

Artificial life: organization, adaptation and complexity from the bottom up

Mark A. Bedau

Department of Philosophy, Reed College, 3023 SE Woodstock Blvd., Portland OR 97202, USA

Artificial life attempts to understand the essential general properties of living systems by synthesizing life-like behavior in software, hardware and biochemicals. As many of the essential abstract properties of living systems (e.g. autonomous adaptive and intelligent behavior) are also studied by cognitive science, artificial life and cognitive science have an essential overlap. This review highlights the state of the art in artificial life with respect to dynamical hierarchies, molecular self-organization, evolutionary robotics, the evolution of complexity and language, and other practical applications. It also speculates about future connections between artificial life and cognitive science.

Artificial life (also known as ‘ALife’) is an interdisciplinary study of life and life-like processes that uses a synthetic methodology. Three broad and intertwining branches of artificial life correspond to three different synthetic methods. ‘Soft’ artificial life creates simulations or other purely digital constructions that exhibit life-like behavior, ‘hard’ artificial life produces hardware implementations of life-like systems, and ‘wet’ artificial life synthesizes living systems out of biochemical substances. Cognitive science and artificial life share some intellectual roots, and their subjects and methodologies are related. Now that artificial life has matured over the past decade or so, it is appropriate to review its achievements and speculate about its future connections with cognitive science.

The origins and methodology of artificial life

The phrase ‘artificial life’ was coined by Christopher Langton, who envisioned a study of life as it could be in any possible setting [2]. Artificial life owes its two deepest intellectual roots to John von Neumann and Norbert Wiener. Von Neumann [3] designed the first artificial-life model (without referring to it as such) when he created his famous self-reproducing, computation-universal cellular automata (see Figure 1). He tried to understand the 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 [4] started applying information theory and the analysis of self-regulatory processes (homeostasis) to the study of living systems. The constructive and abstract methodology of cellular automata still typifies much of artificial life, as

does the abstract and material-independent methodology of information theory.

Artificial life has also been influenced by developments in traditional disciplines. In addition to a wealth of information about the life forms found on Earth, some models originally devised for specific biological phenomenon have been adopted by artificial life. 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. In addition, the use of cellular automata as exemplars of complex systems [5] led directly to

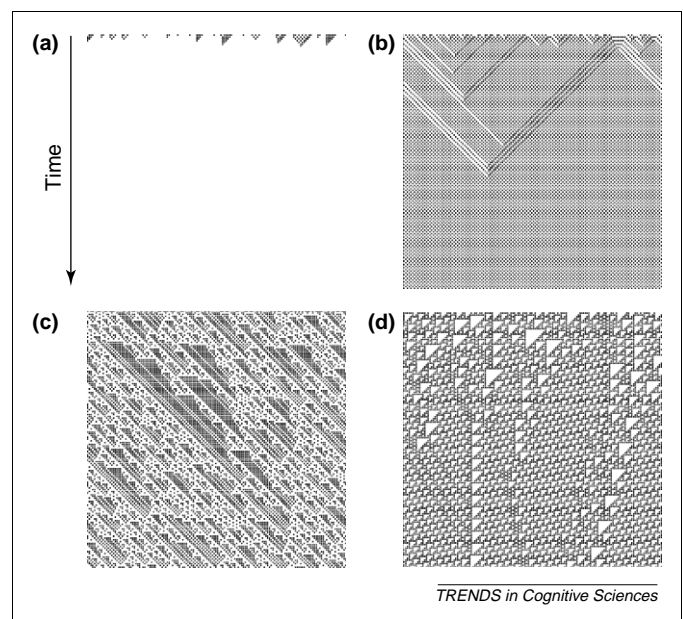


Figure 1. Cellular automata. A cellular automaton is a regular spatial lattice of ‘cells’, each of which can be in one of a finite number of states. (The lattice typically has 1, 2, or 3 spatial dimensions.) The state of each cell in the lattice is updated simultaneously in discrete time steps. Each 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. Typically all cells in the lattice are governed by the same update rule, which is usually deterministic. (a)–(d) show the change over time of four different one-dimensional cellular automata. Each cellular automaton is a strip of 300 cells, each of which is in one of two states (black or white). The initial state of each cell, which is chosen randomly, is at the top of each panel, and each subsequent state of that cell is immediately below the previous state. Thus, each panel is a space–time diagram showing how the cells change state over time given a particular random initial condition. The qualitative behavior exhibited in each case is typical for that cellular automaton; changing the initial condition typically does not change the type of global behavior that emerges. Note though that different simple cellular rules give rise to qualitatively different kinds of global behavior, some of which are quite complex. The cellular automaton in (d) is thought to be computationally universal, that is, capable of precisely modeling any computer algorithm whatsoever.

Corresponding author: Mark A. Bedau (mab@reed.edu).

contemporary artificial life. Complex systems are composed of many elements 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 [6,7], and these are the main focus of artificial life.

