientific developments in artificial life, but more than five percent of the articles published in the proceedings of the last conference [1] concerned the application of artificial life to art and music [2-5]. People in both communities believe that the arts and artificial life have much to offer each other. Given this, it would be useful for the two communities to know each other better. The opportunity to counteract the hype and misleading publicity surrounding artificial life is also welcome. The truth is often more interesting and surprising than fiction, and it is always more valuable.

This paper aims primarily to provide an overview of artificial life, explaining its approach to science and technology and outlining its main open problems, and sketching its broader philosophical implications. It ends with a few words about the implications of artificial life for the arts.

Overview of artificial life

Life is an interconnected web of adaptive systems produced spontaneously by the process of evolution. Living systems exhibit impressively robust and flexible functionality at many levels of analysis. Examples range from the genomic and proteomic regulatory systems that control how biological organisms develop and function, to the evolving ecological networks through which members of different species interact. Man-made adaptive systems like the myriad communication networks that span the globe are beginning to approach the complexity of adaptive systems found in nature. Learning how to engineer flexible and robust adaptive complexity is one of the biggest challenges facing society in the twenty-first century.

Traditionally, adaptive systems of different kinds were modeled independently in different disciplines. Artificial life is now bringing together biologists, physicists, chemists, psychologists, economists, and anthropologists with computer scientists and philosophers to create a unified understanding of adaptive systems of all types. Artificial life (also known as "ALife") studies life

and life-like processes by synthesizing them in artificial media, most often using computer technology. The goals of this activity include modeling and even creating life and life-like systems; the goals also include developing practical applications involving new technologies that exploit intuitions and methods taken from living systems. The phrase “artificial life” was coined by Christopher Langton. He envisioned a study of life as it could be in any possible setting, and he organized the first conference that explicitly recognized this field [6]. There has since been a regular series of conferences on artificial life and a number of academic journals have been launched to publish work in this new field.

Artificial life borrows from other, older disciplines, especially computer science, cybernetics, biology, and the study of complex systems in physics. Its closest intellectual cousin is artificial intelligence (AI). But there is a crucial difference between the modeling strategies AI and ALife typically employ. Most traditional AI models are top-down-specified serial systems involving a complicated, centralized controller that makes decisions based on access to all aspects of global state. The controller’s decisions have the potential to affect directly any aspect of the whole system. On the other hand, many natural living systems exhibiting complex autonomous behavior are parallel, distributed networks of relatively simple low-level “agents” that simultaneously interact with each other. Each agent’s decisions are based on information about only its own local situation, and its decisions directly affect only its own local situation. ALife’s models follow nature’s example. The models themselves are bottom-up-specified parallel systems of simple agents interacting locally. The local interactions are repeatedly iterated and the resulting global behavior is observed. The whole system’s behavior is represented only indirectly. It arises out of the interactions of a collection of directly represented parts.

The synthetic methodology of artificial life has several virtues. The discipline of expressing a theory synthetically, especially in computer code, forces precision and clarity. It also insures that hypothesized mechanisms are feasible. Computer models also facilitate the level of abstraction required for maximally general models of phenomena. The bottom-up architecture of artificial life models creates an additional virtue. Allowing micro-level entities continually to affect the context of their own behavior introduces a realistic complexity that is missing from analytically studied mathematical models. Analytically solvable mathematical models can reveal little about the global effects that emerge from a web of simultaneous nonlinear interactions. The obvious way to study the effects of these interactions is to build bottom-up models and then empirically investigate their emergent global behavior through computer simulations.

Many artificial life models are designed not to represent known biological systems but to generate wholly new and extremely simple instances of life-like phenomena. The simplest example of such a system is the famous cellular automaton called the “Game of Life”, devised by the mathematician John Conway in the 1960s [7]. Computer simulation is crucial for the study of complex adaptive systems. It plays the role that observation and experiment play in more conventional science. The complex self-organizing behavior of the Game of Life would never have been discovered without simulating thousands of generations for millions of sites. The same holds for virtually all other systems studied by artificial life.

