ARTIFICIAL LIFE

Mark A. Bedau

Contemporary artificial life (also known as "ALife") is an interdisciplinary study of life and life-like processes. Its two most important qualities are that it focuses on the essential rather than the contingent features of living systems and that it attempts to understand living systems by artificially synthesizing extremely simple forms of them. These two qualities are connected. By synthesizing simple systems that are very life-like and yet very unfamiliar, artificial life constructively explores the boundaries of what is possible for life. At the moment, artificial life uses three different kinds of synthetic methods. "Soft" artificial life creates computer simulations or other purely digital constructions that exhibit life-like behavior. "Hard" artificial life produces hardware implementations of life-like systems. "Wet" artificial life involves the creation of life-like systems in a laboratory using biochemical materials.

Contemporary artificial life is vigorous and diverse. So this chapter's first goal is to convey what artificial life is like. It first briefly reviews the history of artificial life and illustrates the current research thrusts in contemporary "soft", "hard", and "wet" artificial life with respect to individual cells, whole organisms, and evolving populations. Artificial life also raises and informs a number of philosophical issues concerning such things as emergence, evolution, life, mind, and the ethics of creating new forms of life from scratch. This chapter's second goal is to illustrate these philosophical issues, discuss some of their complexities, and suggest the most promising avenues for making further progress.

1 HISTORY AND METHODOLOGY

Contemporary artificial life became known as such when Christopher Langton coined the phrase "artificial life" in the 1980s. Langton described artificial life as a study of life as it could be in any possible setting and he organized the first conference that explicitly recognized this study [Langton, 1989].

The intellectual roots of contemporary artificial life grow back to the first half of the twentieth century, and the two deepest roots reach to John von Neumann and Norbert Wiener. Von Neumann [1966] designed the first artificial life model (without referring to it as such) when he created his famous self-reproducing, computation-universal cellular automata.[1] Von Neumann tried to understand the

[1]A cellular automaton is a regular spatial lattice of "cells," each of which can be in one of a finite number of states. The lattice typically has 1, 2, or 3 spatial dimensions. The state of

Handbook of the Philosophy of Science. Volume 3: Philosophy of Biology
Volume editors: Mohan Matthen and Christopher Stephens. Handbook editors: Dov M. Gabbay, Paul Thagard and John Woods.

fundamental properties of living systems, especially self-reproduction and the evolution of complex adaptive structures, by constructing simple formal systems that exhibited those properties. At about the same time, Wiener [1948] started applying information theory and the analysis of self-regulatory processes (homeostasis) to the study of living systems. The abstract constructive methodology of cellular automata still typifies much artificial life, as does the abstract and material-independent methodology of information theory.

Artificial life has also been influenced by developments in traditional disciplines. Wet ALife clearly grows out of work in molecular biochemistry on the origin of life, and artificial life in general clearly benefits from a wealth of information about life on Earth. In addition, some models originally devised for specific biological phenomenon have subsequently been adopted and explored for other purposes by artificial life, e.g., the random Boolean networks originally introduced by Kauffman as a model of gene regulation networks.[2] Physics and mathematics, especially statistical mechanics and dynamical systems, have contributed the method of constructing simple model systems that have broad generality and permit quantitative analysis. Furthermore, the use of cellular automata as exemplars of complex systems [Wolfram, 1994] directly led to contemporary artificial life.

Much of the early work on artificial life was showcased at the Santa Fe Institute, an interdisciplinary research institution that helped put the study of complex systems on the map. Complex systems are composed of many elements that are simultaneously interacting with each other. Those in which the rules governing the elements are reshaped over time by some process of adaptation or learning are complex adaptive systems [Holland 1975/1992; 1995]. Artificial life focuses specifically on those complex systems that involve life, and these typically involve adaptation and learning.

Though artificial life differs from artificial intelligence, the two are connected through ALife's deep roots in computer science, especially artificial intelligence (AI) and machine learning. Notable here are John Holland's pioneering investigations of genetic algorithms [1975/1992].[3] The subjects of AI and artificial

---

each cell in the lattice is updated simultaneously in discrete time steps. Each cell is a finite state machine that outputs the next state of the cell given as input the states of the cells within some finite, local neighborhood of the lattice. Typically all cells in the lattice are governed by the same finite state machine, which typically is deterministic.

[2]Random Boolean networks consist of a finite collection of binary (ON, OFF) variables with randomly chosen input and output connections. The state of each variable at each step in discrete time is governed by some logical or Boolean function (AND, OR, etc.) of the states of variables that provide input to it. The network is started by randomly assigning states to each variable, and then the connections and functions in the network determine the successive state of each variable. Since the network is finite, it eventually reaches a state it has previously encountered, and from then on the network will forever repeat the same cycle of states. Different network states can end up in the same state cycle, so a state cycle is called an attractor.

[3]The genetic algorithm is machine learning technique loosely modeled on biological evolution. It treats learning the solution to a problem as a matter of competition among candidate problem solutions, with the best candidate solutions eventually winning. Potential solutions are encoded in an artificial chromosome, and an initial population of candidate solutions is created randomly. The quality or "fitness" of each solution is calculated by application of a "fitness function."

life overlap, since living and flourishing in a changing and uncertain environment requires at least rudimentary intelligence. Their methodologies are also similar, since both study natural phenomena by simulating and synthesizing them.

Nevertheless, there is an important difference between traditional symbolic AI and artificial life. 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 environment, and its decisions directly affect only its own local environment.

