

Open Problems in Artificial Life

Mark A. Bedau^{*,†}
John S. McCaskill[‡]
Norman H. Packard[§]
Steen Rasmussen^{**}
Chris Adami^{††}
David G. Green^{‡‡}
Takashi Ikegami^{§§}
Kunihiko Kaneko^{***}
Thomas S. Ray^{†††}

Abstract This article lists fourteen open problems in artificial life, each of which is a grand challenge requiring a major advance on a fundamental issue for its solution. Each problem is briefly explained, and, where deemed helpful, some promising paths to its solution are indicated.

Keywords

artificial life, evolution, self-organization, emergence, self-replication, autopoiesis, digital life, artificial chemistry, selection, evolutionary learning, ecosystems, biosystems, astrobiology, evolvable hardware, dynamical hierarchies, origin of life, simulation, cultural evolution, information theory

1 Introduction

At the dawn of the last century, Hilbert proposed a set of open mathematical problems. They proved to be an extraordinarily effective guideline for mathematical research in the following century. Based on a substantial body of existing mathematical theory, the challenges were both precisely formulated and positioned so that a significant body of missing theory needed to be developed to achieve their solution, thereby enriching mathematics as a whole.

In contrast with mathematics, artificial life is quite young and essentially interdisciplinary. The phrase “artificial life” was coined by C. Langton [11], who envisaged an investigation of life as it is in the context of life as it could be. Although artificial life is fundamentally directed towards both the origins of biology and its future, the scope and complexity of its subject require interdisciplinary cooperation and collaboration. This broadly based area of study embraces the possibility of discovering lifelike behavior in unfamiliar settings and creating new and unfamiliar forms of life, and its major aim is to develop a coherent theory of life in all its manifestations, rather than an historically contingent documentation bifurcated by discipline. Artificial life’s practical implications are no less sweeping than its scientific implications. It can guide how we use new technologies for extending life and creating new forms of it, including drugs, prosthetics, the Internet, evolvable hardware, and proliferating robots. It has a vast potential for both good and harm, with a serious need for public debate and ethical clarity.

* The first four authors shared the primary responsibility for writing this document.

† Reed College, 3203 SE Woodstock Blvd., Portland, OR 97202, USA, mab@reed.edu

‡ German National Research Center for Information Technology, Schloss, Birlinghoven, D-53754 Germany, mccaskill@gmd.de

§ Prediction Company, 236 Montezuma St., Santa Fe, NM 87501, USA, n@predict.com

** Los Alamos National Laboratory, MS-D450, Los Alamos, NM 87545, USA, steen@lanl.gov

†† California Institute of Technology, Pasadena, CA 91125, adami@krl.caltech.edu

‡‡ School of Environmental and Information Science, Charles Stuart University, Albury NSW 2640 Australia, dgreen@life.csu.edu.au

§§ Institute of Physics, The Graduate School of Arts and Sciences, University of Tokyo, 3-8-1 Komaba, Meguro-ku, Tokyo 153 Japan, ikeg@sacral.c.u-tokyo.ac.jp

*** Department of Pure and Applied Sciences, University of Tokyo, Komaba, Meguro-ku, Tokyo 153 Japan, kaneko@complex.c.u-tokyo.ac.jp

††† Department of Zoology, University of Oklahoma, Norman, OK 73019, USA, tray@ou.edu

Artificial life's youth and mix of disciplines, as well as its practical consequences, only increase the value of publishing a list of its key open problems. Authors at the Seventh International Conference on Artificial Life were explicitly invited to relate their research to the past and future of artificial life (see [1]), and the conference was preceded by widespread interest in formulating a set of grand challenges in artificial life. As a second generation of scientists commences work in this field, many agreed that a publicly recorded list of grand challenges could serve both to better define the core character and objectives of this interdisciplinary activity and to provide starting points for useful new contributions.

The conference organizers took on the task of achieving a coherent and precisely documented list of such challenges. The conference concluded with a round table discussion at which panelists (including all the authors of this document) presented challenges chosen from a larger list that they had collectively suggested. Further contributions followed from the floor, with participation from a large proportion of the conference delegates. Subsequently, all of this information was restructured and significantly extended by the conference organizers (in consultation with the other panelists), and that work culminated in the list presented here. While not completely democratic, we hope our list presents a sensible compromise between representing a diversity of views and constituting a coherent and structured list of challenges.

Artificial life is foremost a scientific rather than an engineering endeavor. Given how ignorant we still are about the emergence and evolution of living systems, artificial life should emphasize understanding first and applications second, so the challenges we list below focus on the former. Although no monetary incentive is provided, the authors believe that the community will vigorously contest the challenges and that their solutions will bring rich rewards in terms of both recognition for those who solve them and scientific accomplishment across the fields within artificial life. As with Hilbert's problems, formulating a clear and fruitful challenge provides direction for, rather than replacing, shorter-term (e.g. five year) research goals, so the challenges reported here are unabashedly long term. However, the potential exists for some significant progress on at least the first ten in the next decade.