Rather than merely producing computer simulations, some artificial life research aims to implement system in the real world. The products of this activity are physical devices such as robots that exhibit characteristic life-like behavior. Some of these implementations are motivated by the concern to engineer practical devices that have some of the useful features of living systems, such as robustness, flexibility, and autonomy. But some of this activity is primarily theoretical, motivated by the belief that the best way to confront the hard questions about how life occurs in the physical world is to study real physical systems. Examples range from evolvable hardware, which attempts to use biologically-inspired adaptive processes to shape the configuration of micro-electronic circuitry, to biologically-inspired robotics, such as using evolutionary algorithms to automate the design of robotic controllers and swarms of robots communicating locally to achieve some collective goal.

Grand challenges in artificial life

A good way to understand a scientific community is to grasp its central aims. The fact that a second generation of scientists is commencing work in artificial life prompted the organizers of the last International Conference on Artificial Life to publish list of grand challenges [8]. Since there is still so much unknown about the emergence and evolution of living systems, the list emphasizes scientific understanding rather than applications, and the challenges are unabashedly long term. This section reviews those challenges.

The challenges fall into three broad categories: the origin of life, life's evolutionary potential, and life's connection to mind and culture.

A. How does life arise from the non-living?

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 arise from inanimate matter.
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 manipulations of 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.

Challenges in the third category are more speculative, and some are interwoven with non-scientific issues. Some areas in which artificial life plays a significant role, such as robotics and art, do not appear on the list. In part this is simply a practical expedient to shorten the list as much as possible. In the rest of this section I will briefly explain a representative selection of these challenges. More information about them all can be found in the original source.

The first challenge involves no less than constructing a novel life form in the laboratory from scratch. The first targets should be the simplest possible forms of life—self-reproducing molecular structures that construct and maintain themselves in a simple environment and evolve. The environment would involve only simple forms of energy and material, and the goal would be to create an encapsulated biochemical system that can derive energy from simple chemicals or light and use information carried in primitive genes. The attempt to create a proto-organism that self-replicates and evolves, using energy and nutrients from its environment, illustrates artificial life's concern with understanding life by synthesizing it. It also shows that artificial life's interests are not just fanciful abstractions. A fundamental understanding of real life in the real world is a key part of what artificial life hopes to provide.

Few questions concerning living systems are as fundamental as the spontaneous generation of life, and the second challenge explores this issue in artificial chemistries. Artificial chemistries are computer-based model systems comprised of objects (abstractions of molecules) which are assembled by collisions among simpler objects according to predefined interaction rules. The chemistry must be constructive rather than merely descriptive, with rules that determine arbitrarily complex products from arbitrarily complex collisions. Furthermore, the chemical interaction rules should be simple compared with the ultimate products that they create. This challenge illustrates artificial life's emphasis on understanding the amazing spontaneous emergence of structure and hierarchy that characterizes life. It also shows how artificial life uses abstraction to capture the essence of such a process.

Life as we know it encodes information about hierarchically organized spatially localized individuals in genetic structures. The third challenge involves determining whether this or any other particular form of organization is necessary for life. The question is relevant to the search for extra-terrestrial life in the universe. Examples of fundamentally different organizations include those without a genetic code, without spatially localized individuals, without hierarchical organization, without a genotype-phenotype distinction or, indeed, without any symbolic representation scheme. The debate about what organizations are "fundamentally different" will clarify our understanding of the nature of life, and pursuing this challenge will expand our horizons and challenge our preconceptions about life.