ALife's models characteristically follow this example from nature. The models themselves are bottom-up-specified parallel systems of simple agents interacting locally. The models are repeatedly iterated and the resulting global behavior is observed. Such lower-level models are sometimes said to be "agent-based" or "individual-based." The whole system's behavior is represented only indirectly. It arises out of interactions among directly represented parts ("agents" or "individuals") and their physical and social environment. This decentralized architecture shares important similarities with some newer trends in AI, including connectionism [Rumelhard and McClelland, 1986], multiagent AI [Rich and Knight, 1991], and evolutionary computation [Holland, 1975/1992; Mitchell, 1996].

An accurate and detailed sense of artificial life's central aims can be found in the unabashedly long-term grand challenges framed by the organizers of Artificial Life VII, the International Conference on Artificial Life that occurred at the new milennium [Bedau *et al.*, 2000]. The challenges fell into three broad categories concerning life's origin, its evolutionary potential, and its connection to mind and culture.

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.

For example, if the problem is to find the shortest route between two cities and a candidate solution is a specific itinerary, then the fitness function might be the sum of the distance of each segment in the itinerary and a solution's fitness is proportional to the reciprocal of its total distance. In effect, the fitness function is the "environment" to which the population adapts. A candidate solution's "genotype" is its chromosome, and its "phenotype" is its fitness. On analogy with natural selection, lower fitness candidates are then replaced in the population with new solutions modeled on higher fitness candidates. New candidates are generated by modifying earlier candidates with "mutations" that randomly change chromosomal elements and "crossover" events that combine pieces of two chromosomes. After reproducing variants of the most fit candidates for many generations, the population contains better and better solutions.

5. Explain how rules and symbols are generated from physical dynamics in living systems.

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.

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.

Some areas of artificial life are missing from the list, notably hard artificial life. This is simply an historical accident of who attended Artificial Life VII.

2 THREE ILLUSTRATIONS OF CONTEMPORARY ARTIFICIAL LIFE

Life exhibits complex adaptive behavior at many different levels of analysis, ranging from individual cells to whole organisms, and even to evolving ecologies. One can get a concrete feel for contemporary artificial life by considering a few illustrations of soft, hard, and wet artificial life. These examples illustrate artificial life's broad interdisciplinary nature, its synthetic methodology, and its concern with understanding the essential properties of living systems.

Artificial cells

The holy grail of "wet" artificial life is to create artificial cells out of biochemicals [Bedau *et al.*, 2000; Rasmmssen *et al.*, 2007]. Such artificial cells would be microscopic autonomously self-organizing and self-replicating entities built from simple organic and inorganic substances [Rasmussen *et al.*, 2004]. Although artificial, for all intents and purposes they would be alive, for they would maintain themselves, spontaneously organize and repair themselves, and adapt in an open-ended fashion to environmental contingencies.

There are two main motivations behind this research. One is pure science. If one could make artificial cells from scratch, especially using materials or methods that are not employed by nature, one would have dramatic proof of the possible molecular foundations of living systems. Artificial cells also have a broad practical appeal. Natural cells are much more complicated than anything yet produced by man, and many people believe that the next watershed in intelligent machines depends on bridging the gap between non-living and living matter [Brooks, 2001]. So, making artificial cells that organize and sustain themselves and continually adapt to their environment would open the door to future technologies with the impressive capacities of living systems.

What will artificial cells do? The initial functionality of these machines will be simply to move through a fluid and process chemicals. To do this flexibly and resiliently involves solving the basic functions of self-maintenance, autonomous control of chemical processing, autonomous control of mobility, and self-replication. Artificial cells will simultaneously solve these tasks by integrating an artificial metabolism with the means of growth and self-reproduction, and localizing these chemical systems by producing some container. Thus, artificial cells will have biochemical systems that construct and repair the system's container (e.g., cell walls), systems that copy the information-bearing molecules that encode and direct cellular processes (genes), and systems that synthesize the materials for cellular self-assembly and regeneration (a metabolism). All life in nature depends on the coordination of these three processes. The first artificial cells will have extremely simple versions of them.

Nobody has yet created an artificial cell, but research toward this goal is actively under way. Two main approaches are being pursued. Human genome pioneer J. Craig Venter and Nobel Prize winner Hamilton Smith recently publicized their intention to create a partly man-made artificial cell, with support from the US Energy Department [Gillis, 2002]. Venter and Smith are using the top-down strategy of simplifying the genome of the simplest existing cell with the smallest genome: *Mycoplasma genitalium* [Fraser *et al.*, 1995; Hutchison *et al.*,1999]. This top-down approach has the virtue that it can simply borrow the biological wisdom embodied in *Mycoplasma* biochemistry. It has the corresponding disadvantage that its insights will be limited by various contingencies of *Mycoplasma's* evolution.

The other approach to making artificial cells is bottom up: to build more and more complex physiochemical systems incorporating more and more life-like properties. Szostak, Bartell, and Luisi [2001] and Pohoril and Deamer [2002] describe bottom-up strategies that are strongly inspired by the lipid bilayer membranes and nucleic acid chemistry found in existing cells. Lipid vesicles have been shown to grow and reproduce in the laboratory [Walde *et al.*, 1994; Menger and Angelova, 1998]. The main challenge of this bottom-up strategy is that there is no known chemical path for synthesizing DNA or RNA that is sufficiently complex to encode the minimal molecular functions needed by such artificial cells. Rasmussen *et al.*,[2003] have proposed a simpler and much less natural bottom-up approach in which PNA chemistry [Nielsen *et al.*,1991] replaces RNA chemistry and lipid micelles replace vesicles.