2 A List of Open Problems

The following is a structured list of key open problems in artificial life. Some goals involve practical synthesis; others involve theoretical analysis. The aim is not necessarily to make each problem wholly independent of the others; some problems overlap or complement others, which may help to triangulate the deeper issues.

The challenges are classified under three broad issues: the transition to life, the evolutionary potential of life, and the relation between life and mind and culture. The challenges falling under the third issue are necessarily more speculative and open-ended, so this whole list may best be viewed as ten challenges plus four areas of investigation. Moreover, some of the questions about mind and culture interweave scientific and nonscientific issues. Those issues are still important, though, not least because addressing them is probably the best way to clarify what in this area can be known scientifically.

A. How does life arise from the nonliving?

1. Generate a molecular proto-organism in vitro.
2. Achieve the transition to life in an artificial chemistry in silico.
3. Determine whether fundamentally novel living organizations can exist.
4. Simulate a unicellular organism over its entire lifecycle.

5. Explain how rules and symbols are generated from physical dynamics in living systems.
- B. What are the potentials and limits of living systems?
6. Determine what is inevitable in the open-ended evolution of life.
 7. Determine minimal conditions for evolutionary transitions from specific to generic response systems.
 8. Create a formal framework for synthesizing dynamical hierarchies at all scales.
 9. Determine the predictability of evolutionary consequences of manipulating organisms and ecosystems.
 10. Develop a theory of information processing, information flow, and information generation for evolving systems.
- C. How is life related to mind, machines, and culture?
11. Demonstrate the emergence of intelligence and mind in an artificial living system.
 12. Evaluate the influence of machines on the next major evolutionary transition of life.
 13. Provide a quantitative model of the interplay between cultural and biological evolution.
 14. Establish ethical principles for artificial life.

This list of challenges could be extended, of course. Some fundamental questions are missing, the most notable being the nature of life itself. This question is presupposed by a number of the problems listed, however, and answering any problem will necessarily involve resolving anything it presupposes. Stating each challenge precisely enough that progress toward its solution can be unambiguously measured is often beyond our grasp. However, where ambiguity remains, the discussion and evaluation of potential solutions can significantly clarify the challenge and contribute toward its eventual solution, in part by identifying and solving precise subproblems of those challenges. The authors are also aware that this list under-represents some areas in which artificial life plays a significant role, such as robotics, games, and art. In addition to promoting work on the issues listed here, the authors seek to prompt others to identify and publicize additional fundamental challenges. Discussion about whether to add new challenges to the list will itself facilitate progress throughout the field and we purposefully avoid the bias that would be introduced by citing what we think is the most important work on each topic.

3 Explanation of the Challenges

The aim of this section is to describe the challenges as precisely as possible and to provide one or two starting points in existing research for those seeking their solution. Existing research is mentioned only to clarify the context of a challenge; an exhaustive review of prior work is well beyond the scope of this paper.

3.1 Generate a Molecular Proto-organism in Vitro

One of artificial life key goals is constructing a life form in the laboratory from scratch, and the first targets in this quest should be the simplest possible forms of life. Molecular life is understood as a self-reproducing organizational form constructing itself in a simple environment and capable of evolution. The environment may be defined either by natural geochemistry or artificially, but should involve only simple forms of energy and

material, so that the goal is to create an encapsulated biochemical system that can derive energy from simple chemicals or light (a simple form of metabolism, e.g. by redox coupling) and use information carried in primitive genes. This proto-organism, which may be held together by a lipid aggregate, should be able to self-replicate, use energy and nutrients from its environment, undergo evolutionary changes over time, and die.

Artificial life research has succeeded in creating both self-replicating (templating) molecules that show limited evolution, and self-replicating compartments (micelles, lipid aggregates). However, combining these in an autonomous, evolvable, self-replicating autopoietic system remains a challenge. Complementing this bottom-up approach, a top-down approach based on the success of genomic research may point the way to a solution while not constituting one. Genomics researchers are whittling down the simplest contemporary cells (e.g., mycoplasmas) to reduce the genetic complexity to a point where synthesis might be possible.

3.2 Demonstrate an Artificial Chemistry in Which the Transition to Life Occurs in Silico

Few questions concerning living systems are as fundamental as the spontaneous generation of life, and a constructive context in which to explore this issue is artificial chemistry. Artificial chemistries are computer-based model systems composed of objects (abstractions of molecules), which are generated by collision between existing objects according to a predefined interaction law. The collisions may be purely random or allow the buildup of spatial correlations, as in diffusion systems. Bimolecular chemistry is assumed to be sufficient to display the transition to life, but this may involve complex structures. The chemistry may be stochastic rather than deterministic, but should be constructive rather than descriptive; that is, an interaction law should predict (like an algorithm) the product molecules for colliding objects of arbitrary complexity. Furthermore, the chemical interaction law should be simple compared with the collective solutions to the problem. Further restriction on reversible chemistries and those satisfying conservation laws (like mass and energy) might make the problem harder but more interesting. Although judging the success of solutions to this challenge requires a clear definition of life, the debate about success in this problem should contribute to our understanding of life. As a starting point, we may use the definition outlined in the first challenge.