The sixth challenge concerns life's contingency. Artificial life is trying to discern the features common to all evolutionary processes, or to broad classes of evolutionary processes. It aims to determine whether different kinds of

evolutionary processes have different potential creativity. Artificial life expects that many of the most fundamental features of the evolution of life on Earth are independent of the physical media that happen to embody the process. Digital information processing in computers is a very different medium from molecular biology, yet artificial life has been building digital organisms based on genetic and cellular principles from its inception. Digital media provide considerable scope to vary the “physics” underlying the evolutionary process, so it is straightforward to investigate evolutionary contingency in that context. But we do not know whether digital systems and physical systems have the same potential for evolutionary innovation. Artificial life’s commitment to a synthetic methodology shows itself here. Not content with mere verbal speculation about kinds of evolutionary creativity, artificial life insists on making systems that actually demonstrate those capacities.

The forms of life that we know all have a complex organization that enables them to act autonomously and in their own interests. Organisms can be transgenically manipulated to express different genes, but the evolutionary consequences and limits of such manipulations are unknown. This raises artificial life’s ninth challenge: determining how well we can predict the evolutionary consequences of making new forms of life. To what extent can one redesign organisms to fulfill novel functions without disrupting their viability? Is there a tradeoff between size of modification and viability? Better understanding of the genetics of development will enable us to create novel multicellular organisms, but they might not flourish or they might unleash unanticipated and uncontrollable ecological consequences. Perhaps major changes to organisms can be perfected only by lengthy coevolutionary optimization. Along with genetic engineering, artificial life must confront questions like these since it unleashes novel autonomous beings with lives of their own. Furthermore, artificial life is ideally poised to address such questions, since it can synthesize all kinds of genetic manipulation in isolated digital contexts.

Once life originated, biological evolution underwent a number of major evolutionary transitions, such as the origin of eukaryotes, the origin of multicellular life, and the origin of culture. Presumably there will be more major transitions in the future. Once culture originates, it has the capacity to evolve on its own. The past century has seen the explosion of technological culture including the creation of computing machines and complex distributed networks connecting them. Many agree it is only a matter of time before artificial life creates machines that are alive, intelligent, reproduce their own kind, have their own purposes, set their own goals, and evolve autonomously. These machines will be part of our world and their evolution will affect our future. Think of how machines currently influence the nature and rate of human communication and interconnection. All this suggests that machines might play an unprecedented role in the next major evolutionary transition. Artificial life’s twelfth challenge is to predict and explain this role. Machines will certainly play at least a supporting role in the next major evolutionary transition since they provide an infrastructure that influences the rate and direction of change. They might even be central players, if autonomously evolving machines have proliferated. This will again push the boundaries of what it means to be alive and embody new forms of the unbounded creativity of evolution.

Culture is one of the products of human existence, and culture itself evolves. Artificial life's thirteenth challenge is to understand the connection between biological and cultural evolution. Examples of cultural evolution include the development of economic markets, the changes in technological infrastructure (see the previous challenge), and growth and revolution in scientific opinion. Some treatments of cultural evolution (e.g., sociobiology and evolutionary psychology) consider how cultural traits evolve due to their impact on biological fitness. But one can also consider how cultural traits evolve in their own right, as Dawkins did when he coined the word "meme" [9]. This sort of "pure" cultural evolution is driven by mechanisms similar to those behind biological evolution, but there are important differences. In each case traits exhibit variation, heritability, and differential fitness, but cultural traits are transmitted not genetically but psychologically, and their fitness concerns not biological survival and reproduction but retention in and proliferation across minds. One question concerns the similarities and differences in the behavior of biological and cultural evolution. Do both exhibit the same kind of creative explosions, and for similar reasons? Another question concerns how they are interconnected. Confronting these questions invites us to reconceptualize life, culture, and technology. Artificial life gives us an increasingly constructive role in our future. Even if we do not try to shape our future to fit our current preconceptions of what is possible, artificial life can help us to understand and appreciate the open-ended creative process in which we are all embedded.