Artificial life also has roots in computer science, especially artificial intelligence (AI) and machine learning. Most notable here are John Holland's pioneering investigations of genetic algorithms and classifier systems [6,8] (see Box 1). The subjects of AI and artificial life overlap, as living and flourishing in a changing and uncertain environment seem to require at least a rudimentary form of intelligence. Their methodology is also similar: both study natural phenomena by simulating and synthesizing them. Nevertheless, there is an important difference between the modeling strategies of 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, and 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. They 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, and arises out of the interactions of directly represented parts ('agents' or 'individuals') with each other and with their physical and social environment. This decentralized architecture shares important similarities with some newer trends in AI,

Box 1. Genetic algorithm

The genetic algorithm is a machine-learning technique loosely modeled on biological evolution; it views learning as a matter of competition among candidate problem solutions. 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'. 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 reciprocal of the sum of the distances of each segment in the itinerary, so shorter-distance routes have higher fitness. 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. By 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 fittest candidates for many generations, the population contains better and better solutions.

including connectionism [9], multi-agent AI [10], and evolutionary computation [6,11].

A sample of the state of the art

Life exhibits complex adaptive behavior at many different levels of analysis: metabolic and genomic networks, single cells, whole organisms, social groups, evolving ecologies, and so forth. This diversity of levels gives artificial life a broad intellectual scope (see Box 2). The following highlights illustrate ascending levels in the hierarchy of living systems.

Dynamical hierarchies

Each level in a dynamical hierarchy consists of the interaction of entities with a distinctive set of autonomous functionalities, and the entities at higher levels are composed of entities at lower levels (for example, think of organs, which are made up of tissues, which are made up of cells, which are made up of organelles, which are made up of molecules, and so on). One of the fundamental open problems in artificial life is to explain how robust, multiple-level dynamical hierarchies emerge solely from the interactions of elements at the lowest-level. This is closely analogous to the problem in cognitive science of explaining how cognitive capacities ultimately emerge from the interactions of non-cognitive elements like neurons.

Box 2. Grand challenges in artificial life

A good way to understand an interdisciplinary science is through its central aims. The challenges of artificial life fall into three broad categories: the origin of life, life's evolutionary potential, and life's connection to mind and culture [1]. (Some areas in which artificial life plays a significant role, such as robotics and art, do not appear on the list, which is not meant to be comprehensive.)

How does life arise from the non-living?

- (1) Generate a molecular proto-organism *in vitro*.
- (2) Achieve the transition to life in an artificial chemistry *in silico*.
- (3) Determine whether fundamentally novel living organizations can arise from inanimate matter.
- (4) Simulate a unicellular organism over its entire lifecycle.
- (5) Explain how rules and symbols are generated from physical dynamics in living systems.

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.

John Conway's famous cellular automaton known as the 'Game of Life' produces many levels of emergent organization, but it is not robust; changing the state of even one micro-level element often destroys the whole hierarchical organization [12]. Living dynamical hierarchies, by contrast, are quite robust. Significant progress on understanding robust dynamical hierarchies would come from a computer model in which the explicit primitive operations govern objects at the lowest level, second-order objects emerge out of the interactions of those primitive objects, third-order objects emerge out of the interactions of the second-order objects, and so on. One step in this direction was taken by Fontana and Buss [13]. Their model first produced different kinds of second-order, self-organizing structures; then, by changing the model's boundary conditions and pooling those second-order structures, a new third-order structure was produced.

More recently, Rasmussen and colleagues [14] showed how to create wholly emergent third-order structures (i.e. without external manipulations by the experimenter) by creating a realistic simulation of the formation of micelles out of monomer molecules in a three-dimensional space. (Monomers are molecules that can combine with others to form polymers, and micelles are aggregations of many polymers.) Their model was the first to exhibit two levels of fully emergent phenomena – that is, one in which interactions among the primitive elements give rise to higher-level emergent entities with emergent properties, such that interactions among those emergent entities give rise to yet higher-level emergent entities with their own emergent properties (Figure 2).

Molecular self-organization

Many cognitive phenomena apparently involve spontaneous self-organization, and self-organization is also a central focus in artificial life. This is covered in this review at several levels (e.g. artificial cells, multicellularity, and swarm intelligence), and here we are concerned with self-organization at the molecular level. Kauffman and colleagues showed that many features of metabolic and genetic networks that are often thought to be adaptations can instead be viewed largely as the result of spontaneous self-organization – or what Kauffman terms 'order for free' [15]. Kauffman makes bold claims about order for free because the models he studies, such as random Boolean networks (see Box 3), are abstract enough to cover a great many real systems. It turns out that the number of variables in the network, the number of connections between the variables, and the character of the Boolean functions determine many biologically crucial properties of these networks. Biologically realistic Boolean networks show promise for explaining several empirically observed features of biological systems, such as how the number of different cell types and cell-replication times vary as a function of the number of genes per cell [15].