Artificial chemistries have been investigated by many authors in spaces of various dimensionalities, with deterministic and probabilistic interaction laws. Molecules have been abstracted using cellular automata, secondary structure folding algorithms, finite state automata, Turing machines, von Neumann machines, and the lambda calculus. Note that the challenge requires a self-organized collection of separate artificial molecules to stably reconstruct, replicate, and evolve (as in the physical case of challenge 1). The spontaneous generation of self-replicating and evolving programs in computer memory seeded with random machine language instructions would fall short of this goal unless the internal organization of these “organisms” and the boundaries separating them from their environment arise and are sustained through the activities of lower-level primitives.

3.3 Determine Whether Fundamentally Novel Living Organizations Can Exist

Artificial life calls special attention to the question of whether there can be forms of life that are fundamentally different from the familiar biological instances. For example, life as we know it encodes information needed to re-create structures that cannot be copied directly (such as proteins) in combinatorial structures that can be copied generically (genes). Such genetic organization may be based on information encoded in various physical degrees of freedom (not necessarily just molecular). This challenge involves

either proving that such organization is necessary for life or producing a fundamentally different living organization either in wet carbon chemistry or in a different medium. As above, the debate about what is “fundamentally different” will help to clarify our definition and understanding of life. Two categories of solution involve novel genetic organizations and non-genetic organizations. Achieving this challenge will be a big step toward discerning the boundaries of life as it could be.

The central dogma of molecular biology and an essentially universal genetic code are examples of strongly conserved organizational themes of living organisms. As an example, Eigen [7] gave a strong argument that autocatalytic sets without genetic underpinning have a problem in principle with evolvability. No proof of this conclusion has been given, but examples to date of autocatalytic sets without replicable elements have failed to exhibit convincing evolution. This may serve as one starting point for those interested in this challenge. Other examples of “fundamentally different” organizations include those without spatially localized individuals, hierarchical organization, genotype-phenotype distinction, or, indeed, any symbolic representation scheme. Macroscopic self-reproducing machines (robots), perhaps living in colonies and exploiting local resources to maintain themselves and grow, might give rise to a significantly different set of organizational possibilities and constraints. The question is also relevant to the recently renewed quest for life in the universe (e.g. at the NASA Astrobiology Institute), since different organizational schemes may further expand the range of possible habitats for life. Additional suggestions for addressing this challenge might be found at the borderline of what we understand as life on our planet, such as evolving populations of clay crystallites [4] and (although this is more controversial at present) the so-called “nanobes” (nano-scale organisms) recently found in Triassic and Jurassic sandstones [18].

3.4 Simulate a Unicellular Organism over Its Entire Life Cycle

In contrast with the first two challenges, this one involves simulating an existing biological organism on the computer, and the goal is for all of the organism’s essential dynamics to emerge from a bottom-up simulation. The artificial organism should exhibit virtually its complete spectrum of behavior, including its ability to evolve. Of course, this behavior will be somewhat unrealistic, because our ability to model complex physical systems is only approximate. Nevertheless, since biological organization is statistical, it is conceivable that the dynamics of a cell can be integrated to produce identical macroscopic behavior under a wide range of stimuli and environmental conditions.

The existence of a complete genome for simple organisms (such as mycoplasmas) definitely takes us one step towards meeting this challenge. The integration of the simulation of many thousands of proteins, and genetic as well as regulatory networks, at the level of deterministic kinetics would already provide important novel quantitative understanding of cell cycle dynamics. However, for moderate completeness, simulating the folding of all biopolymers and their reactions and supramolecular interactions is still a formidable challenge, since current successes in folding are statistical rather than *ab initio*, and vast progress in integrating molecular dynamics on time scales of minutes to hours is needed. Even compared with large current simulations, we would need a factor of at least 10^3 in volume and 10^{12} in time to attain a bottom-up simulation on the time scale of replication. Now $10^{15} \approx 2^{50}$, so that if computer performance continues to improve as in recent years, the simple molecular dynamics version of this goal should be achievable this century. Better algorithms and understanding may well accelerate progress, but those tackling the challenge may also have to deal with the prediction of fluctuation-dependent chemical reactivity, which may require still more computational effort. On the other hand, combinations of (for example) reaction kinetics, molecular dynamics simulations, and lattice gas simulations would be more powerful than any

single simulation approach. Whether the number of configurational variants that need to be sampled in the course of such a simulation will be feasible for capturing essential evolutionary behavior is an open question involving both technical progress (in hardware and software) and scientific progress (in understanding the essential features of macromolecular dynamics for cells).