Philosophical implications of artificial life

Artificial life is not just a scientific and engineering enterprise. It offers a new perspective on the essential nature of many fundamental aspects of reality like life, adaptation, and creation. Thus it has rich implications for a number of broad philosophical issues. In fact, philosophy and artificial life are natural intellectual partners, for a variety of reasons. Both seek to understand phenomena at a level of generality that is sufficiently deep to ignore contingencies and reveal essential natures. In addition, by creating wholly new kinds of life-like phenomena, artificial life continually forces us to reexamine what it is to be alive, intelligent, creative, etc. Furthermore, artificial life's computational methodology is a direct and natural extension of philosophy's traditional methodology of *a priori* thought experiment. In the attempt to capture the simple essence of vital processes, artificial life models abstract away as many details of living systems as possible. These models are thought experiments that are explored with the help of a computer. Like the traditional armchair thought experiments, artificial life simulations attempt to answer "What if X?" questions, but the premises they pose are complicated enough that their implications can be explored only by computer simulation; armchair analysis is simply inconclusive. Synthesizing thought experiments on a computer brings a new kind of clarity and constructive evidence to philosophy. In this section I illustrate artificial life's broad implications for a handful of philosophical topics: emergence, evolution, life, and mind.

Emergence. One of life's amazing features is how the whole is more than the sum of the parts. This is called emergence. In general, emergent phenomena share two broad hallmarks: they both depend on and are autonomous from underlying phenomena. Although apparent emergent phenomena are all around

us, the two hallmarks of emergence seem inconsistent or philosophically illegitimate. How can something be autonomous from underlying phenomena if it depends on them? This is the traditional philosophical problem of emergence. A solution to this problem would both dissolve the appearance of illegitimate metaphysics and show how emergence plays a constructive in scientific explanations of phenomena involving life and mind.

The aggregate global behavior of complex systems studied in artificial life offers a new way to view of emergence. On this view, a system's macrostate is emergent just in case it can be derived from the system's boundary conditions and its micro-level dynamical process but only through the process of iterating and aggregating all the micro-level effects [10]. This new view explains the two hallmarks of emergence. Micro-level phenomena clearly depend on macro-level phenomena; think of how a bottom-up artificial life model works. At the same time, macro-level phenomena are autonomous because the micro-level interactions in the bottom-up models produce such complex macro-level effects that the only way to recognize or predict them is by observing macro-level behavior. This form of emergence is common in complex systems found in nature, and artificial life's models also exhibit it. This view attributes the unpredictability and unexplainability of emergent phenomena to the complex consequences of myriad, non-linear and context-dependent local micro-level interactions. Emergent phenomena can have causal powers on this view, but only by means of aggregating micro-level causal powers. There is nothing inconsistent or metaphysically illegitimate about underlying processes constituting and generating phenomena by iteration and aggregation. Furthermore, this form of emergence is prominent in scientific accounts of exactly the natural phenomena like life and mind that apparently involve emergence.

This shows how artificial life will play an active role in future philosophical debates about emergence and related notions like explanation, reduction, complexity, and hierarchy. Living systems are one of the primary sources of emergent phenomena, and artificial life's bottom-up models generate impressive macro-level phenomena wholly out of micro-level interactions. Artificial life expands our sense of what is possible, and it provides a constructive way to explore it.

Evolution. The evolution of life has shown a remarkable growth in complexity. Simple prokaryotic one-celled life lead to more complex eukaryotic single-celled life, which then lead to multicellular life, then to large-bodied vertebrate creatures with sophisticated sensory processing capacities, and ultimately to highly intelligent creatures that use language and develop sophisticated technology. This illustration of evolution's creative potential leads to a deep question about evolution's creative potential: Does evolution have an inherent tendency to create greater and greater adaptive complexity, or is the complexity of life just a contingent and accidental by-product of evolution?