Another important strand of research in this area concerns the self-organization of self-replicating structures (see below). Models of systems such as cellular automata have shown that the emergence of stable self-replicating structures built out of cooperative molecules succumbs to parasitic invasion unless appropriate membranes or

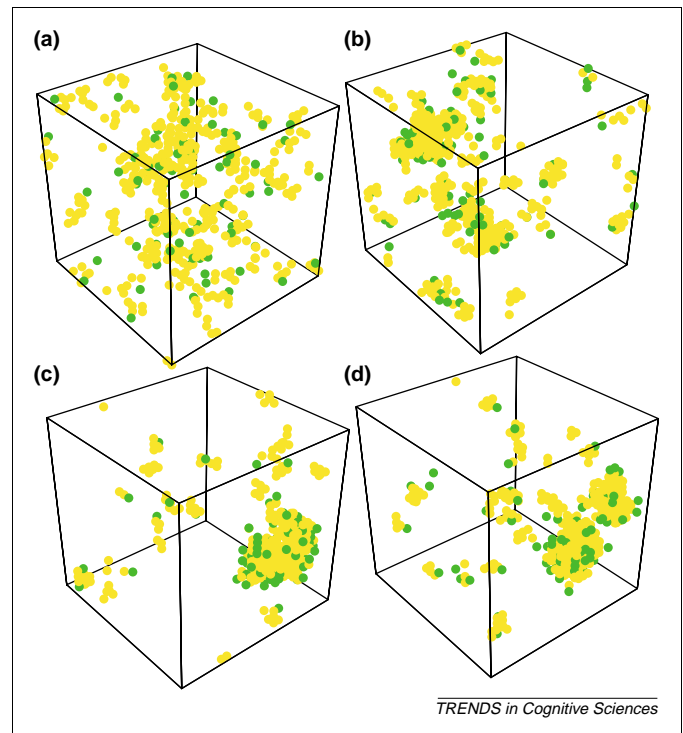


Figure 2. Four stages in the self-assembly of third-order emergent structures, in a three-dimensional lattice in which each side represents about 9 nanometers (see [14]). The time between these snapshots corresponds to hundreds of microseconds, or only tens of microseconds between (c) and (d). (a) shows first-order hydrophilic (yellow) and hydrophobic (green) monomers that are organized into amphiphilic polymers among water molecules (not shown). The second-order polymers in (a) and (b) can be seen to have emergent properties, like elasticity, that the first-order monomers cannot possess. (c) shows the emergence of micelles – third-order structures generated from the dynamics of the second-order polymers. The micelles have an inside and an outside and are permeable – properties that amphiphilic polymers cannot possess. (d) shows one of the micelles dividing (right).

compartments self-organize [16]. Recently, it has been shown that independent replicators can self-organize into cooperative molecular systems *in vitro* [17], and the self-assembly of two-dimensional compartments with flexible, semipermeable, self-organizing membranes has also been demonstrated in computer simulations using reconfigurable hardware [18].

Self-replication

Self-replication is a crucial property of living systems at many levels of analysis. Von Neumann's original self-reproducing cellular automaton [3] contained a universal constructor, capable of constructing any cellular automaton

Box 3. Random Boolean networks

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 discrete step in 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. Because the network is finite, it eventually reaches a state it has previously encountered, and from then on the network will 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.

configuration. Langton noted that natural self-reproducing living systems were not universal constructors, and he produced a vastly simpler self-replicating cellular automaton that lacked a universal constructor [19]. Langton's configuration replicates because a description of it is both copied into, and interpreted in the daughter configuration in order to complete its construction. Langton's result sparked a series of simple self-replicating systems (reviewed in [20,21]). Self-replication is not sufficient for life, of course, and Langton-style self-replicating configurations seem more like growing crystals than living organisms because they are unable to evolve. Sayama has overcome that limitation by creating self-replicating cellular automata that evolve and can become more complex [22], and Ikegami and co-workers have recently demonstrated the emergence of more realistic self-replicating structures [23].

Artificial cells

The holy grail of 'wet' artificial life is to create artificial cells out of biochemicals. Such artificial cells would be microscopic, autonomously self-organizing and self-replicating physical entities that assemble themselves out of non-living materials. Although artificial, they would repair themselves and adapt in an open-ended fashion, so for all intents and purposes they would be alive. The first artificial cells will probably just move through a fluid and process chemicals. To do even this flexibly and robustly, they must solve the functions of self-maintenance, autonomous control of chemical processing, autonomous control of mobility, and self-replication.

Nobody has yet created an artificial cell, but research is pursuing two approaches. Venter and Smith are using the top-down strategy of artificially synthesizing and modifying the genome of the organism with the smallest genome, the bacterium *Mycoplasma genitalium* [24]. The other approach is bottom-up, building more and more complex physiochemical systems that increasingly incorporate life-like properties. Some bottom-up efforts are strongly inspired by the RNA chemistry in existing cells [25,26], whereas others pursue a simpler chemistry that replaces RNA with PNA (peptide nucleic acid, an analog of DNA in which the backbone is a pseudopeptide rather than a sugar) [27]. Laboratory experiments must be compared with simulations of artificial cells, so there is overlap here with systems biology, which aims to represent and simulate the details of real cells [28]. However, artificial life has the broader aim of representing and synthesizing all life-like systems, including those that do not yet exist.