3.5 Explain How Rules and Symbols are Generated from Physical Dynamics in Living Systems

In contrast with popular theories of evolution, in which evolution is defined with respect to given symbolic genes and rules such as for replication or recombination, a proper theory of evolving life must allow evolution to create new symbolic descriptions and rules to manipulate them. In general, these entities are formed by an underlying physical dynamics, and then evolution sharpens and reinforces this symbolic level. How can such systems be modeled effectively where no state transition rules at the new symbolic level are specified in advance, in contrast with current modeling practice. One may try to mitigate this drawback by successively imposing additional constraints, but the challenge is to allow these dynamical constraints to emerge from the underlying dynamical system. This appears essential for modeling spontaneity or autonomy in life. In fact, there are two issues here: how physics can give rise to dynamical systems that can operate on the symbolic level, and how and why physical systems that are alive tend to expand and refine their symbolic platform. The first part of the challenge is to provide a theory explaining how dynamical systems can generate phenomena best understood as novel symbolic and rule-based behavior, that is, creating novel properties independent of detailed fluctuations in configuration. The second part is to specify under what conditions a natural discrete classification of dynamical states is recognized and reinforced by a dynamical system itself to structure its future evolution.

About 30 years ago, Howard Pattee [15, 16] reduced higher rules to non-holonomic constraints in physical systems, but the challenge here is to account for the evolution of these dynamical constraints. Starting with rules at a low level of description, one needs to study the formation of a rule at a higher level. Symbolic states, on which the rules act, must be formed as relatively stable states, like attractors in the underlying dynamical system, but also have inner degrees of freedom allowing them to undergo a transition only when certain other (symbolically characterizable) dynamical conditions are met. A rule that drives other states must be formed as an approximate structure in the system and it appears appropriate to consider the formation of rules and symbols together from a continuous description. The standard dynamical systems approach, albeit seemingly more suitable for this challenge than an automaton approach, may not suffice to fully resolve the issue of rule creation. Mathematicians recognize the difference between (a) dynamical relations between objects and (b) structure that reflects rules for this dynamics (as in modern algebra, category theory for instance), but a novel mathematical formulation will be necessary to show how one gives rise to the other. One promising direction may be provided by function dynamics, where the dynamics of functions rather than states is studied.

Genetic models of evolution start with a discrete state space in the form of heteropolymers of defined sequence, but in which the transmission of information (reading the symbols) is probabilistic. The quasispecies theory of Eigen [7] already demonstrates how information can then be generated from physical dynamics, selection being achievable through differences in rate coefficients for the production and destruction of such polymers. Chemical kinetics also provides a convincing theory to explain the relatively discrete logic of gene expression for example based on allosteric effects, following the tradition of Monod [14]. The new field of DNA computing has already made it clear that symbol-based rules grounded in the physics of DNA can be advanced to complete uni-

versal computation. However, there is currently no clear indication for the conditions under which a discrete symbolic formulation is conducive to inheritance and survival in an evolving system, and a successful response to this challenge would provide an explanation for the apparent tendency towards symbolic knowledge in biology.

3.6 Determine What Is Inevitable in the Open-Ended Evolution of Life

This challenge concerns the contingency of life. In different historical unfoldings of the evolutionary process and in evolution in other media, two related questions arise: (a) What are the features common to all evolutionary processes, or to broad classes of evolutionary processes? (b) Do different evolutionary processes contain fundamentally different evolutionary potential? Resolution of the first question will be analogous to results in dynamical systems theory on genericity and structural stability, and to results in statistical mechanics regarding universality, that is, independence of macroscopic phenomena with respect to microscopic detail.

If the organization of living organisms is their major distinguishing feature within an evolutionary process, and the selection of living organisms is primarily dictated by their interactions with other organisms, then there is some reason to expect that the major features of biology are independent of the underlying physical media supporting life. Digital processing in computers appears to be a very different medium from molecular biology, yet digital organisms based on genetic and cellular principles have been constructed in controlled operating system environments. In digital media one has considerable scope to vary the type of “physics” underlying the evolutionary process, and so it is straightforward to investigate the question of environmental and rule contingency. A key challenge is whether digital systems based on symbolic logic harbor the same potential for evolutionary innovation as physical systems. A preliminary challenge is to unlock the full potential of evolution in digital media. Many believe that digital life today falls far short in this regard, and this issue is starting to be approached quantitatively. Finally, the question of “what would remain the same if the tape of life were replayed” [9] in its narrowest form refers to an identical geophysico-chemical environment and concerns the reproducibility of the evolutionary process. Clearly the time series of many of the innovative discoveries of living systems are subject to extreme variation, but would the character of the ultimate organisms and ecologies be different in a second evolution of life?