Stephen Jay Gould [11] devised a clever way to address this issue: the thought experiment of replaying the tape of life. Imagine that the process of evolution were recorded on a tape. The thought experiment is to rewind the evolutionary process backward in time, erasing the tape, and then playing it forward again but allowing it to be shaped by wholly different contingencies. It is not clear what the outcome of the thought experiment is. Gould himself

suggests that “any replay of the tape would lead evolution down a pathway radically different from the road actually taken.” He concludes that the contingency of evolution destroys any possibility of a necessary growth in adaptive complexity. Daniel Dennett [12] draws exactly the opposite conclusion. He argues that complex features like sophisticated sensory processing provide such a distinct adaptive advantage that natural selection will almost inevitably discover it in one form or another. Dennett concludes that replaying life’s tape will almost inevitably produce highly intelligent creatures that use language and develop sophisticated technology.

Artificial life can make a number of contributions to this debate. Experience in artificial life has shown time and again that armchair expectations about the outcome of thought experiments like replaying life’s tape are highly fallible. The only sure way to know what to expect is to create the relevant system and observe the results of repeated simulation. In fact, artificial life is exactly where the activity of creating and studying such systems occurs. However, we cannot yet conduct the experiment of replaying life’s tape because no one has been able to create a system that exhibits continual open-ended evolution. Achieving this goal is a key open problem in artificial life, related to its sixth grand challenge. All conjectures about evolution’s inherent creativity will remain unsettled until we actually study what happens when the tape of life is replayed.

Life. Philosophers from Aristotle to Kant have addressed the nature of life. But philosophers today ignore the issue, perhaps because it seems too scientific. At the same time, most biologists also ignore the issue, perhaps because it seems too philosophical. The advent of artificial life has revitalized the question. This is partly because one can simulate or synthesize living systems only if one has some idea what life is. Artificial life’s self-conscious aim to discern the essence of life encourages liberal experimentation with novel life-like organizations and processes. Thus, artificial life both fosters a broad perspective on life. In the final analysis, the question of the nature of life will be settled by whatever perspective provides the best explanation of the rich range of natural phenomena that living systems exhibit. Better understanding of how to explain these phenomena will also help resolve a cluster of puzzles about life, such as whether life admits of degrees, how the notion of life applies at different levels in the biological hierarchy, and the relationship between the material embodiment of life and the dynamical processes in which those materials participate.

Artificial life highlights the question whether artificial constructions, especially purely digital systems existing in computers, could ever literally be alive. This question will be easier to answer once there is agreement about the nature of life; but that agreement should not be expected until we have experienced a much broader range of possibilities. So the debate over whether real but artificial life is possible continues. Some people complain that it is a simple category mistake to confuse a computer simulation of life with a real instance of it [13]. A flight simulation for an airplane, no matter how detailed and realistic, does not really fly. A simulation of a hurricane does not create real rain driven by real gale-force winds. Similarly, a computer simulation of a living system produces merely a symbolic representation of the living system. The intrinsic ontological status of this symbolic representation is nothing more than certain electronic states inside the computer (e.g., patterns of high and low

voltages). This constellation of electronic states is no more alive than is a series of English sentences describing an organism. It seems alive only when it is given an appropriate interpretation. But this charge of category mistake can be blunted. Artificial life systems are typically not simulations or models of any familiar living system but new digital worlds. Conway's Game of Life, for example, is not a simulation or model of any real biochemical system but a digital universe that exhibits spontaneous macroscopic self-organization. So, when the Game of Life is actually running in a computer, the world contains a new physical instance of self-organization. Processes like self-organization and evolution are multiply realizable and can be embodied in a wide variety of different media, including the physical media of suitably programmed computers. So, to the extent that the essential properties of living systems involve processes like self-organization and evolution, suitably programmed computers will actually be novel realizations of life.

Mind. Life forms are sensitive to the environment in various ways, and this environmental sensitivity affects their behavior in various ways. So forms of life have broadly mental capacities. Furthermore, the sophistication of these mental capacities seems to correspond to the complexity of those forms of life. So it is natural to ask whether life and mind have some deep connection. Since all forms of life must cope in one way or another with a complex, dynamic, and unpredictable world, perhaps this adaptive flexibility inseparably connects life and mind.