Evolution of genetic code

The genetic code is a nearly universal feature of life on earth, and yet it is difficult to understand how it could evolve. A mutation at a given codon (i.e. three adjacent bases on a strand of DNA or RNA that codes for a particular amino acid) might be a selective advantage, of course, but any change in the code would entail wholesale changes at a vast number of codons. How could this be anything but disastrous? Kaneko, Takagi and colleagues [29] have proposed a solution to this problem by numerically simulating a dynamic system that models an intracellular metabolic-reaction network, interactions between

cells (via chemical diffusion), and mechanisms for cell division (and death) and mutation. The system's variables are concentrations of metabolic chemicals, metabolic enzymes, chemicals for genetic information, and enzymes to translate genetic information into amino acids. The genetic code determines how genetic information is expressed phenotypically, that is, which amino acids are produced, and evolution of the genetic code takes the form of changes in the enzymes used in amino-acid synthesis. Simulations using this model have shown that evolution can create cells with distinct genetic codes.

Origin of multicellularity

The origin of multicellularity is one of the major transitions in the evolution of life, and apparently has evolved independently many times. The developmental processes that create these multicellular organizations share three features. First, development starts with a homogeneous cell (or set of cells) that are multipotent (i.e. they have the ability to differentiate into many different types of specialized cells). Second, the developmental process has an intrinsic temporal direction because the differentiated cells that result are not multipotent. Third, the developmental process is stable in the face of perturbations that destroy clusters of cells.

Alan Turing's pioneering work on cellular inhomogeneity over fifty years ago showed that an instability in a homogeneous chemical system could generate the formation of patterns by a process of chemical reaction and diffusion [30]. Furusawa and Kaneko have now extended Turing's dynamic systems approach to explain the three universal features of multicellular differentiation mentioned above [31]. They simulated a one-dimensional chain of cells governed by randomly generated biochemical reaction networks, and showed that differentiated cells can grow more quickly than undifferentiated cells.

Evolutionary robotics

Much work at the multicellular level has occurred in 'hard' artificial life, which is concerned with various forms of autonomous agents such as robots. This is artificial life's most direct overlap with cognitive science, as its aim is to synthesize autonomous adaptive and intelligent behavior in the real world. One of the tricks is to allow the physical environment to generate the behavior as far as possible. Brooks pioneered a biologically inspired approach to robotics [32–34] (and see [35] for a recent textbook).

Traditional rational design for intelligent autonomous agents is difficult, because it involves sophisticated interconnections among many complex components. The ALife alternative is to follow nature and use an evolutionary design method [36], employing a genetic algorithm [6,11] (see [Box 1](#)). This method can be used to design many aspects of robots, including control systems and sensors [37–39]. In natural autonomous agents, the control system is tightly coupled with morphology. Sims took some steps towards capturing this interconnection when he simultaneously co-evolved simulated creatures' controllers, sensors, and morphology [40], and advances in hardware and software have recently stimulated this line of research [37,38,41]. Jordan Pollack and his students have

demonstrated that co-evolution of controllers and morphology can be constrained by off-the-shelf rapid prototyping technology, so that evolutionary design could be automatically replicated in the real world [42,43].

Evolution of digital organisms

The first significant achievement of spontaneous evolution in a digital medium was Ray's *Tierra* [44]. Studying evolving systems in software is one of the most practical and constructive ways to explore any evolutionary explanations, such as those offered in evolutionary psychology. *Tierra* consists of a population of digital 'creatures' that are actually simple, self-replicating computer programs populating computer memory and consuming CPU time. A *Tierran* genotype consists of a string of machine code, and each *Tierran* creature is a token of a *Tierran* genotype.

A simulation starts when a single self-replicating program, the ancestor, is placed in computer memory and left to replicate (Figure 3). The ancestor and its descendants repeatedly replicate until computer memory is teeming with creatures that all share the same ancestral genotype, so older creatures are continually removed from the system to create space in memory for new descendants. Errors (mutations) sometimes occur, so the population of *Tierra* creatures evolves by natural selection. If a mutation allows a creature 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 hosts 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.

Evolution of complexity

Life has exhibited a remarkable growth in complexity over its evolutionary history. Simple prokaryotic one-celled life developed into 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, sharks), the apparently open-ended growth in complexity of the most complex organisms is intriguing and enigmatic. Much effort in artificial life is directed towards creating a system that shows how this kind of open-ended evolutionary progress is possible, even in principle. Digital evolution in *Tierra* does not show this, because the complexity of individual *Tierran* creatures typically decreases and significant evolutionary change eventually peters out (Figure 3). 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 [45]. Furthermore, when *Tierra* is modified so that creatures are rewarded for performing complex arithmetic operations on numbers they find in their environment, evolution produces the expected increase in genetic complexity [46,47]. However, as with the original version of *Tierra*, these evolutionary progressions eventually stop.