The question of chemical contingency can be answered for local mechanisms by synthetic construction, as has been begun for alternative genetic molecules by varying the sugar moiety (e.g. hexose nucleic acids [8]). Physical contingency is being approached by using very different building blocks, such as plastic disks at interfaces, and evolvable hardware of varying kinds—that is, by shifting to a different physical domain rather than changing any physical laws. Lenski’s work with replicate lines of *E. coli* that are forced to adapt to a change in carbon source illustrate how historical contingency in biological evolution can be addressed experimentally today [17]. Historical contingency in a broader context or over a longer time scale must be answered by simulation, which of course is more tractable for digital life. A significant answer to this question appears achievable by combining these approaches but requires a major effort.

3.7 Determine Minimal Conditions for Evolutionary Transitions from Specific to Generic Response Systems

This challenge concerns the interplay between evolution and learning prior to mind (i.e., between hereditary and acquired information) in the evolutionary formation of response systems. In particular, this challenge distinguishes between specific, special-purpose response systems that embody a narrow range of organismic functionality, and generic response systems that embody a broad range of organismic functionality.

The challenge is to understand how relatively general systems within an individual for sensing and responding to environmental change can arise in the course of evolution.

One example is the development of motile response of an organism to environmental stimuli. A particular example is irritability, which has long been considered a defining property of life. Even single-celled modern organisms show a remarkable ability to sense and respond to their environment. The presence of predators is an obvious example of a major short-term fluctuation in the environment affecting survival that motivates a motile response; another is competition for resources, such as light for a growing leaf affected by the growth of other leaves or plants. The variable geometric relationships in such interactions require a generic response such as directional avoidance. Another example is the development of mechanisms for defense against molecular invasion. For complex multicellular organisms, antibody variability within an immune system provides a near-universal encoding scheme for responding to molecular invasion. Simpler organisms have restriction enzymes that recognize foreign DNA, and metabolic activity is modulated by a system for regulating gene expression in response to resource variations. In this sense, the proteins coded by DNA also provide fairly complete coverage (extended by small cofactors) of the chemical world of catalytic reactions. A final example is the evolution of high-order nervous systems and brains. At the level of cultural evolution, the eye and brain have attained an extraordinarily generic representation scheme enabling the copying of the behavior of other organisms.

One starting point for a solution is to investigate the evolution of artificial life models with short-term fluctuations in survival. There would appear to be a minimum complexity of environmental variation to warrant a generic response system. If other organisms are the major source of such variation, how does the evolution of such generic response mechanisms interact with ecology and population diversity? An indication that such effects are important in evolution theory is given by recent attention to the Baldwin effect [2], highlighting one subtle relationship between acquired characteristics and evolutionary optimization. Although such questions are central to evolutionary biology, they are difficult to address theoretically in biology without resorting to the artificial life context, where greater flexibility in experimental design and increased measurement precision are readily available. Phrased in terms of information flow, this challenge also has bearing on challenge 10 below. Furthermore, it provides a precursor at the molecular level to the more complex issues of challenge 11.

3.8 Create a Formal Framework for Synthesizing Dynamical Hierarchies at All Scales

In biological systems hierarchies with multiple functionalities at different scales can be found everywhere. Clearly we have the fundamental coarse-grained hierarchy that may be expressed as: ecosystems (including social systems), organisms, organs, tissues, cells, organelles, molecules, atoms, which have many refinements and substructures. Obviously, the properties associated with each level are generated by the collective dynamics of the elements in these dynamical hierarchies. One problem here is to create a formal framework for consistently *describing* such hierarchical systems, with some coarse-graining procedure for moving between levels.

A second issue is how these complex and robust functionalities are *generated* in biological (and protobiological) systems. We know that novel functionalities in molecular systems can arise in at least two ways: by assembly and by evolution. Molecular self-assembly processes are probably most central as a mechanism for bridging non-living and living matter, for the transition from prokaryotic to eukaryotic organisms, and perhaps for other major biological transition processes. Self-assembly and self-organization allow limited inherited information to code for complex functionalities, from enzyme catalysts to the brain. It is probably necessary to embrace assembly

and evolution as two complementary and perhaps equally important aspects of how to generate bio-complexity if one is to account for the generation and prevalence of dynamical hierarchies in living systems.

Finally, it is unclear how to *simulate* either the generation or the functioning of such multi-level systems, because their dynamics are characterized by multiple length and time scales and both computational power and a general theory are missing. This is also one of the barriers faced in solving challenge 4.