It is well known in the philosophy of mind and artificial intelligence that the emergent dynamical patterns among human mental states are especially difficult to describe and explain. Descriptions of these patterns must be qualified by "*ceteris paribus*" clauses, as the following example illustrates: If someone wants a goal and believes that performing a certain action is a means to that goal, then *ceteris paribus* they will perform that action. For example, if someone wants a beer and believes that there is one in the kitchen, then he will go get one—unless, as the "*ceteris paribus*" clause signals, he does not want to miss any of the conversation, or he does not want to offend his guest by leaving in midsentence, or he does not want to drink beer in front of his mother-in-law, or he thinks he had better flee the house since it is on fire, etc. This pattern exhibits a special property which I will call "suppleness". Suppleness is involved in a distinctive kind of exceptions to the patterns in our mental lives—specifically, those exceptions that reflect our *ability to act appropriately* in the face of an open-ended range of contextual contingencies. These exceptions to the norm occur when we make *appropriate* adjustment to contingencies. The ability to adjust our behavior appropriately in context is a central component of the capacity for intelligent behavior.

A promising strategy for explaining mental suppleness is to follow the lead set by artificial life [14]. For there is a similar suppleness in vital processes such as metabolism, adaptation, and even flocking. For example, a flock maintains its cohesion not always but only for the most part, only *ceteris paribus*, for the cohesion can be broken when the flock flies into an obstacle (like a tree). In such a context, the best way to "preserve" the flock might be for the flock to divide into subflocks. Artificial life models of flocking exhibit just this sort of supple flocking behavior. Or consider another example concerning the process of adaptation itself. Successful adaptation depends on the ability to explore an

appropriate number of viable evolutionary alternatives; too many or too few can make adaptation difficult or even impossible. In other words, success requires striking a balance between the competing demands for “creativity” (trying new alternatives) and “memory” (retaining what has proved successful). Furthermore, as the context for evolution changes, the appropriate balance between creativity and memory can shift in a way that resists precise and exceptionless formulation. Nevertheless, artificial life models can show a supple flexibility in how they balance creativity and novelty.

Implications for the arts

Artificial life’s central aim is to develop a coherent theory of life in all its manifestations. It embraces the possibility of discovering life in unfamiliar settings and creating unfamiliar forms of life. In the long run artificial life will contribute to the development of practical adaptive systems in many fields of application, such as software development and management, design and manufacture of robots including distributed swarms of autonomous agents, automated trading in financial markets, pharmaceutical design, ecological sustainability, and extraterrestrial exploration. The economic potential of harnessing natural adaptive systems can be compared with cracking the genetic code. Natural adaptive systems vastly exceed the complexity of anything humans have yet created. Understanding and harnessing life’s adaptive creativity will spawn a wealth of new technologies and entrepreneurial opportunities.

Artificial life also has aesthetic applications. There are at least three ways in which artists might find artificial life useful. First, artificial life technology can be used for a variety of aesthetic purposes. They range from commercial applications in computer animations of life forms to new kinds of active art, evolving art, and interactive art (e.g., [3, 4, 15, 16]). Second, artificial life is radically changing human culture and technology, and art often responds to and comments on such changes (e.g., [17, 18]). Third, art has a long tradition of representing and responding to our understanding of nature, so new insights about life revealed by artificial life can spark new aesthetic objects (e.g., [19, 20, 21]).

Just as artificial life can be beneficial for artists, artists can provide complementary benefits to artificial life. For one thing, artists that use artificial life techniques and insights can be counted among the consumers of the product that artificial life produces, and one spur to producing better products is consumer demand. Scientists can also gain a broader perspective on their own scientific activity when artists explore the implications of the science and subject it to commentary and social criticism. Finally, human aesthetic activity is itself one distinctive manifestation of the creative potential contained within life. It would behoove those who want to understand nature’s creative potential to keep an eye on the latest aesthetic developments.

References

1. M. A. Bedau, J. McCaskill, N. Packard, S. Rasmussen, eds., *Artificial Life VII: Proceedings of the Seventh International Conference on the Simulation and Synthesis of Living Systems* (Cambridge: MIT Press, 2000).

