

53 Artificial Life

Wolfgang Banzhaf¹ · Barry McMullin²

¹Department of Computer Science, Memorial University
of Newfoundland, St. John's, NL, Canada

banzhaf@cs.mun.ca

²Artificial Life Lab, School of Electronic Engineering, Dublin City
University, Ireland

barry.mcmullin@dcu.ie

1	<i>Introduction and Historical Overview</i>	1806
2	<i>Fundamental Questions</i>	1808
3	<i>Theory of and Formalisms in Artificial Life</i>	1817
4	<i>Applications</i>	1822
5	<i>Conclusion</i>	1826

Abstract

Artificial life has now become a mature inter-discipline. In this contribution, its roots are traced, its key questions are raised, its main methodological tools are discussed, and finally its applications are reviewed. As part of the growing body of knowledge at the intersection between the life sciences and computing, artificial life will continue to thrive and benefit from further scientific and technical progress on both sides, the biological and the computational. It is expected to take center stage in natural computing.

1 Introduction and Historical Overview

The phrase *artificial life*, referring to a specific coherent research programme, is normally attributed to Chris Langton. He specifically applied the term as a title for the “interdisciplinary workshop on the synthesis and simulation of living systems” that he organized in September 1987, in Los Alamos, New Mexico (Langton 1989). It was then adopted for a biannual international conference series (now up to *Artificial Life XI*, held in Winchester UK, in August 2008), and alternating with the biannual *European Conference on Artificial Life* (ECAL), first held in Paris in 1991. The *Artificial Life journal* (with Langton as first Editor-in-Chief) was founded in 1993. The journal and the two conference series are now formally coordinated through the International Society for Artificial Life (ISAL) (<http://www.alife.org>), which was established in 2001.

Langton’s original announcement of the 1987 workshop defined the new field as follows:

- ▶ Artificial life is the study of artificial systems that exhibit behavior characteristic of natural living systems. This includes computer simulations, biological and chemical experiments, and purely theoretical endeavors. Processes occurring on molecular, cellular, neural, social, and evolutionary scales are subject to investigation. The ultimate goal is to extract the logical form of living systems. (Langton 1987)

It is clear that, from the outset, Langton’s vision was for a very broad, thoroughly interdisciplinary, endeavor, and indeed that is how the field has largely developed. There has been a consistent focus on using tools of computer science to model and simulate biological systems; but also a specific goal of using *synthesis* to start an exploration of the space of *possible* life. Indeed, Langton’s own first use of the phrase “artificial life” was in his paper “*Studying artificial life with cellular automata*” (Langton 1986). This was presented two years before the first artificial life workshop, at the conference *Evolution, Games, & Learning* organized by the Los Alamos Center for Non-Linear Studies (CNLS) in May 1985. In this paper, Langton was already introducing “artificial life” as a direct lineal descendant and inheritor of von Neumann’s seminal research programme, begun in the late 1940s, combining “automata theory” with problems of biological organization, self-reproduction, and the evolution of complexity (von Neumann 1949, 1951, 1966).

But while von Neumann’s influence was certainly strong, there were many other contributions and precursors to the field. The phrase “artificial life” is, in part, a deliberate play on “Artificial Intelligence” or AI. The original premise of AI, dating from Turing’s 1950 paper on “*Computing machinery and intelligence*” (Turing 1950), was that “intelligence” was an intrinsically computational or software phenomenon. It could therefore be divorced from

any particular underlying hardware implementation, including all the particular biological details of terrestrial life. By the 1980s however, after 30 years of intensive and well-resourced efforts, AI was still showing little clear progress toward its central goal of realizing human-level intelligence by means of computer programs. Indeed, as early as 1977, Popper was already declaring:

- "... computers are totally different from brains, whose function is not primarily to compute but to guide and balance an organism and help it to stay alive. It is for this reason that the first step of nature toward an intelligent mind was the creation of life, and I think that should we artificially create an intelligent mind, we would have to follow the same path. (Popper and Eccles 1977)

In 1988, the mounting challenges and criticisms of GOFAI ("good old fashioned AI" (Haugeland 1989)), were summarized in a special edition of *Daedalus*, the Journal of the American Academy of Arts and Sciences (Graubard 1988). This documented a renaissance of interest in the biological underpinnings of intelligence. Artificial life can therefore be seen as just one of the particular results of this splintering of AI, along with subsumption-architecture robotics (Brooks 1985), parallel distributed processing (Rumelhart et al. 1986) (the "new" connectionism), Neural Darwinism (Reeke and Edelman 1988; Edelman 1987), embodied intelligence (Varela et al. 1992), etc. And indeed, in some ways, artificial life continues to function as an integrative term, bridging between these more specialized and disparate research programmes.

A specific debate that artificial life has inherited from its roots in AI, is that between "weak" and "strong" forms. In the AI case, a distinction is made between the claim that an AI system provides a more or less effective *simulation* of intelligence, and the claim that it is "really" intelligent. The distinction is particularly associated with the philosopher John Searle, whose famous "Chinese Room Argument" claimed to prove that however well digital computers might simulate intelligence, computation could never be sufficient to genuinely *realize* intelligence (Searle 1980). Clearly an analogous distinction can be drawn within artificial life, between merely simulating and actually realizing "real" life (Levy 1993). In that case, it is generally focused on whether "life-like" organization, which exists only in a purely computational virtual world, could ever properly be described as genuinely living. However, Searle's abstract philosophical argument cannot be transferred directly into the artificial life domain as it relies on conscious introspection. Perhaps because of this, the issue has received comparatively less intense attention in this field. In any case, while much artificial life research is carried out in computational/virtual worlds, and thus might be subject to this kind of debate, it is generally considered that work in so-called wet artificial life, working in vitro with real chemical systems (e.g., Luisi and Varela 1989; Rasmussen et al. 2008), is certainly immune to this kind of critique.

Other significant influences and inspirations for artificial life include morphogenesis (Turing 1952), cybernetics (Wiener 1965), general systems theory (Bertalanffy 1976), mathematical/relational biology (Rashevsky 1940), theoretical biology (Waddington 1969), self-organization (Yovits and Cameron 1960), hierarchy theory (Pattee 1973), autopoiesis (Maturana and Varela 1973), and genetic and evolutionary algorithms (Fogel et al. 1966; Holland 1975; Rechenberg 1975).

A specific precursor to the modern artificial life conference series was the conference on *Automata, Languages, Development* (Lindenmayer and Rozenberg 1976) organized by Rozenberg and Lindenmayer, March–April 1975 in Noordwijkerhout, Holland. One of the

participants, Alvy Ray Smith, has subsequently commented that (<http://alvyray.com/Papers/PapersCA.htm>):

- ▶ I like to call this conference ALife0 — for Artificial Life 0 conference — since it was the first attempt I know to attempt cross-fertilization between biologists and computer scientists. Many of the players at this conference were present for ALife1, ALife2, etc [. . .] Other participants at ALife0 were Karel Culik, Pauline Hogeweg, John Holland, Aristid Lindenmayer, and Stanislaw Ulam.

In summary, artificial life is now a mature and well-established discipline (or “inter-discipline”) in its own right. The following sections will present the fundamental questions it is concerned with, a selection of specific theories and formalisms it relies on, and we review some applications and conclude with a discussion of the key open problems for future research.

2 Fundamental Questions

2.1 What Is (Artificial) Life?

- ▶ By extending the horizons of empirical research in biology beyond the territory currently circumscribed by life-as-we-know-it, the study of Artificial Life gives us access to the domain of life-as-it-could-be, and it is within this vastly larger domain that we must ground general theories of biology and in which we will discover practical and useful applications of biology in our engineering endeavors. (Langton 1995)

The question of “what is life?” is clearly fundamental to a discussion of the field of artificial life, as the answer (or answers) would be expected to underpin the scope, problems, methods, and results of its investigations.

The starting point is, of course, natural life, the domain of conventional biology. Historically, it was conjectured that the very distinctive properties and characteristics of natural living systems reflected a fundamental *constitutional* distinction: that living systems incorporated (or were literally “animated by”) a distinct and unique “substance” – the so-called *elan vital* that identified this class of theories as forms of *vitalism*. However, from the pioneering inorganic synthesis of urea by Wöhler in 1828 onward, vitalism progressively lost credibility. With the explosive growth of molecular understanding of basic biological mechanisms in the second half of the twentieth century (traditionally dated from the elucidation of the chemical structure of DNA in 1953), vitalism in this sense has now all but disappeared from scientific consideration. There is now a clear scientific consensus that there is no *material* distinction between the domains of living and nonliving systems; rather the differences (which are, of course, still objectively real and substantial) arise from differences in material organization.

However, it is still very much an open question as to what exactly these differences in organization *are*. Conventional molecular biology continues to make rapid progress in teasing out the detailed molecular basis for particular low-level mechanisms of operation in natural organisms; but by focusing on dismantling organisms into molecular components, these investigations, though still valuable in themselves, risk eliminating precisely the higher order organization that is distinctive of natural life. More recently, and partly in reaction to this, there has been growing interest in what is termed *systems biology*: The explicit study of how low-level molecular mechanisms are successfully integrated into higher level “systems” (Wolkenhauer 2007); and also, in *synthetic biology* that investigates the synthesis of entirely novel classes of living systems, either by manipulation of natural organisms and their components

(Smith et al. 2003; Regis 2008) or even by attempting de novo creation of life from entirely nonbiological substrates (Rasmussen 2008).

All these strands of investigation are still essentially premised on the study of existing natural life on Earth; or, as Langton puts it, “life-as-it-is” rather than “life-as-it-could-be” (Langton 1992b) (even synthetic biology, in its current state, largely relies on preexisting biological components, or close analogues of them).