3.9 Determine the Predictability of Evolutionary Manipulations of Organisms and Ecosystems

One of the most pressing practical issues facing humankind is how our activities affect our environment. But before we can evaluate these effects, we need to understand the extent to which they can be predicted. This is part of a larger issue involving the predictability of the effects of manipulating ecosystems of all sizes and kinds. The ecosystems of interest include those as different as the entire global biosphere and individual human immune systems, and ecological manipulations range from industrial pollution, climate change, and large-scale mono-crop agriculture to the introduction of genetically engineered organisms. Although such manipulations have effects on both relatively short and relatively long time scales, the present challenge focuses specifically on evolutionary consequences.

Organisms can be externally directed by manipulations of their environment to perform external tasks (e.g. from catalyzing specific chemical reactions to running artificial mazes) and transgenically manipulated to function differently by expressing different genes, but the evolutionary consequences and limits of such manipulations are not clear. How far can one rationally redesign or rapidly select organisms to fulfill multiple novel criteria without disturbing the viability of the organisms' organization and defense systems? Is there a tradeoff between utility and viability, or between size of modification and duration of organism utilization? These questions become even more of a challenge when it comes to manipulating ecosystems. The importance of assessing the potential costs and benefits of these manipulations (see the final challenge, Section 3.14) is underscored by the fact that these manipulations are already becoming common in restricted contexts.

With increasing understanding of the genetic control of development, it will be possible to create novel multicellular organisms through sequential genetic reprogramming. Do we need long-term evolutionary optimization to support or perfect such major changes to organisms? Widespread application of such manipulations could also rapidly increase biodiversity without creating the variation in the gene pool necessary for longer-term evolution. Speciation and extinction are ubiquitous features of evolutionary history, but the long-term implications of frequent artificial speciation for biodiversity and sustainability are unknown. Addressing this challenge requires combining an understanding of evolution theory with theoretical experimentation in artificial life models that constructively address hypothetical changes to organisms and ecosystems. Progress on this challenge should have an impact on guidelines for intervening in biological and technological contexts with evolutionary implications.

3.10 Develop a Theory of Information Processing, Information Flow, and Information Generation for Evolving Systems

Information plays three distinct roles in evolving systems: transmission, processing, and generation. The challenge refers to all three, and they are listed here in order of the degree to which we presently understand them, and hence inversely with respect to the challenge they represent.

Firstly, there appear to be two complementary kinds of information *transmission* in living systems. One is the conservative hereditary transmission of information through evolutionary time. The other is transmission of information specified in a system's physical environment to components of the system, possibly mediated by the components themselves, with the concomitant possibility of a combination of information processing and transmission. The latter is clearly also linked with the generation of information (to be discussed last). Clarifying the range of possibilities for information transmission, and determining which of those possibilities the biosphere exploits, is a fundamental enquiry of artificial life.

Secondly, the challenge is to unify evolution with information *processing*. One starting point is the observation that components of evolving systems (organisms or groups of organisms) seem to solve problems as part of their existence. It is this capability that led to the creation of the genetic algorithm [10] and its many applications. More generally, theory must address what the capacity of an evolving system's information processing is, and how it changes with evolution. Are there thresholds between levels of information processing during evolution that match the levels identified in automata theory—for example, from finite state machines to universal computation? How do the algorithms employed by organisms classify in terms of their problem solving efficiency?

The third and least-understood role of information is its *generation* during evolution. As evolution takes place, evolving systems seem to become more complex; successfully quantifying complexity and its increase during evolution is one important part of understanding information generation. Another problem in this area is that of understanding how complexity in an evolving system's environment can affect the complexity of the organisms that are evolving within the environment. Nature seems not to have evolved Lamarckian mechanisms to incorporate environmental information from one generation's experience directly into the genome of the next generation, but the characteristics of an environment must inevitably be reflected indirectly in the organisms that evolve within it (a transmission of information from environment to genome). Finally, the environment itself is made up of a physical, nonliving component plus all the living components themselves. The theory envisioned must address the resultant compounding of the informational context continually being built and expanded by the evolutionary process.

3.11 Demonstrate the Emergence of Intelligence and Mind in an Artificial Living System

It is natural to wonder about the relationship between life and mind. They appear coextensive in nature, and the distribution of different kinds of mental capacities (sensitivity to the environment, internal information processing, etc.) across life forms seems roughly to reflect the sophistication of those life forms. Two deep issues in this area arise for artificial life. The first is substantive: whether and, if so, how the natures of life and mind are intrinsically connected. The second is methodological: whether it is most profitable to study mind and intelligence only when embodied in living systems. Both issues motivate artificial life's existing attention to autonomous agents and embodied cognition, and they bear on artificial life's relation to its elder sister, artificial intelligence. Progress on this challenge will shed new light on many current controversies in both fields, such as the extent to which life and mind should be viewed as "computational." A constructive approach to all these concerns is to try to demonstrate the emergence of intelligence and mind in an artificial living system.