Autonomous agents

Much work in artificial life at the level of multicellular organisms has occurred in "hard" artificial life concerned with various forms of autonomous physical agents or robots. This is artificial life's most direct overlap with artificial intelligence. Hard artificial life tries to synthesize autonomous adaptive and intelligent behavior in the real world. It contrasts with traditional artificial intelligence and robotics by exploiting biological inspiration whenever possible, and also by aiming to synthesize behaviors characteristic of much simpler organisms than humans. One of the tricks is to let the physical environment be largely responsible for generating the behavior. Rather than relying on an elaborate and detailed internal representation of the external environment, the behavior of biologically-inspired robotics quite directly depends on the system's sensory input from its immediate environment. With the right sensory-motor connections, a system can quickly and intelligently navigate in complex and unpredictable environments. This so-called "behavior-based" robotics has been pioneered by Rodney Brooks [1989; 1990; 1991]. The initial successes involved insect-like robots and it has since been extended to humanoid robots [Adams *et al.*,2000]. Another trick is to let the physical materials out of which the robot is embodied to automatically provide as much functionality as possible [Pfeifer and Scheier, 2001].

Even with behavior-based robots, design of intelligent autonomous agents is difficult because it involves creating the right interconnections among many complex components. The intelligent autonomous agents found in nature are all alive, and their design was achieved spontaneously through an evolutionary process. So artificial life uses evolution to design autonomous agents [Cliff *et al.*, 1993]. To this end, genetic algorithms have been used to design many aspects of robots, including control systems and sensors [Nolfi and Floreano, 2000; 2002].

In natural autonomous agents, the control system is tightly coupled with morphology. Sims [1994] showed ten years ago how to recreate this interconnection when he simultaneously coevolved simulated creatures' controllers, sensors, and morphology, but he relied on special-purpose software running on extremely expensive supercomputers. More recent advances in hardware and software have enabled this line of research to be pursued with off-the-shelf software running on laptops [Taylor and Massey, 2001]. This work, like Sims's, involves simulations alone. Jordan Pollack and his students have taken the next step and used similar methods to develop actual physical robots. They have connected simulated co-evolution of controllers and morphology with off-the-shelf rapid prototyping technology, allowing their evolutionary design to be automatically implemented in the real world [Lipson and Pollack 2000; Pollack *et al.*, 2001].

Digital evolution

Implementing evolving systems in software is the most practical and constructive way to study many issues about evolving systems, and this "soft" approach has been a dominant trend in artificial life. One of the first significant achievement of spontaneous evolution in a digital medium was Tierra [Ray, 1992], which is simply a population of simple, self-replicating computer programs that exist in computer memory and consume CPU time. A Tierran genotype consists of a string of machine code, and each Tierran "creature" is a instance of some Tierran genotype. A simulation starts when a single self-replicating program, the ancestor, is placed in computer memory and left to replicate. The ancestor and its descendants repeatedly replicate until computer memory is teeming with creatures that all share the same ancestral genotype. Older creatures are continually removed from memory to create space for new descendants. Errors (mutations) sometimes occur, and the population of programs evolves by natural selection. If a mutation allows a program to replicate faster, that genotype tends to spread through the population. Over time, the ecology of Tierran genotypes becomes remarkably diverse. Quickly reproducing parasites that exploit a host's genetic code evolve, and the co-evolution between hoses and parasites spurs the evolution of parasite-resistance and new forms of parasitism. After millions of CPU cycles of this co-evolutionary arms race, Tierra often contains many kinds of creatures exhibiting a variety of competitive and cooperative ecological relationships.

Life has exhibited a remarkable growth in complexity over its evolutionary history. Simple prokaryotic one-celled life led to more complex eukaryotic one-celled life, which led to multicellular life, then to large-bodied vertebrate creatures with complex sensory processing capacities, and ultimately to highly intelligent creatures that use language and develop sophisticated technology — those creatures at the central focus of cognitive science. Although some forms of life remain evolutionary stable for millions of years (e.g, coelacanths and sharks), the apparently open-ended growth in complexity of the most complex organisms is intriguing and enigmatic. Much effort in artificial life is directed toward creating a system that shows how this kind of open-ended evolutionary progress is possible, even in principle. Digital evolution in Tierra does not do this, for significant evolutionary change eventually peters out. Ray has tried to address these limitations by making the Tierra environments much larger and more heterogeneous and by making the ancestral Tierran creatures significantly more complex (in effect, giving them multiple cell types). By allowing Tierran creatures to migrate from machine to machine over the Internet, looking for unused resources and for more favorable local niches, Ray has found signs that they evolve new types of cells [Ray, 2000]. Furthermore, when Tierra is modified so that creatures are rewarded for performing complex arithmetic operations on numbers they find in their local environment, evolution produces the expected increase in genetic complexity [Adami *et al.*, 2000; Lenski *et al.*, 2003]. However, as with the original version of Tierra, these evolutionary progressions eventually stop.

Hillis [1992] demonstrated that co-evolution can spur evolutionary progress, and co-evolutionary arms races might help drive continual evolutionary progression by continually changing the environment for evolution. But the original and most modified versions of Tierra involve some form of co-evolution and yet the environment eventually becomes essentially stable, so there is probably more to the story. Further progress on open-ended evolution would be aided by quantitative comparisons across different artificial and natural evolving systems. Bedau and Packard and their collaborators have taken a step in that direction by defining and studying evolutionary activity statistics. Comparing data from different artificial and natural evolving systems suggests that there are qualitatively different classes of evolutionary dynamics, and no known artificial system generates the kind of evolutionary dynamics exhibited by the biosphere [Bedau *et al.*, 1997; 1998]. We are apparently still missing critical insights about the mechanisms by which evolution continually creates the new kinds of environments that continually elicit new kinds of adaptations.

3 PHILOSOPHICAL IMPLICATIONS OF ARTIFICIAL LIFE

The scientific and engineering of artificial life has rich implications for a number of broad philosophical issues. This section illustrates these implications for a few philosophical issues.