By contrast, the field of artificial life, from its inception, has had a specific goal of investigating the most general possible formulations and instantiations of “living phenomenology.” In particular, it has explicitly brought to bear the possibility of creating computer-synthesized universes, or *virtual worlds*, in which even the most basic laws of “physics” or “chemistry” can be arbitrarily manipulated, in whatever ways may prove most conducive to any particular study.

Of course, this lack of hard boundaries makes “artificial life,” as a field of study, significantly ill-defined. Unlike the case for natural life, there are, as yet, no clear criteria for what virtual world phenomena should qualify as “living” or sufficiently “life-like” to legitimately count as lying within this field. In large measure, this simply reflects the continuing debate and investigation within conventional biology, of what specific organizational (as opposed to material) system characteristics are critical to properly living systems. The key advantage and innovation in artificial life is precisely that it has this freedom to vary and explore possibilities that are difficult or impossible to investigate in natural living systems. In this context, a precise definition of “life” (natural or artificial) is not a necessary, or even especially desirable condition for progress.

2.2 Hierarchy and Emergence

A specific and pervasive feature of natural life is hierarchical organization – from molecules to cells to tissues to organs to multicellular organisms to societies of organisms. Hierarchical design is also common in artificial engineered systems – for example, from semiconductor materials to logic gates to processors to computers to the Internet. At first sight, the hierarchical structure of natural living systems might be viewed as simply the natural outcome of evolutionary optimization for certain engineering benefits: modularity, standardization, ease of repair, robustness, etc. However, there are significant difficulties in such a simplistic account. In particular, evolution by natural selection must operate without foresight: So the opening up of evolutionary “potential” per se, via a new hierarchical level of organization, cannot be a direct locus for immediate selection to operate. The fundamental question here relates to the emergence and establishment of new levels of Darwinian “actors.” This is a process that has evidently occurred repeatedly in natural evolution (the origin of the first protokaryotes; the origin of the eukarotes; multicellularity, sociality, etc.); however, without circular reasoning, this cannot itself simply be assigned to a process of Darwinian selection (not, at least, at the same hierarchical level). In natural biology, this is the focus of investigation into so-called *major transitions in evolution* (Maynard Smith and Szathmáry 1997).

In the case of artificial life, while there have been some examples of model systems demonstrating some degree of emergence of hierarchical organization (e.g., Rasmussen et al. 2001b), the achievement to date has been limited, both in terms of evolutionary potential and in number of hierarchical levels, and is the subject of significant ongoing debate within the

community (Gross and McMullin 2001; Rasmussen et al. 2001a). This can be expected to be a particularly active area of continuing research (Baas 1994).

2.3 Constructive and Autopoietic Systems

A problem closely associated with emergence and hierarchy, yet somewhat distinct, is that of “self-constructing” systems. Natural life, at all hierarchical levels, exhibits a characteristic ability to maintain its system organization while simultaneously turning over components at the lower hierarchical levels. Thus, a cell can stably maintain itself (and even grow and divide) while continuously regenerating all its significant molecular components (metabolizing environmental substrates as necessary); similarly, a multicellular organism maintains its organization while constantly replacing its component cells; and a social insect colony may long outlive most, if not all, of its constituent individual organisms. While this is a typical phenomenon of natural life, it is very different from the behavior of conventional engineered systems. The latter may well have complex hierarchical organization, but this organization is typically static: The system is assembled once, and then retains its fixed compositional structure for its functional lifetime. Indeed, it is typical of engineered systems that failure of any component, at any hierarchical level, will at least significantly impair system level function, and will commonly result in full system failure. Although so-called “fault-tolerant” systems may be engineered to incorporate significant redundancy, this often involves static assembly of additional components at manufacture time. Functional lifetime is extended as failing components are automatically removed from service; but once this pre-built redundancy is exhausted, system failure follows from any further component failure. Further, even such redundant systems are generally vulnerable to failures in the “failure-detecting” components that cannot also simply be duplicated without potentially entering into infinite regress.

This is a complex problem to address, but natural life provides a “proof of principle” that such self-constructing organizations are possible. Artificial life certainly offers a potentially very fruitful avenue for further investigation. In particular, artificial life allows “virtual worlds” to be formulated in which problems of self-construction can be simplified, and the core organizational mechanisms can be exposed and subjected both to mathematical analysis and experimental exploration. While a variety of work in artificial life bears on this, some of the most direct contributions might be summarized as follows:

- The Chilean biologists Maturana and Varela pioneered the abstract concept of *autopoiesis* (literally “self-production”) as a description of the core mechanism of organizational self-maintenance in biological cells (Varela et al. 1974). Almost 15 years before the modern computational study of artificial life was even named (Langton 1989), they were already using a molecular level, agent-based, abstract chemistry to give a concrete demonstration of this abstract theory of cellular organization. In essence, the proposal is that biological cells are dynamic, self-sustaining, chemical networks (an abstraction of cell metabolism), which also create and maintain a spatial boundary (an abstraction of a cell membrane), where the contained network and the boundary reciprocally rely on each other for stability. While some technical deficiencies were identified in the original presentation of this work (McMullin and Varela 1997), it has given rise to a sustained and continuing programme of active research (McMullin 2004).

- Also in the 1970s, Holland proposed the so-called α -universes (Holland 1976) and, independently, Hofstadter described the *typogenetics* system (Hofstadter 1979). Both involved one-dimensional fragments of computer code, which could interact with each other. In both cases, they were explicitly inspired by aspects of the molecular replication and translation machinery of biology, and were concerned with understanding the reflexive interactions that arise when the “same” class of entities (molecules) can sometimes function as executable “code” (enzyme/protein) and sometimes as literal “data” (nucleic acid), and the possibility of collective self-reproduction and maintenance. These were, however, restricted to theoretical/analytic treatments, and were not subject to empirical investigation until much later (McMullin 1992; Snare 1999; Kvasnicka et al. 2001).
- This same tension between syntax and semantics in materially instantiated dynamic systems is also at the core of Pattee’s analysis of what he terms “semantic closure” (Rocha 2000; Cariani 1992).
- A related, but independent line of investigation was pioneered by Kauffman and others, in the form of “collectively autocatalytic sets” (Farmer et al. 1986; Kauffman 1993). Such a set is formally similar (though more strictly defined) to the contained chemical network of an autopoietic system, but without the requirement for a self-generated spatial boundary. The key result here was the demonstration, in completely abstract virtual “chemistries” that such collectively autocatalytic sets can spontaneously arise, quickly and under relatively weak constraints on the underlying chemistry.
- Fontana and Buss self-consciously launched a mathematical and computational investigation of “constructive dynamical systems,” based around the so-called *alchemy* system (Fontana and Buss 1994). These systems deliberately diverge from the classical concept of “dynamical systems” by relaxing the normally strict demarcation between “state variable” and “dynamic law.” The authors demonstrated and analyzed a number of organizational phenomena, including forms of “self-sustaining closure” where more-or-less complex dynamic aggregates of components successfully sustain themselves, even as all the individual components are continuously diluted or degraded.
- As a comprehensive overview treatment, Dittrich et al. (2001) provide a summary of the general approach of building computational models of arbitrary “artificial chemistries,” as a platform for investigating these and other problems. More recently, Dittrich and Speroni have developed a primarily algebraic mathematical analysis in a comprehensive way, in the form of “*Chemical organisation theory*” (Dittrich and di Fenizio 2007). However, while an important advance, the difficulties of incorporating dynamics (chemical kinetics) and the interaction of chemical self-production and spatial demarcation (the self-constructed “boundary” of autopoietic theory) remain substantial.
- A quite different and more radical approach to the problem of biological self-maintenance was proposed by the theoretical biologist Robert Rosen over an extended series of works (Rosen 1959, 1972, 1985, 1991). In brief, Rosen argues that the self-constructing organization characteristic of natural living systems, which he termed “closure under efficient causation” transcends the possibilities of conventional dynamical systems in a fundamental way; with the consequence that it cannot be realized in any “computational” universe. This is a complex and contentious theoretical claim which, not surprisingly, continues to be the focus of considerable critique and criticism (e.g., Chu and Ho 2006). Rosen’s work was also a strong influence in Kampis’ development of a comprehensive mathematical treatment of what he terms “self-modifying” systems (Kampis 1991).

2.4 Complexity (and Its Growth)

- There is a concept which will be quite useful here, of which we have a certain intuitive idea, but which is vague, unscientific, and imperfect . . . I know no adequate name for it, but it is best described by calling it “complication.” It is effectivity in complication, or the potentiality to do things. I am not thinking about how involved the object is, but how involved its *purposive operations* are. In this sense, an object is of the highest degree of complexity if it can do very difficult and involved things. (von Neumann 1949, p. 78, emphasis added)

Arguably the most fundamental problem in the theory of biology is that of the growth of “complexity”; with “complexity” understood in the sense explained by von Neumann above, of the ability to do complicated things. While this is certainly an informal definition, it suffices to clearly demarcate this notion from purely syntactic, combinatorial, or computational concepts of complexity.

Von Neumann is a key figure in the early development of abstract computational modeling approaches to understanding biological phenomena; indeed, in this respect he can be considered as having instigated the earliest investigations in artificial life, as the term is now understood (McMullin 2000a). In particular, in the late 1940s, (von Neumann 1949) he started to draw attention to an apparent paradox arising from the contrast between any mechanistic (i.e., not vitalistic) understanding of living organisms and common experience of engineering artificial mechanisms or automata. While it is quite generally possible to design machines that construct other machines, this process is normally *degenerative in complexity*: a machine of a given complexity (such as an automated factory) can only construct machines of comparatively lower complexity (such as cars or phones or televisions etc.). Whereas, if the theory of Darwinian evolution is accepted, then biological “machines” (organisms) must be capable of constructing descendant machines of greater complexity. Granted, these increases in complexity may occur only in very small increments, and even then only in some lineages and accumulating over geological periods of time; but nonetheless it must be possible, in principle, for biological organisms to give rise to offspring more complex than themselves.