The history of artificial intelligence shows how controversial it is to settle whether a model demonstrates intelligence or mind, and addressing this controversy from within the self-referential context provided by evolving systems is one point of this challenge. To make progress, one must have a method to detect intelligence and mind when they

are present in a system. Consciousness is the most difficult aspect of mind to detect, and initial progress is certain to be somewhere else. A more tractable aspect of mind to detect is meaning, that is, internal states that have semantic or representational significance for the entity and that influence the entity's behavior by means of their semantic content. One guide for progress in this vein is artificial life's extensive current research on the evolution of communication. The easiest aspect of mind and intelligence to detect is flexible adaptive behavior, i.e., the capacity to act appropriately in a complex dynamic environment. Maturana and Varela [12], among others, view this capacity as the hallmark of both cognition and life. This capacity obviously comes in degrees, and most existing artificial life work illustrates it to some extent. Further progress involves, among other things, clearly delineating the different degrees or stages of this capacity, and having increasingly sophisticated forms of this capacity emerge from increasingly impoverished initial conditions.

3.12 Evaluate the Influence of Machines on the Next Major Evolutionary Transition of Life

Once life originated, biological evolution underwent a number of major evolutionary transitions, and presumably there are more to come. Even though the identification of "major" transitions involves some inescapable arbitrariness, probably everyone will agree that past transitions include the origin of eukaryotes, the origin of multicellular life, and the origin of language and culture. Once culture originates, it has the capacity to evolve on its own, and the past century has shown an explosion of technological culture, including notably the creation of computing machines and complex distributed informational networks connecting them. Many agree it is only a matter of time before artificial life creates machines that are alive, are intelligent, reproduce their own kind, have their own purposes, set their own goals, and evolve autonomously. These machines will be as much a part of the natural world as features in the landscape or existing forms of life, and their evolution will affect the course of existing forms of life. Think of how machines are currently ramifying the nature and rate of human communication and interconnectedness. All this suggests that machines might play an unprecedented role in the next major evolutionary transition, and the challenge here is to predict and explain this role. Machines may well be the central players in the transition, as will be the case if autonomously evolving machines get established in the natural world. Machines might also play a supporting role in the transition by providing an infrastructure that influences its rate or direction of change.

One place where one might seek clues for how to address this challenge is current theories of major evolutionary transitions [13], though there is little consensus on this subject. A second source of clues would be artificial life models that exhibit a "major" evolutionary transition, though there is no consensus whether any such model yet exists. A third source is to study humankind's increasingly symbiotic relationship with increasingly autonomous and intelligent machines. Today industrial robots are used to produce industrial robots and Internet search engines adapt and take initiative. Relevant here also are studies of self-programmable and evolvable machines and studies of the robustness of informational networks. Theories of technological evolution (see the next challenge) are a fourth possible source of clues.

3.13 Provide a Quantitative Understanding of the Interplay between Cultural and Biological Evolution

Culture, in all its forms, is one of the products of human existence, and culture itself evolves. Examples of cultural evolution can be found throughout the social sciences, including the development of economic markets, the changes in technological infras-

tructure (see the previous challenge, Section 3.12), and growth and revolution in the community of scientific beliefs. Some treatments of cultural evolution (e.g., sociobiology and evolutionary psychology) consider how cultural traits evolve due to their effect 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” [6]. This sort of “pure” cultural evolution is driven by mechanisms similar to those behind biological evolution, but there are important differences. Evolution happens in each case because 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. Psychological transmission, like genetic transmission, appears to depend on the ability to reconstruct complex behavior from the copyable structures of communication. One fundamental question addressed by a quantitative model of the interplay between cultural and biological evolution concerns the similarities and differences in the behavior of biological and cultural evolution. For example, do the two kinds of evolutionary processes exhibit similar generic trends (perhaps punctuated equilibria and increasing diversification), and if so, for similar reasons? Another fundamental question concerns how biological and cultural evolution are interconnected and influence each other.

Psychological information transmission can involve speaking, writing, or non-linguistic forms of behavior. Explaining psychological information transmission in terms of standard information transmission is one preliminary step toward meeting this challenge, and subsequently relating psychological and genetic (evolutionary) information transmission (recall challenge 10) is one strategy for doing so. Existing quantitative theories of cultural evolution [3, 5] are another starting point for pursuing this topic. These theories adapt analytic population biology to the context of cultural evolution, and they include treatment of how cultural evolution and biological evolution interact. However this approach has the characteristic limitations of analytical population biology, the very limitations that drive the pursuit of synthetic bottom-up models in artificial life.

3.14 Establish Ethical Principles for Artificial Life

Both the process of pursuing artificial life research and the scientific and practical products of that research process raise complicated ethical issues. These issues fall into four broad categories: (a) the sanctity of the biosphere, (b) the sanctity of human life, (c) the responsible treatment of newly generated life forms, and (d) the risks of exploitation of artificial life.