Philosophy and artificial life are natural intellectual partners, for three reasons. By creating wholly new kinds of life-like phenomena, artificial life continually forces us to reexamine and reassess what it is to be alive, adaptive, intelligent, creative, etc. In addition, both philosophy and artificial life seek to understand phenomena at a level of generality that ignores contingencies and reveals essential natures.

Finally, artificial life's computational methodology is a direct and natural extension of philosophy's traditional methodology of *a priori* thought experiment. Aiming to capture the simple essence of vital processes, artificial life abstracts away as many details of living systems as possible. The resulting artificial life models are thought experiments that are explored by actually syntehsizing instances of the models. Like the traditional armchair thought experiments, artificial life simulations attempt to answer "What if X?" questions, but the premises they pose are too complicated to be understood except by synthesizing them. These synthetic methods are often computational (in soft artificial life), but they sometimes involve constructing novel hardware (in hard artificial life) or even construcing novel systems by biochemical means (in wet artificial life). In each case, the motivation is the same: the behavior of the system cannot be determined except through direct experience. These constructive thought experiments bring a new kind of clarity and constructive evidence to philosophy.

Emergence

One of life's amazing features is how the whole is more than the sum of the parts. This is called emergence [Bedau and Humphries, 2007]. As a general definition, emergent phenomena are macro and micro phenomena that are related so that the macro both depends on and is autonomous from the underlying micro 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 give emergence a constructive role in scientific explanations of emergent macro phenomena like life and mind.

The aggregate global behavior of complex systems studied in artificial life offers a new view of emergence, so-called "weak" emergence [Bedau, 1997; 2003], in contrast to the "strong" emergence that involves in principle irreducibility of macro from micro [Kim, 1999]. 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 potentially all of the micro-level effects. 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 by driving only the local micro processes. 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. Weak emergence is common in complex systems found in nature, and artificial life's models also exhibit it. The unpredictability and unexplainability of weak emergent phenomena comes from the myriad, non-linear and context-dependent local micro-level interactions that drive the systems. Emergent phenomena can have causal powers on this view, but only by aggregating micro-level causal powers. There is nothing inconsistent or metaphysically illegitimate about underlying processes constituting and generating phenomena in this way by iteration and aggregation. Furthermore, weak emergence is rampant in scientific explanations of exactly the natural phenomena that apparently involve emergence, like life and mind.

This shows how artificial life will play an active role in future philosophical debates about emergence, as well as related notions like explanation, reduction, and hierarchy. Living systems are a paradigm example 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 can emerge from what by constructively exploring what is possible.

Evolution

As noted above, the evolution of life has produced 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 raises a deep question about evolution's creative potential: Does evolution have an inherent tendency to create greater and greater adaptive complexity, or is the increasing complexity of life just a contingent and accidental by-product of evolution? This question has attracted the attention of both philosophers and biologists.

Stephen Jay Gould [1989] 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 [1995] 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.

I am dubious about both answers, for the same reason. Gould's thought experiment of replaying the tape of life is exactly the right way to investigate the scope of contingency and necessity in evolution. But neither Gould nor Dennett actually carry out the experiment. Instead, they just speculate about what would happen were one to do so. Extensive experience in artificial life has shown time and again that armchair speculations about the outcome of such thought experiments are highly fallible.

We cannot actually replay life's tape, of course, since we cannot roll back time to an earlier biosphere. But we can do the next best thing and synthesize artificial biospheres that are like the real biosphere in relevant respects, and then observe their behavior. The easiest artificial biospheres to construct are simply software systems. The behavior of vast numbers of instances of these software systems can be observed, and very robust generalizations discovered. Obviously, soft artificial life can constructively contribute to this project for it is precisely in the business of creating and studying such systems.

Of course, there is no way to recreate all the conditions of early life on Earth, including the right environment and distribution of species (including the absence of humans). But replaying life's tape does not require returning to life's actual origin. Instead, the subsequent evolution of an entirely different biosphere would provide even more information about evolution's inherent creative potential, as long as that biosphere's creative evolutionary potential was sufficiently open. So artificial life software systems that are analogous to Earth's early life in relevant respects could serve to replay life's tape.

It is far from trivial to create systems displaying the richness of real life. In fact, no one has yet devised a system that exhibits the continual open-ended evolution that seems to be happening in the biosphere (recall above). Achieving this goal is a key open problem in artificial life, related to its sixth grand challenge. The final evaluation of conjectures like Gould's and Dennett's about evolution's inherent creativity must await artificial life's progress on replaying the tape of life.

Life

Life seems to be one of the most basic categories of actual natural phenomena. Yet it is notoriously difficult to say what life is, exactly. The fact is that today we know of no set of individually necessary and jointly sufficient conditions for life. Nevertheless, there is broad agreement about the distinctive hallmarks that life forms share. These hallmarks include being complex adaptive organization sustained by metabolic processes, produced by natural selection through a process involving random variation and historical contingency, and producing qualitative and unpredictable phenomena involving unique and variable individuals containing unique macromolecules [Mayr, 1982]. The characteristic coexistence of these hallmarks is striking, and it is a reason to suspect there is a unified explanation of life. But appearances might be deceptive. Vital phenomena might have no unified explanation and life might not be a basic category of natural phenomena. Skeptics like Sober [1992] think that the question of the nature of life, in general, has no interesting answer. But one should retreat to skepticism, if at all, only as a last resort. Those searching for extraterrestrial life, those searching for the origin of life on Earth, and those attempting to synthesize life in artificial media typically are betting that there is an interesting explanation of life in general.

Philosophers from Aristotle to Kant have investigated 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 is especially revitalized the question today. One can simulate or synthesize living systems only if one has some idea what life is. Artificial life's aim to discern the essence of life encourages liberal experimentation with novel life-like organizations and processes. Thus, artificial life 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 hallmarks 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 [Bedau, 1998]. And artificial life provides a constructive setting in which to explore the empirical implications of different conceptions of life.