Von Neumann quickly developed an outline resolution of this paradox, in the abstract form of what he termed a “general constructive automaton.” This was inspired by Turing’s earlier formulation of a universal computing automaton (Turing 1936), and also by von Neumann’s own contemporaneous involvement in designing and building the earliest electronic stored program digital computers (von Neumann 1945). In essence, a general constructive automaton is a *programmable* constructor or assembler, capable of constructing an indefinitely large set of target automata – loosely, a “universal” set comprised of all automata that could be assembled from a given, finite, set of primitive components. Significantly, von Neumann hypothesized that a general constructive automaton, relative to a specific set of components, might itself be realizable as an assembly of these same components. With this conceptual architecture and some relatively minor technical elaboration, von Neumann showed that this would then give rise to an indefinitely large set of automata, spanning an indefinite large range of complexity, all of which would be fully connected by a network of heritable “mutations” (essentially, chance variations in the construction programs). That is, there would exist mutational pathways leading from the simplest to arbitrarily complex members of this set; with all these machines also, incidentally, being capable of self-reproduction. (While terminology in the field is not completely consistent, in this

chapter “(self-)reproduction” is reserved to mean this von Neumann style process, involving separate actions of copying and decoding; whereas “(self-)replication” is used to denote a process of copying only.)

In retrospect, this was already an astonishing achievement. Von Neumann effectively described the abstract architecture of biological self-reproduction, based on separate processes of syntactic copying (replication) and semantic “decoding” (translation) of a passive information carrier; he showed how this architecture supported self-reproduction, heritable mutation, and thus evolutionary growth of complexity; and he presented all these, at least in outline form, in 1948, five years before the chemical structure of DNA was even identified (von Neumann 1949).

However, as yet this was still only a sketch of a solution. To make it properly convincing, he needed to present a concrete set of “primitive parts” and show that with these, it would actually be possible to realize an example of a general constructive automaton. This proved to be a complicated and potentially intractable problem: any plausible “real world” set of components would introduce many ancillary complications – mechanics, thermodynamics, etc. Following a suggestion from Stanislaw Ulam, von Neumann instead formulated an artificial, virtual, world that would simplify away these ancillary complications. He proposed a two-dimensional “tessellation automaton,” or, as it is now called, a *cellular automaton* (see Sect. 3.1.1 for a general introduction to CA). Within this simplified virtual world, von Neumann then successfully developed a fully detailed design for an example of a general constructive automaton. He had planned this as a first step in a much more general and comprehensive “theory of automata.” However, he put his unfinished manuscript aside in 1953; and due to his untimely death in 1957, was never to return to this project. The manuscript was subsequently edited by Burks and formally published in 1966 (Burks 1966).

Von Neumann’s achievement, both in theoretical insight and technical implementation, was considerable, and it gave rise to an extended program of further research. Many different cellular automata worlds have been formulated and investigated, and many variants on von Neumann’s architecture for evolvable self-reproducing automata have been demonstrated (Thatcher 1970; Codd 1968; Langton 1984b). Conway’s *Game of Life* has particularly popularized the study of life-like phenomena in cellular automata worlds (Berlekamp et al. 1982). More recently, it has become feasible to actually implement von Neumann’s original designs on real computers (Pesavento 1995).

Nonetheless, seminal as this work was, it also explicitly left certain questions still very much unresolved. In particular, as von Neumann himself was clearly aware, the mere existence of mutational pathways from simple to complex automata does not mean that evolution will actually follow such pathways. Indeed, in von Neumann’s original cellular automaton world, the embedded self-reproducing machine was extremely fragile. It could only operate successfully at all on the condition that the rest of the world was completely quiescent; and, of course, once it completes one round of reproduction, that could no longer be the case. Parent and offspring would then interfere with each other’s structure and operation, and both would rapidly distintegrate. While this problem could be superficially avoided by constraining each complete, successively constructed automaton to restrict its further operation to separate regions of the space, it rules out any possibility of natural selection; and, in any case, cannot be sustained in the long term if there is any possibility of stochastic, component level, malfunction, or failure (which itself is actually *desirable*, as a source of the random variation that is the grist to the mill of Darwinian evolution).

2.5 Coreworlds: Spontaneous Evolution of Computer Programs

- ▶ Nothing in Biology makes sense except in the light of evolution. (Dobzhansky 1973)
- ▶ Discovering how to make such self-reproducing patterns more robust so that they can evolve to increasingly more complex states is probably the central problem in the study of artificial life. (Farmer and d'A. Belin 1992)

Although the cellular automata models of von Neumann and his successors have not, to date, demonstrated extended evolutionary dynamics, various other artificial life models have addressed evolution more directly, through attempts to demonstrate evolution among populations of virtual (software) agents in virtual worlds. The most systematically investigated framework is to envisage these as small machine code programs, each executed by a separate, parallel processor, but all sharing a single main memory (“core”). Generically, these approaches can be referred to as *coreworlds*. In principle, the parallelism could be implemented directly with a sufficiently large pool of hardware processing elements, dynamically allocated as agents are created or destroyed; but it is typically just realized with conventional timeslice-based multithreading of a much smaller pool of hardware CPUs (often just one).

The earliest work adopting this methodological approach can be traced to the seminal investigations of Nils Barricelli, carried out between 1953 and 1963 (Barricelli 1957, 1963). In fact, this work was carried out on the original computer designed and built at the Princeton Institute for Advanced Studies (IAS) between 1946 and 1952 under the direction of von Neumann. Barricelli was a visitor to the institute at von Neumann’s invitation; and took the opportunity to conduct an investigation of the role of “sybiogenesis” in the origin of life. He programmed the IAS computer to directly model patterns of numbers that could interact with each other within a circular array (effectively the main memory system of the machine), according to some fixed local rules. This is conceptually similar to a one-dimensional cellular automaton. Although these patterns of numbers did not have the general computing ability of Turing machines, Barricelli demonstrated the existence of patterns that could self-replicate, as well as various forms of competitive and cooperative (symbiotic) interaction between patterns. Somewhat strangely, in his published description of this work, Barricelli only obliquely related it to von Neumann’s own contemporary work on self-reproducing and evolving automata. In any case, the insight and innovation of this work was not widely recognized at the time, and has only recently been properly reinstated through George Dyson’s investigations of the early history of digital computing at the IAS (Dyson 1997, 1998).

The first attempt to study the dynamics of competition among co-resident *general purpose programs* in a shared memory was pioneered by Vyssotsky et al. at Bell Labs in the early 1960s (though not published until 1972), and appears to have been independent of Barricelli’s study (Vyssotsky 1972). This was effectively created as a form of computer programming game in which different programmers provided hand-coded programs, which were then instantiated in a common core, and executed to see which would survive longest. Again, self-replication was a key feature. An interesting aspect was that, unlike almost all later work, these programs were implemented in the native (IBM 7090) machine code, rather than in a virtual machine code to be executed by an interpreter. This allowed comparatively much higher execution speed, but with some associated limitations. For example, as programs could not be written to be position independent, successful self-replication required active relocation processing, as opposed to simple self-copying. In this early implementation, a “lethal” program was relatively quickly developed by R.H. Morris, and the original game was put aside.

The α -universes described by Holland in the mid-1970s (Holland 1976) may again be mentioned here. Although formulated independently, and with quite different motivations, this proposal was somewhat similar to the Vyssotsky system. In particular, it again envisaged an essentially linear core memory inhabited by concurrently executing fragments of computer code. However, in other ways this was actually much closer to Barricelli's work in that the specific intention was to investigate the spontaneous emergence of crude collective self-reproduction activity, rather than simple competition between programs precoded by human programmers; but, in any case, as already mentioned in [Sect. 2.3](#), this system was not experimentally investigated until much later (McMullin 1992).

The more direct descendent of the Vyssotsky system was the *Core War* game, developed by Dewdney and others in the early 1980s. This now relied on an interpreter and offered much more varied gameplay opportunities (Dewdney 1984). Following the establishment of an international tournament (Dewdney 1987), *Core War* has had a sustained following, and remains active to this day (<http://corewars.org/>).

Although the *Core War* framework is predicated on the idea of human programmers coding the competing programs, it does also naturally lead to the question of whether, under conditions of potential mutation during self-replication, there could be an autonomous and sustained evolutionary process, not relying on programmer intervention. This was explicitly studied by Rasmussen et al. in a system they called the *coreworld* (Rasmussen et al. 1990). In this instantiation, however, the system suffered from a similar limitation to that of the von Neumann style cellular automata systems. While the world could be seeded with an initial self-replicating program, as its offspring filled up the core they quickly began overwriting each other and the self-replicating functionality and thus the potential for sustained Darwinian evolution was lost. This issue was addressed by Ray, in the *Tierra* system, developed and investigated in the early 1990s (Ray 1992). *Tierra* was based on the idea of competing and interacting self-replicating programs in a shared core memory. Programs were now given an ability to employ “memory protection” to limit overwriting by other programs. On its own, this would simply lead to the core filling with programs once, and then no further self-replication (much less evolution) would be possible. To overcome this, Ray added a mechanism for random removal of programs (the so-called “reaper”). This addition allowed for continuing self-replication and the possibility for longer term evolution. He also adopted a significantly different instruction set, incorporating a biologically inspired “template” addressing mode. With these innovations, Ray was able to demonstrate an extended process of evolution, with multiple episodes of Darwinian natural selection. The system produced a surprising array of evolutionary phenomena, including optimization of program length, obligate parasitism, and facultative “hyper-parasitism” (as a defense against obligate parasites).