- a. The autonomous proliferation and evolution of engineered biological systems, whether natural or artificial, could rapidly and irreversibly change the biosphere, and the impact of such changes is difficult to assess. This raises the question whether there is an ethical right to introduce such systems to solve immediate problems. Similar concerns also affect proliferation of artificial organisms on increasingly indispensable artificial systems such as the Internet. Existing computer viruses wreak havoc as it is, but imagine how much worse they would be if they spontaneously evolved like artificial life systems. On the positive side, artificial life is helping to develop new measures to protect our informational networks from such threats, including those modeled on the immune systems found in nature. Furthermore, the investigation of ecosystem changes in artificial environments may help to establish guidelines for current and future interference with the biosphere, be they by chemicals (e.g. CO₂, drugs), genetic engineering, or artificial organisms.

- b. Much of current ethics is based on the sanctity of human life. Research in artificial life will affect our understanding of life and death and the relationship between life and mind. This, like the theory of evolution, will have major social consequences for human cultural practices such as religion. Artificial life can also provide positive insight, through long-term simulations, into necessary regulations and freedoms in dealing with our evolving species. The issue of our future increasing dependence on artificial life systems also falls under this category.
- c. It is worth noting that public protocols govern the responsible treatment of human and animal research subjects. The lack of analogous protocols in artificial life may be no serious problem today, but as we create more sophisticated living entities we will have to face the responsibility of treating them appropriately.
- d. The practical uses of the products of artificial life research will create various costs and benefits. For example, it is easy to imagine the military and commercial uses of automatic adaptive control processes. Such uses of artificial life research may create conflicts of interest for artificial life researchers, as may responsibilities toward sources that fund artificial life research.

Artificial life's ethical issues somewhat resemble those concerning animal experimentation, genetic engineering, and artificial intelligence. The extensive literature on the ethical issues raised in those three fields may provide some guidance for exploring 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.

Acknowledgments

Thanks to the audience at Artificial Life VII for helpful suggestions and critical comments during the round table discussion out of which this paper grew. Special thanks to Inman Harvey for his participation in the conference round table and his critical and constructive comments on this paper.

References

1. Bedau, M. A., McCaskill, J., Packard, N., & Rasmussen, S. (Eds.) (2000). *Artificial life VII: Proceedings of the Seventh International Conference on the Simulation and Synthesis of Living Systems*. Cambridge, MA: MIT Press.
2. Belew, R., & Mitchell, M. (Eds.) (1996). *Adaptive individuals in evolving populations: Models and algorithms*. Redwood City, CA: Addison-Wesley.
3. Boyd, R., & Richardson, P. (1985). *Culture and the evolutionary process*. Chicago: University of Chicago Press.
4. Cairns-Smith, A. G. (1982). *Genetic takeover and the mineral origins of life*. Cambridge, UK: Cambridge University Press.
5. Cavalli-Sforza, L., & Feldman, M. (1981). *Cultural transmission and evolution: A quantitative approach*. Princeton: Princeton University Press.
6. Dawkins, R. (1976). *The selfish gene*. Oxford: Oxford University Press.
7. Eigen, M. (1971). Self-organization of matter and evolution of biological macromolecules. *Naturwissenschaften*, 58, 465–523.
8. Eschenmoser, A. (1994). Chemistry of potentially prebiological natural products. *Origins of Life and Evolution of the Biosphere*, 24, 389–423.
9. Gould, S. J. (1989). *Wonderful life: The Burgess Shale and the nature of history*. New York: Norton.

10. Holland, J. H. (1975). *Adaptation in natural and artificial systems: An introductory analysis with applications to biology, control, and artificial intelligence*. Ann Arbor: University of Michigan Press. Expanded edition, 1992. Cambridge, MA: MIT Press.
11. Langton, C. G. (1986). Studying artificial life with cellular automata. *Physica D*, 22, 120–149.
12. Maturana, H., & Varela, F. (1980). *Autopoiesis and Cognition*. Dordrecht: Reidel.
13. Maynard Smith, J., & Szathmary, E. (1995). *The major transitions in evolution*. New York: Freeman.
14. Monod, J. (1970). *Le hasard et la nécessité*. Paris: Éditions du Seuil. English translation (1971): *Chance and necessity*. New York: Knopf.
15. Pattee, H. H. (1967). Quantum mechanics, heredity and the origin of life. *Journal of Theoretical Biology*, 17, 410–420.
16. Pattee, H. H. (1996). Physical theories, automata, and the origin of life. In H. H. Pattee et al. (Eds.), *Natural automata and useful simulations*. Washington: Spartan Books.
17. Travisano, M., Mongold, J. A., Bennet, A. F., & Lenski, R. E. (1995). Experimental tests of the roles of adaptation, chance, and history in evolution. *Science*, 267, 87–90.
18. Uwins, P. J. R., Webb, R. I., & Taylor, A. P. (1998). Novel nano-organisms from Australian sandstones. *American Mineralogist*, 83, 1541–1550.