Motivated partly by experience in artificial life, Bedau [1996; 1998] has recently argued for the admittedly unintuitive view that life in the most fundamental sense is displayed by a system that is continually exhibiting creative evolution. Organisms would then be explained as alive in a derivative sense, by virtue of their connection with and role in an evolving system. One virtue of the conception of life as evolution is that it explains why Mayr's hallmarks of life coexist in nature. We would expect life to involve the operation of natural selection producing complex adaptive organization in historically connected organisms with evolved genetic programs. The random variation and historical contingency in the evolutionary process explains why living phenomena are especially qualitative and unpredictable and involve unique and variable individuals with frozen accidents like chemically unique macromolecules. This view can also explain why metabolism is so important in living systems, for a metabolism is a physically necessary prerequisite in any system that can sustain itself long enough to adapt and evolve. In addition, this view accounts for four of the main puzzles about life [Bedau, 1998].

There are two main objections to this view of life. First, one might think it is entirely contingent that life forms were produced by an evolutionary process. The Biblical story of Adam and Eve shows that is easy to imagine life forms in the absence of any evolutionary process. But it is not clear that this is anything more than a philosophical fantasy, unrelated to what would actually happen anywhere in the real world. A second objection calls attention to the fact that some evolving systems seem devoid of life. Viruses and prions evolve but are dubiously alive, and cultural and technological evolution provides even starker counterexamples. One response to this sort of worry is to bite the bullet and claim that these kinds of evolving systems actually deserve to be considered to be alive, at least in the primary sense. It is important to realize that the project of uncovering the nature of life is not simply to analyze our concept of life. Our concepts are historical accidents that might be unsuited to the underlying categories in nature. It could turn out that the fundamental process that produces the familiar phenomena of life is essentially the same as the process that produces phenomena that we do not today recognize to involve life. If so, then learning this would reveal a new deep truth about life.

Artificial life has called special attention to the question whether 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 [Pattee, 1989]. 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

All forms of life have mental capacities, broadly speaking [Dennett, 1997]. They are sensitive to the environment in various ways, and this environmental sensitivity affects their behavior in various ways. Furthermore, the sophistication of these mental capacities seems to correspond to the complexity of those forms of life. So it is natural to ask if there is an interesting connection between life and mind. For example, life and mind would be strikingly unified if Beer [1990, p. 11] is right that "it is adaptive behavior, the ... ability to cope with the complex, dynamic, unpredictable world in which we live, that is, in fact, fundamental [to intelligence itself]" (see also [Maturana and Varela, 1987/1992]). Since all forms of life must cope in one way or another with a complex, dynamic, and unpredictable world, perhaps this adaptive flexibility intrinsically connects life and mind. Understanding the ways in which life and mind are connected is one of the basic puzzles about life.

Many mental capacities are certainly adaptations produced by the process of evolution of living organisms. This is sufficient for a certain shallow connection between life and mind. Aristotle's view that there is an intrinsic conceptual unity of life and mind goes much deeper. For Aristotle, an organism's mental activity consists of the exercise of various internal capacities and potentialities (its "soul"), and being alive consists of the exercise of those very same capacities and potentialities [Code and Moravcsik, 1992]. The theory of life as continual creative evolution (recall above) implies a related view, according to which the mind as an expression of a process (creative evolution) that is also the definitive feature of life. One specific way to make this argument is by appealing to the suppleness of life and mind [Bedau, 1977a; 1999].

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 that 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 of artificial life, because 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 [Bedau, 1999]. The suppleness of both life and mind suggests that they might be two different manifestations of essentially the same kind of underlying process, two sides of the same coin. This suggestion is a very open question today, but it shows how artificial life might deeply unify life and mind.

Ethics

Both the process of pursuing artificial life research and the scientific and practical products of that research process raise complicated ethical issues [Bedau *et al.*, 2000]. These issues include four broad categories: (i) the sanctity of the biosphere, (ii) the sanctity of human life, (iii) the responsible treatment of newly generated life forms, and (iv) the risks of using artificial life technology.

Artificial life's ethical issues somewhat resemble those concerning animal experimentation, genetic engineering, and artificial intelligence, and the extensive literature on those topics may guide exploration of the ethical issues in artificial life. On the other hand, creating novel forms of life and interacting with them in novel ways will place us in increasingly uncharted ethical terrain.

Perhaps the most vivid ethical issues arise from wet artificial life efforts aimed ultimately at making new forms of life in the laboratory from scratch [Bedau and Parke, 2007]. These efforts can be expected to generate public concern. Some will object that creating artificial cells is unnatural or fails to give life due respect [Kass, 2002; Cho, 1999], or that it involves playing God [Cho, 1999]. One main driver for these ethical concerns is the fact that creating new forms of life will inevitably involve what I call deciding "in the dark" [Bedau and Triant, 2007]. Decisions "in the dark" are those we have to make even though we are largely ignorant about their possible consequences. New and revolutionary technologies, such as genetic engineering and nanotechnology, are allowing us to change our environment at an accelerating rate. Much of this change is being driven by the private economic interests of large international corporations. But the unprecedented nature of these technological innovations makes their implications for human health and the environment extremely difficult to forecast.