Tierra has led to a wide variety of related work, such as:

- The *Avida* system of Adami and Brown that introduced a more conventional 2-D spatial world in which programs would colonize nodes, and also be evaluated for their ability to complete extrinsically defined computing tasks (Adami and Brown 1994).
- A proposed extension to *Tierra* into a multiprocessor networked environment ultimately to be distributed across the wider Internet, under suitable “sandbox” execution controls (Ray 1995).
- *Nanopond* (<http://adam.ierymenko.name/nanopond.shtml>), a highly simplified implementation of a 2-D spatial “program evolution” system, drawing on elements of both *Tierra* and *Avida*, but implemented in less than 1,000 lines of C source code.

- The *Amoeba* system, developed by Pargellis, which specifically demonstrated the possibility of spontaneous emergence of a self-replicating program in coreworld-like systems (Pargellis 2001).

Evolution in coreworlds can also usefully be compared and contrasted with several other related but divergent research fields and investigations:

- There has been a separate line of research into the logic of so-called *Quine* programs. These are defined as computer programs that produce their own source code as output. The original discussion is generally credited to Bratley and Millo (1972), though the term “Quine program” is usually attributed to Hofstadter (1979). The existence of Quine programs in general-purpose programming languages essentially follows from Kleene’s recursion theorem. But while the formulation of specific Quine programs in different languages has been a hobby among programmers, these programs have not generally been embedded within an execution environment in which exponential growth, Darwinian selection, and evolution could actually take place. It has not therefore, particularly influenced research approaches in artificial life. That said, through its definitional restriction to programs that must self-reproduce without having direct access to self-inspection, the study of Quine programs draws attention to the logical distinction between proper von Neumann style genetic “self-reproduction” and the “self-replication” by copying normally adopted in coreworlds. It is, in fact, a very open question as to what the effect of this architectural choice is on evolutionary potential (McMullin et al. 2001).
- The replication and propagation of programs in controlled coreworld environments obviously also lead to the idea of self-replicating programs propagating in open, networked, computer environments. This is the realm of computer malware – computer viruses, worms, etc. And indeed, this connection was explicitly made by Dewdney already in the immediate aftermath of the Morris Internet worm incident in 1988 (Dewdney 1989). However, while malware development certainly involves an arms race between the human developers on both sides, there is to date no evidence of effective *autonomous* evolution of “free-living” malware.
- There is significant overlap in inspiration between evolution in coreworld-like systems, and work in evolutionary algorithms, and, especially, *genetic programming* (GP) (Koza 1992; Willis et al. 1997; Banzhaf et al. 1998). This, in turn, has a long history extending back at least to Samuel’s late 1950s investigations in machine learning (Samuel 1959). However, the major distinction is that GP is generally driven, directly or indirectly, by an externally provided evaluation function (used as an extrinsic “fitness” to drive an imposed evolutionary algorithm) rather than the coreworld approach, which is to investigate spontaneous autonomous evolutionary dynamics, in which the software agents are responsible for their own replication and relative fitness emerges from their bottom-up ecological interactions.
- Investigations into the evolution of cooperation by Axelrod and others should also be mentioned here (Axelrod 1984, 1987). Although not generally using a coreworld-like framework, and focused on a relatively narrowly defined pattern of interaction (the iterated prisoner’s dilemma), using an extrinsically applied genetic algorithm, this examination of the problem of cooperation, and coevolution between “cooperating” and “defecting” strategies has been very influential in the wider fields of political science, ecology, and complexity theory.

To summarize, the general experience of the investigation to date of evolution in core-world-like systems is that the evolutionary potential of these systems is interesting but still strictly limited. Indeed, they can be viewed as formally similar to pure artificial “replicator chemistries,” comparable to the systems already mentioned in [Sect. 2.3](#) above, and exhibiting similar phenomena of collective autocatalytic closure and self-maintenance (McMullin 2000b). Thus, both von Neumann-style genetic self-reproduction in cellular automaton worlds, and replication by simple copying or self-inspection in coreworlds, naturally lead directly into the problems of biological robustness, self-maintenance, and hierarchical organization already discussed above. The integration of self-reproducing or self-replicating programs with self-maintenance and individuation (autopoiesis), and the demonstration of sustained evolutionary growth of complexity in a purely virtual world remains perhaps *the* key “grand challenge” problem in the field of artificial life.

3 Theory of and Formalisms in Artificial Life

A number of formalisms have been effective and widely used in artificial life systems. Among them are cellular automata, rewriting systems, and complex dynamical systems. Other formalisms have been used to analyze artificial life systems, like network analysis. In this section, the most important synthetic and analytic tools will be described.

3.1 Automata

The theory of abstract automata defines finite state automata as behavioral models of machines that possess a finite number of states, a set of transitions between those states, and machine actions (such as “read input” or “write output”). Based on this notion, an entire field in computer science (automata theory) has been formulated, which addresses their mathematical features and the use of these entities. One variant of automata of particular relevance to artificial life are cellular automata.

3.1.1 Cellular Automata

A cellular automaton (CA) is composed of a large interconnected collection of component finite state automata. These possess, besides states and transitions, an additional feature, a location within a neighborhood, often constituted as a grid. The neighboring “cells” (although the word “cell” here gives rise to the term “cellular automaton,” note that it is *not* being used with its specifically biological sense; it rather just means a discrete, spatially demarcated, component of a larger, static, locally interconnected, array) or “nodes”, which are themselves finite state automata with a different location, receive the state of the particular automaton as input, and also, reciprocally, provide their own states as inputs that inform a particular automaton which state transition to realize. CA are typically defined to be structurally homogenous, that is all cells contain the same finite state automaton and all share the same local neighborhood pattern. Of course, cell *states* vary dynamically in time and across the cells of the automaton; and it is such spatiotemporal patterns of states that are used to represent so-called *embedded* machines or automata. Such a model can thus be used to model spatially extended systems and study their dynamical behavior. CA have also been used as a modeling

tool to capture behavior otherwise modeled by partial differential equations (PDEs); and indeed can be regarded as a discrete time, discrete space, analogue of PDEs.

Cellular automata play a particularly useful role, in that they possess the twin features of being applicable for synthesis of system behavior, and for analysis of system behavior. They have been thoroughly examined in one-dimensional and two-dimensional contexts (Wolfram 1986, 1994), and have been considered in the fundamental contexts of computability (Wolfram 1984) and parallel computing (Toffoli and Margolus 1987; Hillis 1989).

As mentioned previously (● Sect. 2.4), von Neumann, following a suggestion of Ulam, originally introduced CAs as a basis to formulate his system of self-reproducing and evolvable machines. The von Neumann CA was two dimensional with 29 states per cell. According to Kemeny, von Neumann's core design for a genetically self-reproducing automaton would have occupied a configuration of about 200,000 cells (a size dominated by the "genome" component, which would stretch for a linear distance of about 150,000 cells) (Kemeny 1955, p. 66).

In fact, von Neumann's contributions stand at the cradle of the field of artificial life by showing, for the first time, a mathematical proof that an abstract "machine-like" entity can be conceptualized, which can construct itself, with that copy being able to achieve the same feat in turn, while also having the possibility of undergoing heritable mutations that can support incremental growth in complexity. This was a breakthrough showing that above a certain threshold complexity of an entity, this entity is constructionally powerful enough not only to reproduce itself (and thus maintain that level of complexity) but to seed a process of indefinite complexity growth.

While von Neumann achieved this result in principle, a large number of scientists have since succeeded in pushing down the complexity threshold. In the meantime, much simpler CAs (with smaller number of states per cell) have been shown to support self-reproducing configurations, many of those as 2D loops (Langton 1984a; Sipper 1998). In fact, the smallest currently known example in this form is a cellular automaton with eight states per cell where the self-reproducing configuration is of size five cells (Reggia et al. 1993). However, in many of these cases, the core von Neumann requirement to embody a "general constructive automaton" in the self-reproducing configuration has been relaxed: So the potential for evolutionary growth of complexity is correspondingly impoverished.

Cellular automata also have been used in the design of the Game of Life (Gardner 1970, 1971), which draws inspiration from real life in that replication and transformation of moving entities, again in a 2D CA, were observed (for a comprehensive introduction, see Berlekamp et al. (1982) and Conway's Game of Life, <http://en.wikipedia.org/wiki/conway> and online resources therein).

The discrete nature of states, space, and time available in cellular automata has led to a number of applications of these tools, many of which could be considered artificial life applications in the widest sense. For example, the patterning of seashells has been modeled using discrete embodiments of nonlinear reaction-diffusion systems (Meinhardt 2003). Scientists in Urban Planning have started to use the tool of cellular automata to model urban spread (White and Engelen 1993).

3.2 Rewriting Systems

Rewriting systems are a very general class of formal systems in mathematics, computer science, and logic, in which rules for replacing formal (sub)structures with others are

repeatedly applied. Rewriting systems come in many forms, such as term rewriting, string rewriting, equation rewriting, graph rewriting, among others, and have been at the foundation of programming languages in computer science and automated deduction in mathematics. String rewriting systems, first considered by Axel Thue in the early part of the twentieth century (Dershowitz and Jouannaud 1990; Book and Otto 1993), have a close relationship with Chomsky's formal grammars (Salomaa 1973). Many rewriting systems have the goal of transforming particular expressions into a normalized form, deemed to be stable in some sense. Prior to reaching this stable form from a starting expression, the rewriting process goes through a series of transient expressions. For example, mathematical equations can be formally rewritten until they reach a form in which it is possible to read off solutions.