Hillis demonstrated that co-evolution can spur evolutionary progress [48], and co-evolutionary arms races might be necessary to drive continual evolutionary progression by continually changing the environment. Even so, the original and modified versions of *Tierra* all involve co-evolution and yet the environment eventually becomes essentially stable, so there must be 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 [49,50]. We still are missing some insight about the mechanisms by which evolution continually creates new kinds of environments that elicit new kinds of adaptations.

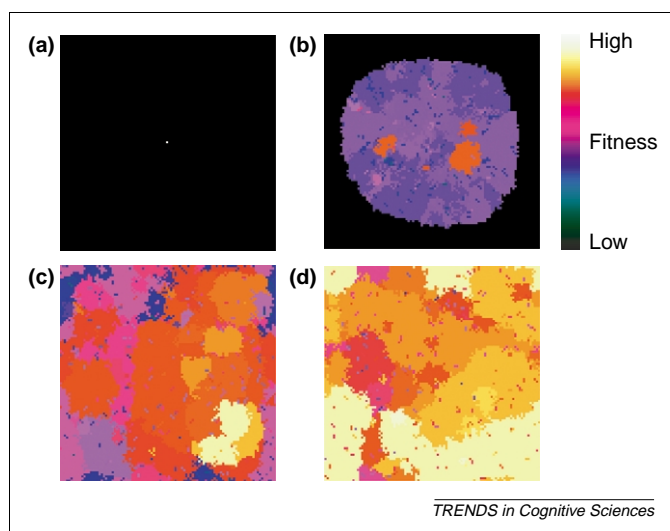


Figure 3. Four snapshots of the evolutionary dynamics exhibited by a typical run of *Avida* [46,47], a derivative of *Tierra* [44]. *Avida* consists of a population of self-replicating computer programs, each of which occupies one location in a two-dimensional grid with periodic boundary conditions (i.e. the left and right edges of the grid are connected, as are the top and bottom edges). The system is started by seeding the grid with one self-replicating program that is designed by hand (a). As time evolves, this program replicates and spreads on the grid. Random mutations change the fitness of offspring programs, and higher fitness programs spread more quickly (b–d). Fitness is determined by replication rate and success at performing various logical operations. The evolving population of programs eventually fills the 100 × 100 grid (c,d). Regions containing high-fitness programs typically contain miscellaneous isolated low-fitness programs that arise by deleterious mutations and quickly die.

Evolvability

Evolvability – the capacity of evolution to create new adaptations – depends on a system’s ability to produce adaptive phenotypic variation, and this hinges on both the extent to which phenotype space contains adaptive variation and the ability of evolutionary search to find it. For evolutionary search to explore a suitable variety of viable evolutionary pathways, genetic operators must generate enough appropriate evolutionary novelty. At the same time, evolutionary memory is needed to retain incremental improvements discovered over time. Evolvability requires successfully and flexibly balancing these competing demands for novelty and memory; this is known as the ‘explore–exploit’ trade-off in the machine learning literature [6].

One way to address this trade-off is to study models in which second-order evolution controls the genetic mechanisms that structure first-order evolution, thus studying the evolution of evolvability. The simplest such genetic mechanism is the mutation rate. Optimal and evolving mutation rates are studied in the evolutionary computation literature, where the primary focus is on quickly solving optimization problems (see [51] for a recent review). As this work presumes that evolution is driven by a fixed and externally-specified fitness function, it might not illuminate the evolvability in settings with implicitly specified fitness functions that unpredictably change over time. (See [52] for the most recent work on evolvability in this context.)

Swarm intelligence

Many organisms live in social groups, and artificial life uses bottom-up models to explore how the structure and behavior of social groups arises and is controlled. The simplest examples concern the social organization of social insects. Distributed networks of relatively simple insects give rise to complex collective behaviors, involving foraging, nest building, transporting resources, and the like (for a review, see [53]). These collective behaviors are remarkably flexible, robust and autonomous.

The attempt to design algorithms or distributed problem-solving methods inspired by the collective behavior of insect societies has come to be called ‘swarm intelligence’, and it has obvious relevance to distributed methods in cognitive science. Recent advances in swarm intelligence include a mathematical theory [54] of how groups of robots work together to solve group goals, enabling quantitative and well as qualitative comparison between theory and experiment involving robot swarms, and an explicit formal correspondence [55] between some swarm intelligence algorithms and the technique known as stochastic gradient descent which is extensively used in machine learning [56].

Artificial economics

At the opposite end of the social behavior spectrum is the behavior of groups of humans. Although humans have vastly more sophisticated cognitive capacities than insects, bottom-up models of the behavior of human groups are prevalent in artificial life, especially in application to economics. Decentralized economic markets consist of

large numbers of adaptive agents involved in parallel local interactions. Those local interactions give rise to macroeconomic regularities such as price structures that themselves influence the local interactions, and so on. The result is a complex adaptive system with recurrent causal chains connecting individual agent behavior with macroscopic market regularities.