Decision theory [Raiffa, 1968; Resnick, 1987] has a well-developed arsenal for confronting what are known as decisions "under risk" and decisions "under ignorance or uncertainty," but it is unequipped to help with decisions in the dark. Decision theory approaches a decision in a given context by tabulating the different possible actions that could be made in that context, determining the likely consequences of each action, determining the likely social utility of each consequence, and then analyzing this table by calculating such things as each action's expected utility. Decisions "under risk" are those in which the likely consequences of the actions are uncertain and can only be assigned a probability, and decisions "under ignorance or uncertainty" are those in which even the probabilities of the consequences are unknown. In both kinds of decisions, however, the consequences of different courses of action can be tabulated. Decisions "in the dark" are different in just this respect: We are ignorant about even the possible outcomes of our actions, so we cannot even construct a decision table. So contemporary decision theory has no advice to offer about such decisions.

Yet technological innovations are increasingly forcing society to make decisions in the dark. Genetic engineering and nanotechnology are two examples. Recombinant DNA technology and advances in self-assembling molecular systems are now realizing undreamt of new bio- and nanotechnologies, and governments in most developed countries are betting vast economic stakes on the bio-nano future. But at the same time, their risks are also causing growing alarm. Genetically modified foods are now anathema throughout Europe, and Bill Joy created a stir when he described the dangers of combining biotechnology with nanotechnology in such things as artificial cells [Joy, 2000]. Because of the revolutionary novelty of these technologies, it is impossible for us to know the likely consequences of their development. Yet we nevertheless face choices today about whether and how to develop them, whether and how to regulate them, etc. We have to make these decisions in the dark.

Society today has two main methods for tackling decisions in the dark: risk analysis and the Precautionary Principle. Growing out of decision theory, risk analysis is the primary method by which large organizations and public agencies (e.g., the EPA and the FDA) make decisions with major social and economic implications [Morgan and Henrion, 1990; Wilson and Crouch, 2001, Ropeik and Gray, 2002]. For example, top officials in the U.S. Department of Agriculture cited a Harvard Center for Risk Analysis study to justify FDA inaction about mad cow disease. But it is unclear whether risk analysis can adequately overcome decision theory's shortcomings regarding decisions in the dark.

Much contemporary discussion of genetic engineering and nanotechnology is influenced by the Precautionary Principle, which states that we should ban new technologies that might create significant risks even if we lack clear evidence of such risks [Raffensperger and Tickner, 1999; Morris, 2000]. The Precautionary Principle is designed to apply precisely to situations in which society is in the dark, and it is playing an increasing role in international law, appearing in over a dozen international treaties and agreements (e.g., the Rio Declaration from the 1992 United Nations Conference on Environment and Development). But the Principle is controversial because it seems to ignore the possible benefits of new technologies.

The creation of new forms of life from scratch will create exciting new opportunities. It will also create new responsibilities. The choices society will confront will be especially difficult, because they will require deciding in the dark. Philosophers have a special expertise for helping think through these novel and consequential issues raised by wet artificial life.

4 CONCLUSIONS

This brief survey of the scientific and philosophical implications of contemporary artificial life should allay some pervasive misconceptions. The primary activity in artificial life today is not to produce toy models superficially reminiscent of life. Indeed, software creations comprise only one of its three synthetic methods. Artificial life does aim to create life-like behavior in artificial systems, to be sure, but the point of this is to uncover the essential properties of living systems, wherever they might exist in nature. The potential fruits of such insights are not just theoretical; they also promise to unlock the door to what could be literally called "living technology." Pursing this goal involves interdisciplinary collaboration as well as connection with the traditional sciences such as biology and chemistry. Increasingly empirical and rigorous, artificial life has made incremental advances toward a broad and ambitious agenda. But the extent to which it will achieve this agenda remains an open question.

Artificial life is not just science and engineering. It is also an important new tool for philosophy. In fact, the interests and methods of artificial life and philosophy overlap in a natural way, illustrating how the sciences and the humanities can work together in the pursuit of shared goals. If artificial life is successful in creating wholly new forms of life, it will also have a hand in changing the nature of the world in which we live. In any case, it is clear that artificial life will continue to have a significant and distinctively constructive impact on a wide variety of old and new philosophical questions.

BIBLIOGRAPHY

[Adami *et al.*, 2000] C. Adami, C. Ofria, and T. C. Collier. Evolution of biological complexity. *Proc. Nat. Acad. Sci. USA* 97, 4463-4468, 2000.

[Adams *et al.*, 2000] B. Adams, C. Breazeal, R. A. Brooks, and B. Scassellati. *IEEE Intell. Syst. Applic.* 15, 25-31, 2000.

[Bedau, 1996] M. A. Bedau. The nature of life. In Margaret Boden, ed., *The philosophy of artificial life* (New York: Oxford University Press), pp. 332-357, 1996.

[Bedau, 1997] M. A. Bedau. Weak emergence. *Philosophical Perspectives* 11, 375-399, 1997.

[Bedau, 1997a] M. A. Bedau. Emergent models of supple dynamics in life and mind. *Brain and Cognition* 34: 5-27, 1997.

[Bedau, 1998] M. A. Bedau. Four puzzles about life. *Artificial Life* 4: 125-140, 1998.

[Bedau, 1999] M. A. Bedau. Supple laws in biology and psychology. In V. Hardcastle, ed., *Where biology meets psychology: philosophical essays* (Cambridge: MIT Press) pp. 287-302, 1999.

[Bedau, 2003] M. A. Bedau. Downward causation and autonomy in weak emergence. *Principia Revista Inernacional de Epistemologica* 6, 5-50, 2003.

[Bedau *et al.*, 1997] M. A. Bedau, e. Snyder, C. T. Brown, and N. H. Packard. A comparison of evolutionary activity in artificial evolving systems and the biosphere. In *Proceedings of the Fourth European Conference on Artificial Life* (Husbands, P. and Harvey, I., eds.), pp. 125-134, MIT Press, 1997.