It can be seen, therefore, that rewriting allows a sort of dynamics of symbolic systems. It is worth noting that rewriting systems provide the basis for the λ -calculus (Barendegt 1984), which has been used to implement artificial chemistries.

In a more general artificial life context, transformations based on rewriting systems can be used to describe growth and developmental processes. For example, besides transforming complex expressions into normal form, another approach uses rewriting to generate complex objects from simple start objects. This has provided the formal basis to studies of biological development through what were later called L-systems (Lindenmayer 1968). The difference between L-systems simulating growth and development and formal grammars lies in the order of rewriting. While formal grammars consider rewriting in a sequential fashion, L-systems consider parallel (simultaneous) rewriting, much closer to the asynchrony of the growth of real cells. But central to all these models is the recursive nature of rewriting systems (Herman et al. 1975). Lindenmayer and, later, Prusinkiewicz and Lindenmayer (1990) have pioneered the use of turtle commands of computer languages like LOGO to visualize plant patterning and growth.

P-systems (Paun 1998; Paun and Rozenberg 2002) are a more complex class of rewriting systems used in artificial life. They are similar to artificial chemistries in that production of symbolic chemicals is prescribed by reaction rules that could also be considered as rewriting rules. In addition to reactions, however, P-systems contain membranes, which limit the availability and direct the flow of material. This provides a higher-level means of structuring environments, something artificial chemistries with well-stirred reactions are unable to do (Dittrich et al. 2001).

In the meantime, a large number of applications of rewriting systems have appeared. For example, such classical artificial life systems as *Tierra* have been reformulated as rewriting systems (Sugiura et al. 2003); Giavitto et al. (2004) describe applications of rewriting systems in modeling biological systems, and virtual cities have been modeled with the help of L-systems (Kato et al. 1998), just to name three examples.

3.3 Dynamical Systems Modeling with ODE

Artificial life systems are often couched in terms of the mathematical language of dynamical systems, captured by ordinary differential equations. This formalism is used for describing the continuous change of state of a system via the change in time of n quantified state variables $\mathbf{q}(t) = \{q_1(t), q_2(t), \dots, q_n(t)\}$. A very simple and linear example of a dynamical system is a system

consisting of two-state variables q_1, q_2 , which develop according to the following differential equation

$$\frac{dq_1}{dt} = a_{11}q_1 + a_{12}q_2 \quad (1)$$

$$\frac{dq_2}{dt} = a_{21}q_1 + a_{22}q_2 \quad (2)$$

Coefficients a_{11}, \dots can be lumped together into a matrix A , and the stability of the dynamical system can be examined based on a solution to the equation $d\mathbf{q}/dt = A \times \mathbf{q}(t) = \mathbf{0}$.

3.3.1 Nonlinear Dynamical Systems

More interesting than this basic dynamical system are those which provide for nonlinear interactions between components and the environment. This can be achieved by higher order terms in the coupling of state variables. For instance

$$\frac{dq_1}{dt} = a_{11}q_1 + a_{12}q_2^2 \quad (3)$$

$$\frac{dq_2}{dt} = a_{21}q_1^2 + a_{22}q_2 \quad (4)$$

would already provide such an example.

Chaotic Systems

Nonlinear couplings are the hallmark of chaotic systems, as nonlinearities in the coupling of state variables tend to amplify small fluctuations (e.g., noise in state variables) into large global state changes. Of course, all these depend on the size of the fluctuations. It is well known that as long as fluctuations are small relative to the strength of nonlinear interactions, systems act in quasi-linear mode and can be formally linearized. However, once either the fluctuation signal is big enough, or the nonlinear interactions become stronger, systems tend to behave in less predictable ways, for example, start oscillating, and might even transit into a chaotic regime characterized by short-term predictability and long-term unpredictability (Argyris et al. 1994). In the artificial life community, the notion of chaos has played an important role as the extreme case of unpredictability beyond life, much as the predictability of dead matter lies at the other end of the spectrum. This notion was first developed in the context of cellular automata (Langton 1992a), but later extended to other systems.

Constructive Systems

Dynamical systems can be extended by considering a notion of change in the number of state variables. For instance, if each of the variables characterizes the concentration of a chemical substance, and some of the substances are consumed or produced for the first time, this can be captured by changing the number and/or assignment of state variables. Simulation might proceed by alternating between a low-level simulation of, for example, individuals of species, and high-level events that signal the birth or extinction of a particular species (Ratze et al. 2007). While couched here in ecological terms, the same approaches apply to the substances of artificial chemistries. Constructive dynamical systems are an active area of research whose exploration has barely begun.

3.4 Complex Systems

The relation between complex systems and artificial life is complicated. Some might argue that artificial life is a subfield of complex systems. This opinion draws from the argument, valid in both complex systems research and artificial life, that these kinds of systems work by producing emergent effects through the interaction of many entities. As such, living systems (and their artificial counterparts) are complex systems, but not the only ones, since, for instance, societies, economies, or legal and norm systems might legitimately be called complex systems too. On the other hand, one might argue that if systems like societies or economies are behaving as complex systems, making use of emergent phenomena and top-down causation, they really can be considered to be alive. If that argument is followed it could be stated that complex systems are part of the research in artificial life.

In this section, it is preferred not to resolve this tension, but rather to focus on one area of research that one legitimately can consider as at the intersection of the two: networks. Networks and their associated formalisms are peculiar since they provide evidence for a phenomenon frequently seen in artificial life systems, as well as in complex systems: There is no easy way to dissect these systems into cause and effect relationships. Indeed, networks allow us to formalize the concept that causes and associated effects have a many-to-many relationship.

If a simple cause–effect relationship is visualized as an edge between two nodes or vertices in a graph, while the nodes themselves symbolize the entities that might or might not be in these relationships, one can easily see that networks, with nodes of in-degree and out-degree larger than one will be able to capture multiple cause–multiple effect relationships. This is a key characteristic of complex systems, as opposed to simple systems, where simple pathways exist between causes and their associated effects.

3.5 Networks

The recent surge in network research (Albert and Barabasi 2002) can be attributed to the realization that networks are the quintessential structure to capture emergent phenomena. Without network infrastructure, emergent behavior of a system would be transient and very quickly disappear, like fluctuations. It is when a network infrastructure is in place that originally weak and transient signals of emergence can be captured and stabilized into identifiable phenomena.

Network studies have revealed widespread characteristics in complex systems, such as small-world features or scale-free connectivity (Watts 1999; Strogatz 2001) that have been ascribed to the generation or evolution of these networks. The dynamics of network evolution is an active area of research (Dorogovtsev and Mendes 2003).

But even the much simpler formalisms of static networks provide useful insights into complex systems and artificial life systems, because there is a deep relationship between autocatalytic sets (Farmer et al. 1986) and closed paths on networks, which in mathematical terms is the result of a connection between graph theory and algebra. Networks can be formally represented by directed graphs $G = G(N, E)$, defined by the set of N nodes and the set of E edges or connections between those nodes. Edges or links between nodes are often represented by an ordered pair of nodes, called the connectivity matrix. If connections are only present or absent, symbolized by a “0” or “1” entry in the matrix, this is called the adjacency matrix A of a graph G . The adjacency matrix has a set of eigenvalues and eigenvectors, which can be easily computed by solving the characteristic equation of A : $|A - \lambda I| = 0$, where λ are

eigenvalues of A and I is the unit matrix. Given that A is a nonnegative matrix, the Perron–Frobenius theorem states that the largest eigenvalue λ_1 of A is real. If this eigenvalue is $\lambda_1 \geq 1$, then there exists a closed path through the graph (Rothblum 1975). Closed paths are deeply related to autocatalytic sets, which were first introduced in chemical systems where each sort of molecules present is catalyzed by at least one other sort of molecules (Eigen 1971). This relation would be considered equivalent to an edge in the graph of all molecular sorts (nodes). Autocatalytic sets correspond therefore to closed paths in that graph.

In the wider context of a theory of organizations (Fontana and Buss 1994; Speroni di Fenizio et al. 2000), autocatalytic sets or closed and self-maintaining sets (organizations) are the key to life-like processes, in that they maintain whatever structure or organization has come into existence. This is the core of the previous statement that emergent phenomena need the stabilizing infrastructure of a network (and its closed paths) to move from the status of a fluctuation in a system to that of an identifiable phenomenon.

The deep connections between algebra and graph theory indicated here would merit a much more detailed discussion, yet space is a constraint here. Suffice it to say that networks are the stage for a dynamics of the chemical (or otherwise) sorts of agents symbolized by the nodes. Thus, it can be said that a dynamical system can be imposed on these networks, and once this dynamics starts to interact with the structure of the network, the most interesting phenomena will be observed (Jain and Krishna 2001).

3.6 Other Formalism

The list of formalisms discussed here is not exhaustive. For instance, the area of statistical mechanics has also contributed valuable insights into artificial life (see, e.g., Adami 1998). Category theory, chemical organisation theory, and other approaches can also be legitimately mentioned as contributing to our understanding of artificial life phenomena.

4 Applications

A huge number of applications can be considered to make use of artificial life techniques in the wider sense (Kim and Cho 2006). It is difficult to even list these applications comprehensively. Instead, a few applications have been selected here to demonstrate the width of the applicability of its concepts.