2. Gary R. Greenfield, "Art and Artificial Life—A Coevolutionary Approach," in M. A. Bedau, J. S. McCaskill, N. H. Packard, and S. Rasmussen, eds., *Artificial Life VII* (Cambridge: MIT Press, 2000), pp. 529-536.
3. Eleonora Bilotta, Pietro Pantano, and Valerio Talarico, "Synthetic Harmonies: An Approach to Musical Semiosis by Means of Cellular Automata," in M. A. Bedau, J. S. McCaskill, N. H. Packard, and S. Rasmussen, eds., *Artificial Life VII* (Cambridge: MIT Press, 2000), pp. 537-546.
4. Christa Sommerer and Laurent Mignonneau, "Modeling Emergence of Complexity: The Application of Complex System and Origin of Life Theory to Interactive Art on the Internet," in M. A. Bedau, J. S. McCaskill, N. H. Packard, and S. Rasmussen, eds., *Artificial Life VII* (Cambridge: MIT Press, 2000), pp. 547-554.
5. Arantza Etxeberria, "Artificial Evolution: Creativity and the Possible," in M. A. Bedau, J. S. McCaskill, N. H. Packard, and S. Rasmussen, eds., *Artificial Life VII* (Cambridge: MIT Press, 2000), pp. 555-562.
6. C. G. Langton, "Artificial Life," in C. G. Langton, ed., *Artificial Life* (Redwood City, CA: Addison-Wesley, 1989).
7. E. R. Berlekamp, J. H. Conway, and R. K. Guy, *Winning ways for your mathematical plays, Vol. 2: Games in particular* (New York: Academic Press, 1982).
8. M. A. Bedau, J. S. McCaskill, N. H. Packard, S. Rasmussen, C. Adami, D. G. Green, T. Ikegami, K. Kaneko, and T. S. Ray, "Open Problems in Artificial Life," *Artificial Life* 6, 363-376 (2000).
9. Dawkins, R. 1976. *The Selfish Gene*. Oxford: Oxford University Press.
10. M. A. Bedau, "Weak Emergence," *Philosophical Perspectives* 11, 375-399 (1997).
11. S. J. Gould, *Wonderful Life: The Burgess Shale and the Nature of History* (New York: Norton, 1989).
12. D. C. Dennett, *Darwin's Dangerous Idea: Evolution and the Meanings of Life* (New York: Simon and Schuster, 1995).
13. Pattee, H. H. 1989. *Simulations, Realization, and Theories of Life*. In C. G. Langton, ed., *Artificial Life* (Redwood City, CA: Addison-Wesley, 1989), pp. 63-78.
14. Bedau, M. A. "Emergent Models of Supple Dynamics in Life and Mind," *Brain and Cognition* 34, 5-27 (1997).
15. Laurent Mignonneau and Christa Sommerer, "Creating Artificial Life for Interactive Art and Entertainment," *Leonardo*, Vol. 34, No. 4, 303 – 308 (2001).
16. Troy Innocent, "The Language of Iconica," *Leonardo*, Vol. 34, No. 3, 255–260 (2001).
17. Ursula Huws, "Nature, Technology and Art: The Emergence of a New Relationship?," *Leonardo*, Volume 33, Issue 1, 33-40 (2000).
18. Arike Ando, "What are People for?," *Leonardo*, Vol. 34, No. 5, 447-452 (2001).
19. Amy M. Youngs, "The Fine Art of Creating Life," *Leonardo*, Vol. 33, No. 5, 377-380 (2000).
20. Jane Prophet, "TechnoSphere: 'Real' Time, 'Artificial' Life," *Leonardo*, Vol. 34, No. 4, 309–312 (2001).
21. Mitchell Whitelaw, "The Abstract Organism: Towards a Prehistory for A-Life Art," *Leonardo*, Vol. 34, No. 4, 345–348 (2001).