Agent-based, bottom-up models of economic systems – what is sometimes called ‘artificial economics’ – posit a population of economic agents (e.g. traders or financial institutions) governed by internal procedures (e.g. rules for when to buy or sell, and beliefs about the conditions of the market). This approach has been applied to a variety of complex phenomena associated with decentralized economies, including inductive learning, the formation of trade networks, the evolution of behavioral norms, and the open-ended evolution of economic markets (for a review, see [57]).

Evolution of language

The most complicated behavior exhibited by humans is probably language. Indeed, language has been termed the most complex natural system [58], because it results from the interaction of three complex adaptive systems that operate on different timescales: language learning on an ontogenetic timescale, language evolution on a historical timescale, and the evolution of the brains of language users on a phylogenetic timescale. It is no surprise, then, that bottom-up, artificial life methods are increasingly being used to explain many aspects of language, including phonetics and phonology, language acquisition, language change, the evolution of signaling systems, the grounding of symbols and the evolution of meanings, the emergence of complex structured languages, and the co-evolution of languages and language learning mechanisms (for reviews see [58,59]).

Practical applications of artificial life

One measure of a scientific field’s success is its usefulness for solving practical problems. By this criterion, artificial life is a success today mainly because of applications that exploit genetic algorithms [6,11] and offshoots like genetic programming [60]. But biologically-inspired methods are increasingly applied to technological problems, such as using immune-system principles and mechanisms to protect computer systems against attacks by computer viruses and worms [61], and designing novel strategies for navigation of autonomous flight systems [62]. In addition, our increased understanding of real biological systems is enabling us to control them better. For example, artificial life is helping to illuminate why normal cells evolve into cancerous cells [63]. Finally, artificial life is used for a variety of aesthetic purposes. There are artificial-life approaches to music composition [64], and the technologists’ journal *Leonardo* regularly publishes papers concerning ALife.

The connection between life and mind

Living systems exhibit various forms of ‘cognitive’ capacities in its broadest sense. Plants, bacteria, insects and mammals all are sensitive to their environment in ways

that affect their behavior, and they all exhibit memory and inter-organism communication. Furthermore, the relative sophistication of such capacities seems roughly to reflect and explain a life form's relative biological complexity.

Evolution presumably explains much about the origin and nature of the mental capacities found in nature, as it does for any adaptation. But life and mind might have a deeper and more distinctive connection, and emergence might be a clue to it. Cognitive science tries to explain how cognitive capacities ultimately emerge from the behavior of biological materials like neurons, with qualitatively different properties, and artificial life tries to explain how life emerges from a non-living substrate of molecules. This methodological affinity suggests an affinity of subject matter. However, physics explains how the solid, liquid and gaseous phases of water emerge out of the behavior of water molecules that individually are neither solid, liquid or gaseous, so it cannot be said that emergence alone entails a close connection between life and mind.

Adaptation might explain a tighter link. Open-ended adaptability is a hallmark of life, at least when considered on an evolutionary time scale [65]. Similarly, following Maturana and Varela [66], many consider the ability to cope with a complex, dynamic, unpredictable environment to be a defining feature of cognitive and intelligent systems [32,34,65,67,68]. This line of argument, which is being explored by the embodied cognition and dynamical systems approaches to cognition [34,68,69], implies a fundamental similarity in the key mechanisms behind living and cognitive systems. This conclusion is supported by Brooks's conclusion that the future of both AI and artificial life hinges on bridging the gap between non-living and living matter [70]. If Brooks is right, then we can expect 'wet' artificial life to become inseparably intertwined with 'soft' and 'hard' artificial life, and we can expect all three to merge with cognitive science.

Conclusions

Artificial life is an interdisciplinary investigation into one of the most fundamental aspect of the natural world – life itself. Its synthetic methodology is making incremental progress on a wide range of issues, from dynamical hierarchies and artificial cells to the evolution of complexity and even language. Its ambitious agenda for the future involves explaining how life arises from non-living substrates, determining the potentials and limits of living systems, and understanding how life is related to mind, machines, and culture. This overlapping agenda means that artificial life is likely to change the future face of cognitive science in significant ways.

Acknowledgements

For many helpful suggestions on earlier drafts of this review, thanks go to the editor and anonymous referees.