[Bedau *et al.*, 1998] M. A. Bedau, E. Snyder, and N. H. Packard. A classification of long-term evolutionary dynamics. In *Artificial Life VI* (Adami, C. et al., eds.), pp. 228-237, MIT Press, 1998.

[Bedau *et al.*, 2000] M. A. Bedau, J. S. McCaskill, N. H. Packard, S. Raasmussen, C. Adami, D. G. Green, T. Ikegami, K. Kaneko, and T. S. Ray. Open problems in artificial life. *Artificial Life* 6, 363-376, 2000.

[Bedau and Humphries, 2007] M. A. Bedau and P. Humphries, eds. *Emergence: contemporary readings in science and philosophy.* MIT Press, 2007.

[Bedau and Parke, 2007] M. A. Bedau and E. Parke. *Our Future with Protocells: The Social and Ethical Implications of the Creation of Living Technology.* Cambridge: MIT Press, 2007.

[Bedau and Triant, 2007] M. A. Bedau and M. Triant. Social and ethical implications of artificial cells. In M. A. Bedau and E. Parke, eds., *Social and ethical perspectives on protocells.* MIT Press, 2007.

[Beer, 1990] R. D. Beer. *Intelligence as adaptive behavior: an experiment in computational neuroethology* (Boston: Academic Press), 1990.

[Brooks, 2001] R. Brooks. The relationship between matter and life. *Nature* 409, 409-411, 2001.

[Cho *et al.*, 1999] M. K. Cho, D. Magnus, A. L. Caplan D. McGee, and the Ethics of Genomics Group. Ethical considerations in synthesizing a minimal genome. *Science* 286, 2087, 1999.

[Cliff *et al.*, 1993] D. Cliff, I. Harvey, and P. Husbands. Explorations in evolutionary robotics. *Adaptive Behavior* 2, 73-110, 1993.

[Code and Moravcsik, 1992] A. Code and J. Moravcsik. Explaining various forms of living. In Nussbaum and Rorty, *Essays on Aristotle's* De Anima (Oxford: Clarendon Press), pp. 129-145, 1992.

[Dawkins, 1976] R. Dawkins. *The selfish gene.* Oxford: Oxford University Press, 1976.

[Dennett, 1995] D. C. Dennett. *Darwin's dangerous idea: evolution and the meanings of life* (New York: Simon and Schuster), 1995.

[Dennett, 1997] D. C. Dennett. *Kinds of Minds: Towards an Understanding of Consciousness*. New York: Basic Books, 1997.

[Fraser *et al.*, 1995] C. M. Fraser, J. D. Gocayne, O. White, M. D. Adams, R. A. Clayton, R. D. Fleischmann, C. J. Bult, A.R. Kerlavage, G. Sutton, J. M. Kelley, J. L. Firtchman, J. F. Weidman, K. V. Small, M. Sandusky, J. Furhmann, D. Nguyen, T. R. Utterback, D. M. Saudek, C. A. Phillips, J. M. Merrick, J.-F. Tomb, B. A. Dougherty, K. F. Bott, P.-C. Hu, T. S. Lucier. S. N. Peterson, H. O. Smith, C. A. Hutchinson III, and J. C. Venter. The minimal gene component of *Mycoplasma genetalium*. *Science* 270, 397-403, 1995.

[Gillis, 2002] J. Gillis. Scientists planning to make new form of life. *Washington Post*, November 21, A01, 2002.

[Gould, 1989] S. J. Gould. *Wonderful life: the Burgess shale and the nature of history* (New York: Norton, 1989).

[Hillis, 1992] D. Hillis. Co-evolving parasites improve simulated evolution as an optimization procedure. In *Artificial Life II* (Langton, C.G. et al., eds.), pp. 313-324, Addison-Wesley, 1992.

[Holland, 1975/1992] J. H. Holland. *Adaptation in natural and artificial systems: an introductory analysis with applications to biology, control, and artificial intelligence*, University of Michigan Press; second expanded edition, MIT Press, 1975/1992.

[Holland, 1995] J. H. Holland. *Hidden order: how adaptation builds complexity*, Helix Books, 1995.

[Hutchison *et al.*, 1999] C. A. Hutchison III, S. N. Peterson, S. R. Gill, R. T. Cline, O. White, C. M. Fraser, H. O. Smith, and J. C. Venter. Global transposon mutagenesis and a minimal Mycoplasma genome. *Science* 286, 2165-2169, 1999.

[Joy, 2000] Joy, B. 2000. Why the future does not need us. *Wired* 8 (April). Available on the web at `http://www.wired.com/wired/archive/8.04/joy.html` [last accessed June 2003].

[Kass, 2002] L. R. Kass. . *Life, liberty, and the defense of dignity: the challenge for bioethics*. San Francisco: Encounter Books, 2002.

[Kim, 1999] J. Kim. Making sense of emergence. *Philosophical Studies* 95, 3-36, 1999.

[Langton, 1989] C. Langton, ed. *Artificial life* (Redwood City, CA: Addison-Wesley), 1999.

[Lenski *et al.*, 2003] R. E. Lenski, C. Ofria, R. T. Pennock, and C. Adami. The evolutionary origin of complex features. *Nature* 423, 139-144, 2003.

[Lipson and Pollack, 2000] H. Lipson and J. B. Pollack. Automatic design and manufacture of robotic lifeforms. *Nature* 406, 974-978, 2000.

[Maturana and Varela, 1987/1992] H. R. Maturana and F. J. Varela. *The tree of knowledge: the biological roots of human understanding*, 1987. Revised edition (Boston: Shambhala), 1992.

[Mayr, 1982] E. Mayr. *The growth of biological thought* (Cambridge, MA: Harvard University Press), 1982.