4.1 Biology

The most obvious application of artificial life is in biology. It was already noted that there is a close relation between artificial life models and synthetic biology (Ray 1994) and that artificial life models in fact are a useful tool for the exploration of questions in biology (Taylor and Jefferson 1994). For a more recent example, see Strand et al. (2002).

4.1.1 Synthetic Biology

The current state of synthetic biology does not acknowledge this connection extensively, despite, historically speaking, emerging from artificial life research. Key questions in this

area today are, how to generate biological entities (cells, proteins, genes, organisms) that have some purposefully designed function (Ferber 2004). The question how systems function that can be ascribed to be “living” has been separated off into the field of “systems biology,” another newcomer to biology with inheritance from artificial life (Kitano 2002). Artificial life was always concerned with “Life as it could be,” that is, alternative designs that can be termed living, designed for a purpose (as engineers would do), or for scientific inquiry, and it was following systems thinking long before systems biology was established (Langton 1997).

4.1.2 Health/Medicine

The modern health-care system is so complex now, that complex systems thinking and ideas from artificial life are ripe for application (Plsek and Greenhalgh 2001). This will entail the understanding of system behavior and response to attempted changes, financial issues such as exploding health-care costs and redesign of the system to become more adaptive.

Medicine has long benefited from its relation to biology, which will result in new applications for artificial life methods as well (Coffey 1998; Hamarneh et al. 2009). Notably, regenerative medicine will be the beneficiary of model systems provided by artificial life (Semple et al. 2005).

4.1.3 Environmental Science and Ecology

In environmental science, the *Gaia* theory (Lovelock 1989) has benefited from simulations provided by artificial life models (Gracias et al. 1997; Downing 2000; Lenton 2002). Artificial life approaches have been used to design sustainable architecture, at least on a local level (Magnoli et al. 2002). In the design of other systems, for example, for recycling, artificial life models have been successfully applied (Okuhara et al. 2003).

Ecosystem research in general has made use of artificial life models, from the setting up of “artificial ecosystems” whose behavior can then be explored in the computer (Lindgren and Nordahl 1994) to the examination of characteristics necessary to provide opportunity for ecosystem formation (Wilkinson 2003). Further, artificial life models have been applied to questions of astrobiology (Centler et al. 2003), another offshoot with some roots in artificial life research.

4.2 Engineering

A number of applications in the engineering subdisciplines have been examined, with autonomous robotics being a very prominent one.

4.2.1 Autonomous Robotics

While it is natural to expect artificial life to contribute to biology, its contribution to the field of engineering is less to be expected. But given the tight connection between behavior and autonomous robotics, it is perhaps less a surprise. In fact, an entire conference series has been

termed “from animals to animats” (Meyer and Wilson 1991; Asada et al. 1998), carrying with it the notion of machines that behave adaptively. Recent reviews and summaries of autonomous robotics as the field pertains to artificial life can be found in Eaton and Collins (2009), Harvey et al. (2005), and Pfeifer et al. (2005).

Often, emergence of behavior or functionality is studied in behaving artifacts like robots, using the bottom-up principles of artificial life. Notably, the interaction of software or hardware parts of robots can be assumed to be closely influenced by artificial life thinking. For instance, the subsumption architecture of Brooks (1990, 1991) has been a closely studied subject in the interaction of behavioral modules. Different sorts of swarm-like behavior (Reynolds 1987; Bonabeau et al. 1999) and, more recently, specifically constructed swarm-bots on the hardware side (Mondada et al. 2004; Groß et al. 2006) have been studied extensively.

Finally, groups of robots cooperating as social robots, and interacting and cooperating with each other, have been studied (Fong et al. 2003). All these applications emphasize the real-world character of robotics problems, and the embodiment of agents (Steels and Brooks 1995).

4.2.2 Transport

Traffic and our societies’ transportation systems are collective systems of an enormous complexity. Their modeling, simulation, and design stand to win substantially from new biological principles (Lucic and Teodorovic 2002; Wang and Tang 2004). Notably, principles gleaned from the organization of insect societies (Deneubourg et al. 1994), and from insect behavior (Theraulaz and Bonabeau 1999), that even give rise to new optimization algorithms (Dorigo and Stützle 2004), are relevant here.

Artificial life applications are not restricted to the movement of physical objects on roads, but are also widespread in other infrastructure networks, such as communication networks (Tanner et al. 2005), notably the Internet (Prokopenko et al. 2005). Early on, it was already proposed to consider computer viruses as a form of artificial life (Dewdney 1989; Spafford 1994) spreading in these communication networks (Bedau 2003).

Already there is a large number of engineering applications for artificial life, but this number is only bound to grow over the coming years (Ronald and Sipper 2000). The reader is referred to (Kim and Cho 2006) for a recent overview of artificial life in engineering.

4.3 Computer Graphics, Virtual Worlds, and Games

Terzopoulos (1999) has summarized a number of applications of artificial life in computer graphics.

4.3.1 Computer Graphics

In particular, Sims’ work is worth mentioning (Sims 1991, 1994). He used virtual organisms that coevolved in competition with each other to demonstrate the power of artificial evolutionary systems for developing naturally looking moving behavior. Terzopoulos built animated systems resembling simulated (“artificial”) fishes, again studying locomotion (Terzopoulos et al. 1994).

The breve visualization environment (Klein 2002) offers a very natural way to interact with artificial life simulations (Spector et al. 2005).

In general, the aim of these tools is animation of agents equipped with virtual sensors and actors, in a life-like fashion (Miranda et al. 2001; Conde and Thalmann 2004). This development was able to close the gap to games (Maes 1995; Reynolds 1999; Rabin 2002).

4.3.2 Virtual Worlds

Magenat-Thalmann and Thalmann (1994) provide an early overview of how artificial life ideas could be put to use in creating virtual worlds. One key idea, also found in other application areas, is to make use of unpredictability in the form of random events, as they are also present in the real world of life through mutation and chance encounters. This line of inquiry is continued with further work on multi-sensor integration in virtual environments, see, for example, Conde and Thalmann (2004). The more complex the inner workings of agents are, and the more complex the environment is providing stimuli, the more intelligence is required to allow for realistic simulations. Aylett et al. discuss the connection between an adaptive environment and the intelligence of the agents/avatars populating it (Aylett and Luck 2000).

More recently, artificial life virtual worlds have been put to use in social scenarios, such as urban environments. For instance, Shao and Terzopoulos (2007) integrate motor, perceptual, behavioral, and cognitive components in a model of pedestrian behavior. Building on Reynolds (1987, 1999), this allows much more realistic simulation of multi-individual interactions in complex urban settings than was possible previously. An entire area, crowd simulation has now emerged using these techniques (Lee et al. 2007).

Virtual worlds are bound to expand enormously in the future, and artificial life applications will continue to proliferate.

4.3.3 Games

Artificial life helped to grow the area of (evolutionary) game theory (Alexander 2003; Gintis 2009) and provided valuable insights into more traditional board games (Pollack et al. 1997). But its most notable and unique success is in areas where traditional games have not been able to contribute much: multiuser games, distributed over the Internet.

One now classical example, and one of the earliest games that made headlines and really was a commercial success for some time, was Steve Grand's game "Creatures" (Grand et al. 1996). The idea behind this game was very simple: Evolution through breeding of creatures using genomes of characteristics that could be combined in multiple ways. By exchanging creatures over the Internet, a multiuser sphere was created setting in motion something akin to real evolution, with implicit fitness, not controlled by any central agency (Grand 2001). Prior to this commercial application, Ray's *Tierra* was distributed over the net in a large experiment (Ray 1998).

4.4 Art

The connection between artificial life and art can be easily seen when considering applications in computer graphics. However, there is a much deeper connection, having to do with

creativity and surprise, which is one of the hallmarks of evolutionary systems that, like artificial life, concern themselves with emergent phenomena. A discussion of these connections is provided in Whitelaw (2004). In a 1998 special section of the journal *Leonardo* (Rinaldo 1998), the connection was examined from various viewpoints. Bentley and Corne (2002) and references therein look at creative processes from an evolutionary perspective, which might be argued to be a superset of artificial life approaches. Arts applications range from painting (Todd and Latham 1992, 1994) (even painting by robots (Moura 2004)) via video installations to sculpture (Baljko and Tenhaaf 2006). Creativity is also considered in the area of sound generation and music (Bilotta and Pantano 2002; Berry and Dahlstedt 2003). Again, artificial life work is often put together with evolutionary work in composition (Romero and Machado 2007).

4.5 Artificial Societies and Artificial Economies

Agent-based modeling using artificial life ideas has also been applied to formulate systems best described as artificial societies and artificial economies.

The area of artificial societies was spawned by the seminal work of Axtell and Epstein (Epstein and Axtell 1996), which used the idea of bottom-up effects in social systems to generate models of social systems that could be simulated in a computer. These ideas have been further developed by Epstein (2007). Social simulation in general (Gilbert and Troitzsch 2005) has now moved to center stage in the social sciences, and a new journal has been established to examine artificial societies (<http://jasss.soc.surrey.ac.uk/JASSS.html>). In the context of social science, emergence (Goldspink and Kay 2007) and learning (Gilbert et al. 2006) have become important topics of discussion.

These very same techniques can also be applied to markets, at which point an artificial society mutates into an artificial economy (Zenobia et al. 2008). The formulation of economies as complex adaptive systems is already somewhat older, and Kaufmann (1993) points out models of economies based on ideas of evolution of complex systems. The power of these ideas lies in their natural ability to treat nonequilibrium phenomena, a topic for the most part carefully avoided by many traditional economists, since it would have to be based on nonlinear systems (Arthur 2006). Recently, another subfield of artificial life, artificial chemistries, have been applied to model economies (Straatman et al. 2008).