References

- 1 Bedau, M.A. *et al.* (2000) Open problems in artificial life. *Artif. Life* 6, 363–376
- 2 Langton, C. ed. (1989) *Artificial Life* Addison-Wesley
- 3 Von Neumann, J. (1966) *Theory of Self-Reproducing Automata*, University of Illinois Press
- 4 Wiener, N. (1948) *Cybernetics, or Control and Communication in the Animal and the Machine*, Wiley
- 5 Wolfram, S. (1994) *Cellular Automata and Complexity*, Addison-Wesley
- 6 Holland, J.H. (1975/1992) *Adaptation in Natural and Artificial Systems: An Introductory Analysis with Applications to Biology, Control and Artificial Intelligence* (2nd edn, expanded, 1992), University of Michigan Press
- 7 Holland, J.H. (1995) *Hidden Order: How Adaptation Builds Complexity*, Helix Books
- 8 Holland, J.H. (1986) Escaping brittleness: the possibilities of general-purpose learning algorithms applied to parallel rule-based systems. In *Machine Learning II* (Michalski, R.S. *et al.*, eds), pp. 3–25, Morgan Kaufmann
- 9 Rumelhart, D.E. and McClelland, J.L. (1986) *Parallel Distributed Processing: Explorations in the Microstructure of Cognition* (Vols 1 & 2), MIT Press
- 10 Rich, E. and Knight, K. (1991) *Artificial Intelligence*, McGraw-Hill
- 11 Mitchell, M. (1996) *An Introduction to Genetic Algorithms*, MIT Press
- 12 Berlekamp, E.R. *et al.* (1982) *Winning Ways* (Vol. 2), Academic Press
- 13 Fontana, W. and Buss, L.W. (1994) The arrival of the fittest: toward a theory of biological organization. *Bull. Math. Biol.* 56, 1–64
- 14 Rasmussen, S. *et al.* (2001) Ansatz for dynamical hierarchies. *Artif. Life* 7, 329–353
- 15 Kauffman, S.A. (1993) *The Origins of Order: Self-Organization and Selection in Evolution*, Oxford University Press
- 16 Boerlijst, M.C. and Hogeweg, P. (1991) Spiral wave structures in prebiotic evolution: hypercycles stable against parasites. *Physics D* 48, 17–28
- 17 McCaskill, J.S. (1997) Spatially resolved *in vitro* molecular evolution. *Biophys. Chem.* 66, 145–158
- 18 Breyer, J. *et al.* (1998) Evolving reaction-diffusion ecosystems with self-assembling structures in thin films. *Artif. Life* 4, 25–40
- 19 Langton, C.G. (1986) Studying artificial life with cellular automata. *Physica D* 22, 120–149
- 20 Sipper, M. (1998) Fifty years of research on self-replication: an overview. *Artif. Life* 4, 237–257
- 21 Sipper, M. and Reggia, J.A. (2001) Go forth and replicate. *Sci. Am.* 285, 34–43
- 22 Sayama, H. (2000) Self-replicating worms that increase structural complexity through gene transmission. In *Artificial Life VII* (Bedau, M.A. *et al.*, eds), pp. 21–30, MIT Press
- 23 Ono, N. and Ikegami, T. (2003) Selection of catalysts through cellular reproduction. In *Artificial Life VIII* (Standish, R. *et al.*, eds), pp. 57–64, MIT Press
- 24 Hutchison, C.A. *et al.* (1999) Global transposon mutagenesis and a minimal Mycoplasma genome. *Science* 286, 2165–2169
- 25 Szostak, J.W. *et al.* (2001) Synthesizing life. *Nature* 409, 387–390
- 26 Pohorille, A. and Deamer, D. (2002) Artificial cells: prospects for biotechnology. *Trends Biotechnol.* 20, 123–128
- 27 Rasmussen, S. *et al.* (2003) Bridging nonliving and living matter. *Artif. Life* 9, 269–316
- 28 Kitano, H. (2002) Systems biology: a brief overview. *Science* 295, 1662–1664
- 29 Takagi, H. *et al.* (2000) Evolution of genetic codes through isologous diversification of cellular states. *Artif. Life* 6, 283–306
- 30 Turing, A. (1952) The chemical basis of morphogenesis. *Philos. Trans. R. Soc. B Biol. Sci.* 237, 37–72
- 31 Furusawa, C. and Kaneko, K. (2002) Origin of multicellular organisms as an inevitable consequence of dynamical systems. *Anat. Rec.* 268, 327–342
- 32 Brooks, R.A. (1989) A robot that walks: emergent behaviours from a carefully evolved network. *Neural Comput.* 1, 253–262
- 33 Brooks, R.A. (1990) Elephants don't play chess. *Robots Auton. Syst.* 6, 3–15
- 34 Brooks, R.A. (1991) Intelligence without representation. *Artif. Intell. J.* 47, 139–160
- 35 Pfeifer, R. and Scheier, C. (1999) *Understanding Intelligence*, MIT Press
- 36 Cliff, D. *et al.* (1993) Explorations in evolutionary robotics. *Adapt. Behav.* 2, 73–110
- 37 Nolfi, S. and Floreano, D. (2000) *Evolutionary Robotics: the Biology, Intelligence and Technology of Self-Organizing Machines*, MIT Press