[Menger and Angelova, 1998] F. M. Menger and M. I. Angelova. Giant vesicles: imitating cytological processes of cell membranes. *Accounts of Chemical Research* 31, 789-797, 1998.

[Mitchell, 1996] M. Mitchell. *An introduction to genetic algorithms* (Cambridge: MIT Press), 1996.

[Morgan and Henrion, 1990] M. G. Morgan and M. Henrion. *Uncertainty: A Guide to Dealing with Uncertainty in Quantitative Risk and Policy Analysis*. Cambridge: Cambridge University Press, 1990.

[Morris, 2000] J. Morris, ed. *Rethinking risk and the Precautionary Principle*. Oxford: Butterworth-Heinemann, 2000.

[Nielsen *et al.*, 1991] P. E. Nielsen, M. Egholm, R. H. Berg, and O. Buchardt. Sequence-selective recognition of DNA by strand displacement with a thymine-substituted polyamide. *Science* 254, 1497, 1991.

[Nolfi and Floreano, 2000] S. Nolfi and D. Floreano. *Evolutionary robotics: the biology, intelligence, and technology of self-organizing machines* (Cambridge: MIT Press), 2000.

[Nolfi and Floreano, 2002] S. Nolfi and D. Floreano. Synthesis of autonomous robots through evolution. *Trends in Cognitive Science* 6, 31-37, 2002.

[Pattee, 1989] H. H. Pattee. Simulations, realization, and theories of Life. In C. G. Langton, ed., *Artificial life* (Redwood City, CA: Addison-Wesley), pp. 63-78, 1989.

[Pfeifer and Scheier, 2001] R. Pfeifer and C. Scheier. *Understanding intelligence*, 2nd printing (paperback edition). Cambridge, Mass.: MIT Press, 2001.

[Pohorille and Deamer, 2002] A. Pohorille and D. Deamer. Artificial cells: prospects for biotechnology. *Trends in Biotechnology* 20, 123-128, 2002.

[Pollack *et al.*, 2001] J. B. Pollack, H. Lipson, G. Hornby, and P. Funes. Three generations of automatically designed robots. *Artificial Life* 7, 215-223. 2001.

[Raiffa, 1968] H. Raiffa. *Decision Analysis: Introductory Lectures on Choices under Uncertainty*. Reading, Mass: Addison-Wesley, 1968.

[Raffensperger and Tickner, 1999] C. Raffensperger and J. Tickner, eds. *Protecting public health and the environment: implementing the Precautionary Principle*. Washington, D.C.: Island Press, 1999.

[Rasmussen *et al.*, 2003] S. Rasmussen, L. Chen, M. Nilsson, and S. Abe. Bridging nonliving and living matter. *Artificial Life* 9, 269-316, 2003.

[Rasmussen *et al.*, 2004] S. Rasmussen, L. Chen, D. Deamer, D. Krakauer, N. H. Packard, P. F. Stadler, and M. A. Bedau. Transitions from nonliving to living matter. *Science* 303: 963-965, 2004.

[Rasmussen *et al.*, 2007] S. Rasmussen, M. A. Bedau, S., L. Chen, D. Deamer, D. Krakauer, N. H. Packard, and P. F. Stadler. *Protocells: Bridging Nonliving and Living Matter.* MIT Press, 2007.

[Ray, 1992] T. S. Ray. An approach to the synthesis of life. In *Artificial life II* (Langton, C.G. et al., eds.), pp. 371-408, Addison-Wesley, 1992.

[Ray, 2000] T. S. Ray. In *Artificial life VII* (Bedau, M.A. et al., eds.), MIT Press, 2000.

[Resnick, 1987] M. Resnick. *Choices: an introduction to decision theory.* Minneapolis: University of Minnesota Press, 1987.

[Rich and Knight, 1991] E. Rich and K. Knight. *Artificial intelligence*, McGraw-Hill, 1991.

[Ropeik and Gray, 2002] D. Ropeik and G. Gray. *Risk: a practical guide to deciding what's really safe and what's really dangerous in the world around you.* Boston: Houghton Mifflin, 2002.

[Rumelhart and McClelland, 1986] D. E. Rumelhart and J. L. McClelland. *Parallel distributed processing: explorations in the microstructure of cognition*, 2 Vols., MIT Press, 1986.

[Sims, 1994] K. Sims. In *Artificial life IV* (Brooks, R. and Maes, P., eds.), pp. 28-39, MIT Press, 1994.

[Sober, 1992] E. Sober. Learning from functionalism — prospects for strong artificial life. In Langton et al. (1992), *Artificial life II* (Redwood City, CA: Addison Wesley), pp. 749-765, 1992.

[Szostak *et al.*, 20011] J. W. Szostak, D. P. Bartel, and P. L. Luisi. Synthesizing life. *Nature* 409, 387–390, 2001.

[Taylor and Massey, 2001] T. Taylor and C. Massey. Recent developments in the evolution of morphologies and controllers for physically simulated creatures. *Artificial Life* 7, 77-87, 2001.

[von Neumann, 1966] J. Von Neumann. *Theory of self-reproducing automata* (Urbana-Champagne: University of Illinois Press), 1966.

[Walde *et al.*, 1994] P. Walde, R. Wick, M. Fresta, A. Mangone, and P. L. Luisi. Autopoietic self-reproduction of fatty acid vesicles. *J. Am. Chem. Soc.* 116, 11649-116454, 1994.

[Wiener, 1948] N. Wiener. *Cybernetics, or control and communication in the animal and the machine*, Wiley, 1948.

[Wilson and Crouch, 2001] R. Wilson and E. A. C. Crouch. *Risk-benefit analysis.* Cambridge: Harvard University Press, 2001.

[Wolfram, 1994] S. Wolfram. *Cellular automata and complexity*, Addison-Wesley, 1994.