5 Conclusion

As we have seen in this chapter, artificial life is a vivid research area, spawning interesting research directions, and providing inspiration for a large number of applications. While the boundaries of the area are fuzzy, there is no doubt that over its 20 years of existence as a named field, artificial life has exerted substantial influence on a whole number of other research areas.

To judge the perspectives of a field, it is perhaps best to focus on its main objectives. The proposal of Bedau (Bedau et al. 2001; Bedau 2003) is followed and slightly changed in dividing the central goals of artificial life into three broad thrusts, all posited here as questions:

- (A) What is the origin of life, or how does life arise from the nonliving?
- (B) What are the potentials and limits of living systems?
- (C) How does life relate to intelligence, culture, society, and human artifacts?

These questions belong to the most challenging and most fascinating in all of science today. From each of these questions derives an entire research programme within artificial life that can be formulated with a non-comprehensive number of tasks:

- (A1) Generate a molecular proto-organism *in vitro*.
- (A2) Achieve the transition to life in an artificial chemistry *in silico*.
- (A3) Determine the threshold of “minimal” life in the organic world *in vivo*.
- (A4) Determine whether fundamentally novel living organizations can arise from inanimate matter.
- (A5) Simulate a unicellular organism over its entire life-cycle.
- (B1) Determine what is inevitable in the open-ended evolution of life.
- (B2) Establish a set of mechanisms for the generation of novelty and innovation.
- (B3) Determine minimal conditions for evolutionary transitions from specific to generic response systems.
- (B4) Create a formal framework for synthesizing dynamical hierarchies at all scales.
- (B5) Determine the predictability of evolutionary manipulations of organisms, ecosystems, and other living systems.
- (B6) Develop a theory of information processing, information flow, and information generation for evolving systems.
- (C1) Explain how rules and symbols are generated from physical dynamics in living systems.
- (C2) Demonstrate the emergence of intelligence and mind in an artificial living system.
- (C3) Evaluate the influence of machines on the next major evolutionary transition of life.
- (C4) Provide predictive models of artificial economies and artificial societies.
- (C5) Provide a quantitative model of the interplay between cultural and biological evolution.
- (C6) Establish ethical principles for artificial life.

Inasmuch as life constitutes a large part of the natural world, artificial life will constitute a large part and exert heavy influence on the area of natural computing discussed in this handbook.

Acknowledgments

WB would like to thank the Canadian Natural Science and Engineering Research Council (NSERC) for providing continuous operating funding since 2003 under current contract RGPIN 283304-07. BMCm acknowledges the support of Dublin City University, and of the European Union FP6 funded project ESIGNET, project number 12789.

References

- Adami C (1998) *Artificial life, an introduction*. Springer, Berlin
- Adami C, Brown CT (1994) Evolutionary learning in the 2D artificial life system “Avida.” In: Brooks RA, Maes P (eds) *Artificial life IV: Proceedings of fourth international workshop on the synthesis and simulation of living systems*. MIT Press, Cambridge, MA, pp 377–381, <http://citeseer.ist.psu.edu/101182.html>
- Albert R, Barabasi AL (2002) Statistical mechanics of complex networks. *Rev Mod Phys* 74:47–97
- Alexander J (2003) Evolutionary game theory. In: *Stanford encyclopedia of philosophy*. <http://plato.stanford.edu/>

- Argyris J, Faust G, Haase M (1994) An exploration of chaos. North-Holland, New York
- Arthur W (2006) Out-of-equilibrium economics and agent-based modeling. In: Handbook of computational economics, vol 2. North-Holland, Amsterdam, The Netherlands pp 1551–1564
- Asada M, Hallam J, Meyer J, Tani J (eds) (2008) From animals to animats: Proceedings of the tenth international conference on simulation of adaptive behavior, Osaka, Japan, July 2008. Springer, Berlin
- Axelrod R (1984) The evolution of cooperation. Basic Books, New York
- Axelrod R (1987) The evolution of strategies in the iterated prisoner's dilemma. In: Davis L (ed) Genetic algorithms and simulating annealing. Pitman, London
- Aylett R, Luck M (2000) Applying artificial intelligence to virtual reality: intelligent virtual environments. *Appl Artif Intell* 14(1):3–32
- Baas N (1994) Hyperstructures. In: Artificial life III, proceedings of the second workshop on artificial life. Addison-Wesley, Redwood City, CA, Santa Fe, NM, June, 1992
- Baljko M, Tenhaaf N (2006) Different experiences, different types of emergence: A-life sculpture designer, interactant, observer. In: Proceedings of AAAI fall 2006 symposium on interaction and emergent phenomenon in societies of agents. Arlington, VA, October 2006, pp 104–110
- Banzhaf W, Nordin P, Keller R, Francone F (1998) Genetic programming – An introduction. Morgan Kaufmann, San Francisco, CA
- Barendegt H (1984) The lambda calculus, its syntax and semantics. Elsevier-North Holland, Amsterdam, The Netherlands
- Barricelli N (1957) Symbiogenetic evolution processes realized by artificial methods. *Methodos* 9(35–36): 143–182
- Barricelli N (1963) Numerical testing of evolution theories. *Acta Biotheor* 16(3):99–126
- Bedau M (2003) Artificial life: organization, adaptation and complexity from the bottom up. *Trends Cogn Sci* 7(11):505–512
- Bedau MA et al. (February 2001) Open problems in artificial life. *Artif Life* 6(4):363–376
- Bentley P, Corne D (eds) (2002) Creative evolutionary systems. Academic Press, San Diego, CA
- Berlekamp ER, Conway JH, Guy RK (1982) What is life? In: Winning ways for your mathematical plays, vol 2. Academic Press, London, chap 25, pp 817–850
- Berry R, Dahlstedt P (2003) Artificial life: Why should musicians bother? *Contemp Music Rev* 22(3):57–67
- Bertalanffy LV (1976) General system theory: foundations, development, applications. George Braziller, New York
- Bilotta E, Pantano P (2002) Synthetic harmonies: Recent results. *Leonardo* 35:35–42
- Bonabeau E, Dorigo M, Theraulaz G (1999) Swarm intelligence: from natural to artificial systems. Oxford University Press, New York
- Book R, Otto F (1993) String-rewriting systems. Springer Verlag, London
- Bratley P, Millo J (1972) Computer recreations; self-reproducing automata. *Software Pract Experience* 2:397–400
- Brooks R (1990) Elephants don't play chess. *Robot Autonomous Syst* 6:3–15
- Brooks R (1991) Intelligence without representation. *Artif Intell* 47:139–159
- Brooks RA (1985) A robust layered control system for a mobile robot. Technical Report A. I. Memo 864, Massachusetts Institution of Technology, Artificial Intelligence Laboratory, <http://people.csail.mit.edu/brooks/papers/AIM-864.pdf>
- Burks AW (ed) (1966) Theory of self-reproducing automata [by] John von Neumann. University of Illinois Press, Urbana, IL
- Cariani P (1992) Some epistemological implications of devices which construct their own sensors and effectors. In: Toward a practice of autonomous systems-proceedings of the first European conference on artificial life Paris, France, December 1991. MIT Press, Cambridge, MA, pp 484–493
- Centler F, Dittrich P, Ku L, Matsumaru N, Pfaffmann J, Zauner K (2003) Artificial life as an aid to astrobiology: testing life seeking techniques. In: Banzhaf W, Christaller T, Dittrich P, Kim J, Ziegler J (eds) Advances in artificial life. Proceedings of ECAL 2003, Dortmund, Germany, September 2003. Lecture notes in computer science, vol 2801. Springer, Berlin, pp 31–40
- Chu D, Ho WK (2006) A category theoretical argument against the possibility of artificial life: Robert Rosen's central proof revisited. *Artif Life* 12(1):117–134. doi: 10.1162/106454606775186392, <http://www.mitpressjournals.org/doi/abs/10.1162/106454606775186392>, <http://www.mitpressjournals.org/doi/pdf/10.1162/106454606775186392>
- Codd EF (1968) Cellular automata. ACM Monograph Series. Academic Press, New York
- Coffey D (1998) Self-organization, complexity and chaos: the new biology for medicine. *Nat Med* 4(8):882–885
- Conde T, Thalmann D (2004) An artificial life environment for autonomous virtual agents with multi-sensorial and multi-perceptive features. *Comput Anim Virtual Worlds* 15:311–318
- Deneubourg J, Clip P, Camazine S (1994) Ants, buses and robots-self-organization of transportation systems. In: Proceedings of the From Perception to Action Conference, PerAc-94, IEEE Press, New York, pp 12–23
- Dershowitz N, Jouannaud J (1990) Rewrite systems. In: Handbook of theoretical computer science.