- 38 Nolfi, S. and Floreano, D. (2002) Synthesis of autonomous robots through evolution. *Trends Cogn. Sci.* 6, 31–37
- 39 *Artificial Life 7* (2), Special issue on sensor evolution (Dautenhahn, K. et al., eds)
- 40 Sims, K. (1994) Evolving 3D morphology and behavior by competition. In *Artificial Life IV* (Brooks, R. and Maes, P., eds), pp. 28–39, MIT Press
- 41 Taylor, T. and Massey, C. (2001) Recent developments in the evolution of morphologies and controllers for physically simulated creatures. *Artif. Life* 7, 77–87
- 42 Lipson, H. and Pollack, J.B. (2000) Automatic design and manufacture of robotic lifeforms. *Nature* 406, 974–978
- 43 Pollack, J.B. et al. (2001) Three generations of automatically designed robots. *Artif. Life* 7, 215–223
- 44 Ray, T.S. (1992) An approach to the synthesis of life. In *Artificial Life II* (Langton, C.G. et al., eds), pp. 371–408, Addison-Wesley
- 45 Ray, T.S. (2000) Evolution of differentiation in multithreaded digital organisms. In *Artificial Life VII* (Bedau, M.A. et al., eds), pp. 132–140, MIT Press
- 46 Adami, C. et al. (2000) Evolution of biological complexity. *Proc. Natl. Acad. Sci. U.S.A.* 97, 4463–4468
- 47 Lenski, R.E. et al. (2003) The evolutionary origin of complex features. *Nature* 423, 139–144
- 48 Hillis, D. (1992) 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
- 49 Bedau, M.A. et al. (1997) A comparison of evolutionary activity in artificial systems and in the biosphere. In *Proceedings of the Fourth European Conference on Artificial Life* (Husbands, P. and Harvey, I., eds), pp. 125–134, MIT Press
- 50 Bedau, M.A. et al. (1998) A classification of long-term evolutionary dynamics. In *Artificial Life VI* (Adami, C. et al., eds), pp. 228–237, MIT Press
- 51 Eiben, A.E. et al. (1999) Parameter control in evolutionary algorithms. *IEEE Trans. Evol. Comput.* 3, 124–141
- 52 *BioSystems* 69 (2–3) Special double issue on evolvability
- 53 Bonabeau, E. et al. (1999) *Swarm Intelligence: From Natural to Artificial Systems*, Oxford University Press
- 54 Lerman, K. et al. (2001) A macroscopic analytical model of collaboration in distributed robotic systems. *Artif. Life* 7, 375–393
- 55 Meuleau, N. and Dorigo, M. (2002) Ant colony optimization and stochastic gradient descent. *Artif. Life* 8, 103–121
- 56 Mitchell, T. (1997) *Machine Learning*, McGraw Hill
- 57 Tesfatsion, L. (2002) Agent-based computational economics: growing economies from the bottom up. *Artif. Life* 8, 55–82
- 58 Kirby, S. (2002) Natural language from artificial life. *Artif. Life* 8, 185–215
- 59 Christiansen, M. and Kirby, S. (2003) Language and evolution: consensus and controversies. *Trends Cogn. Sci.* 7, 300–307
- 60 Koza, J.R. (1992) *Genetic Programming: On the Programming of Computers by Means of Natural Selection*, MIT Press
- 61 Kephart, J. (1994) A biologically inspired immune system for computers. In *Artificial Life IV* (Brooks, R. and Maes, P., eds), pp. 130–139, MIT Press
- 62 Thakoor, S. et al. (2002) Bioinspired engineering of exploration systems for NASA and DoD. *Artif. Life* 8, 357–369
- 63 Maley, C.C. and Forrest, S. (2000) Exploring the relationship between neutral and selective mutations in cancer. *Artif. Life* 6, 325–346
- 64 Todd, P.M. and Werner, G.M. (1999) Frankensteinian methods for evolutionary music composition. In *Musical Networks: Parallel Distributed Perception and Performance* (Griffith, N. and Todd, P.M., eds), pp. 313–340, MIT Press
- 65 Bedau, M.A. (1998) Four puzzles about life. *Artif. Life* 4, 125–140
- 66 Maturana, H. and Varela, F. (1973) *Autopoiesis: the Organization of the Living*, Reidel
- 67 Beer, R.D. (1990) *Intelligence as Adaptive Behavior: An Experiment in Computational Neuroethology*, Academic Press
- 68 Port, R. and van Gelder, T. (1995) *Mind as Motion: Explorations in the Dynamics of Cognition*, MIT Press
- 69 *Artificial Life* Special issue on embodied cognition and the new robotics (in press)
- 70 Brooks, R.A. (2001) The relationship between matter and life. *Nature* 409, 409–411

The World Color Survey (WCS) data are now available on the web at
<http://www.icsi.berkeley.edu/wcs/data.html>

These data represent the color naming responses of an average of 24 (mode 25) respondents in each of 110 unwritten languages worldwide. The color stimuli for the WCS were 330 Munsell color chips representing 40 equally spaced Hues at eight levels of Value (lightness) and maximum Chroma (saturation) plus 10 Neutral (achromatic) chips in equal Value steps. Each respondent named each chip and then picked out from the full stimulus array the best (focal) example(s) of each term that the field worker judged to be of major importance.

The data were gathered in the period 1976–1979 by field linguists of the Summer Institute of Linguistics, with direction by Brent Berlin, Paul Kay and William Merrifield. The web archive was created by Richard Cook, with direction by Paul Kay and Terry Regier.

The data in this archive embody significant corrections of earlier versions of the data, including corrections based on communications with the original field workers.

No earlier version of these data should be considered valid for scientific research.