- Volume B, Formal Methods and Semantics. North-Holland, Amsterdam, the Netherlands, pp 243–320
- Dewdney AK (1984) Computer recreations: in a game called Core War hostile programs engage in a battle of bits. *Sci Am* 250:14–22
- Dewdney AK (1987) Computer recreations: a program called mice nibbles its way to victory at the first Core Wars tournament. *Sci Am* 256(1):8–11
- Dewdney AK (1989) Computer recreations: of worms, viruses and Core War. *Sci Am* 260(3):90–93
- Dittrich P, di Fenizio PS (2007) Chemical organization theory. *Bull Math Biol* 69(4):1199–1231
- Dittrich P, Ziegler J, Banzhaf W (2001) Artificial chemistries – a review. *Artif Life* 7(3):225–275
- Dobzhansky T (1973) Nothing in biology makes sense except in the light of evolution. *Am Biol Teach* 35: 125–129. http://www.pbs.org/wgbh/evolution/library/10/2/text_pop/l_102_591.html
- Dorigo M, Stützle T (2004) Ant colony optimization. MIT Press, Cambridge, MA
- Dorogovtsev S, Mendes J (2003) Evolution of networks: from biological nets to the Internet and WWW. Oxford University Press, Oxford
- Downing K (2000) Exploring Gaia theory: artificial life on a planetary scale. In: Artificial life VII: proceedings of the seventh international conference on artificial life. MIT Press, Cambridge, MA, p 90
- Dyson G (1997) Darwin among the machines; or, the origins of [artificial] life. Presentation hosted by the Edge Foundation, Inc., 8 July 1997, http://edge.org/3rd_culture/dyson/dyson_p1.html
- Dyson GB (1998) Darwin among the machines: the evolution of global intelligence. Basic Books, New York
- Eaton M, Collins J (2009) Artificial life and embodied robotics: current issues and future challenges. *Artif Life Robot* 13(2):406–409
- Edelman G (1987) Neural Darwinism: the theory of neuronal group selection. Basic Books, New York
- Eigen M (1971) Self-organization of matter and the evolution of biological macro-molecules. *Naturwissenschaften* 58:465–523
- Epstein J (2007) Generative social science: studies in agent-based computational modeling. Princeton University Press, Princeton, NJ
- Epstein J, Axtell R (1996) Growing artificial societies: social science from the bottom up. MIT Press, Cambridge, MA
- Farmer J, Kauffman S, Packard N (1986) Autocatalytic replication of polymers. *Physica D* 22:50–67
- Farmer JD, d'A Belin A (1992) Artificial life: the coming evolution. In: Langton CG, Taylor C, Farmer JD, Rasmussen S (eds) Artificial Life II. Santa Fe Institute studies in the sciences of complexity, vol X. Addison-Wesley, Redwood City, CA, pp 815–838. Proceedings of the workshop on artificial life, Santa Fe, NM, February 1990
- Ferber D (2004) Synthetic biology: microbes made to order. *Science* 303:158–161
- Fogel LJ, Owens AJ, Walsh MJ (1966) Artificial intelligence through simulated evolution. Wiley, New York
- Fong T, Nourbakhsh I, Dautenhahn K (2003) A survey of socially interactive robots. *Robot Autonomous Syst* 42(3–4):143–166
- Fontana W, Buss L (1994) The arrival of the fittest: toward a theory of biological organization. *Bull Math Biol* 56:1–64
- Gardner M (1970) Mathematical games: The fantastic combinations of John Conway's new solitaire game "Life." *Sci Am* 223(4):120–123
- Gardner M (1971) Mathematical games: on cellular automata, self-reproduction, the garden of Eden and the game "Life." *Sci Am* 224(2):112–117
- Giavitto J, Malcolm G, Michel O (2004) Rewriting systems and the modeling of biological systems. *Comparative Funct Genomic* 5:95–99
- Gilbert G, Troitzsch K (2005) Simulation for the social scientist. Open University Press, London
- Gilbert N, den Besten M, Bontovics A, Craenen B, Divina F et al. (2006) Emerging artificial societies through learning. *J Artif Soc Soc Simulation* 9(2)
- Gintis H (2009) Game theory evolving: a problem-centered introduction to modeling strategic interaction, 2nd edn. Princeton University Press, Princeton, NJ
- Goldspink C, Kay R (2007) Social emergence: distinguishing reflexive and non-reflexive modes. In: AAAI fall symposium: emergent agents and socialities: social and organizational aspects of intelligence. Washington, DC, November 2007
- Gracias N, Pereira H, Lima J, Rosa A (1997) Gaia: An artificial life environment for ecological systems simulation. In: Artificial life V: proceedings of the fifth international workshop on the synthesis and simulation of living systems. MIT Press, Cambridge, MA, pp 124–134
- Grand S (2001) Creation: life and how to make it. Harvard University Press, Cambridge, MA
- Grand S, Cliff D, Malhotra A (1996) Creatures: artificial life autonomous software agents for home entertainment. In: Proceedings of the first international conference on autonomous agents, Marina del Rey, CA, February 1997. ACM Press, New York
- Graubard SR (ed) (1988) Daedalus, Winter 1988: Artificial intelligence. American Academy of Arts and Sciences, issued as *Proc Am Acad Arts Sci* 117(1)
- Gross D, McMullin B (2001) Is it the right ansatz? *Artif Life* 7(4):355–365
- Groß R, Bonani M, Mondada F, Dorigo M (2006) Autonomous self-assembly in swarm-bots. *IEEE Trans Robot* 22(6):1115–1130
- Hamarneh G, McIntosh C, McInerney T, Terzopoulos D (2009) Deformable organisms: an artificial life

- framework for automated medical image analysis. Chapman & Hall/CRC, Boca Raton, FL, p 433
- Harvey I, Paolo E, Wood R, Quinn M, Tuci E (2005) Evolutionary robotics: a new scientific tool for studying cognition. *Artif Life* 11(1–2):79–98
- Haugeland J (1989) Artificial intelligence: the very idea. MIT Press, Cambridge, MA
- Herman G, Lindenmayer A, Rozenberg G (1975) Description of developmental languages using recurrence systems. *Math Syst Theory* 8:316–341
- Hillis W (1989) The connection machine. MIT Press, Cambridge, MA
- Hofstadter DR (1979) Gödel, Escher, Bach: An eternal golden braid. Penguin Books, London. First published in Great Britain by The Harvester Press Ltd 1979. Published in Penguin Books 1980
- Holland JH (1975) Adaptation in natural and artificial systems. The University of Michigan Press, Ann Arbor MI
- Holland JH (1976) Studies of the spontaneous emergence of self-replicating systems using cellular automata and formal grammars. In: Lindenmayer A, Rozenberg G (eds) Automata, languages, development. North-Holland, New York, pp 385–404. Proceedings of a conference held in Noordwijkerhout, Holland, 31 March–6 April 1975
- Jain S, Krishna S (2001) A model for the emergence of cooperation, interdependence and structure in evolving networks. *PNAS* 98:543–547
- Kampis G (1991) Self-modifying systems in biology and cognitive science. IFSR international series on systems science and engineering, vol 6. Pergamon Press, Oxford, NY. Editor-in-Chief: George J. Klir
- Kato N, Okuno T, Okano A, Kanoh H, Nishihara S (1998) An ALife approach to modelling virtual cities. In: IEEE Conference on Systems, Man and Cybernetics, vol 2, San Diego, CA, October 1998. IEEE Press, New York, vol 2, pp 1168–1173
- Kauffman SA (1993) The origins of order: self-organization and selection in evolution. Oxford University Press, Oxford
- Kemeny JG (1955) Man viewed as a machine. *Sci Am* 192(4):58–67
- Kim K, Cho S (2006) A comprehensive overview of the applications of artificial life. *Artif Life* 12(1): 153–182
- Kitano H (2002) Systems biology: a brief overview. *Science* 295:1662–1664
- Klein J (2002) Breve: a 3D environment for the simulation of decentralized systems and artificial life. In: Proceedings of artificial life VIII, the eighth international conference on the simulation and synthesis of living systems, Sydney, Australia, December 2002. MIT Press, Cambridge, MA, pp 329–334
- Koza J (1992) Genetic programming. MIT Press, Cambridge, MA
- Kvasnicka V, Pospíchal J, Kaláb T (2001) A study of replicators and hypercycles by typogenetics. In: ECAL 2001: Proceedings of the sixth European conference on advances in artificial life, Prague, Czech Republic, September 2001. Springer-Verlag, Berlin, pp 37–54
- Langton CG (1984a) Self-reproduction in cellular automata. *Physica D* 10:135–144
- Langton CG (1992a) Life at the edge of chaos. In: Langton C, Taylor C, Farmer J, Rasmussen S (eds) Artificial life II: Proceedings of the second workshop on artificial life. Addison-Wesley, Redwood, CA
- Langton CG (1995) What is artificial life? In: Brown T (ed) comp.ai.alife Frequently Asked Questions (FAQ), November 18th 1995. <http://www.faqs.org/faqs/ai-faq/alife/>
- Langton CG (1997) Artificial life: an overview. MIT Press, Cambridge, MA
- Langton CG (1984b) Self-reproduction in cellular automata. *Physica D* 10:135–144
- Langton CG (1986) Studying artificial life with cellular automata. *Physica* 22D:120–149
- Langton CG (1987) Announcement of artificial life workshop. http://groups.google.com/group/news.announce.conferences/browse_thread/thread/776a2af47daf7a73/7ab70ed772c50437, posted to news.announce.conferences, April 20 1987, 7:22 pm
- Langton CG (ed) (1989) Artificial Life. Sante Fe Institute studies in the sciences of complexity, vol VI. Addison-Wesley, Redwood City, CA. Proceedings of an interdisciplinary workshop on the synthesis and simulation of living systems, Los Alamos, NM, September 1987
- Langton CG (1992b) Preface. In: Langton CG, Taylor C, Farmer JD, Rasmussen S (eds) Artificial life II. Santa Fe Institute Studies in the Sciences of Complexity, vol X. Addison-Wesley, Redwood City, CA. Proceedings of the workshop on artificial life, Santa Fe, NM, February 1990
- Langton CG, Taylor C, Farmer JD, Rasmussen S (eds) (1992) Artificial life II. Sante Fe Institute Studies in the sciences of complexity, vol X. Addison-Wesley, Redwood City, CA. Proceedings of the workshop on artificial life Santa Fe, NM, February 1990
- Lee K, Choi M, Hong Q, Lee J (2007) Group behavior from video: a data-driven approach to crowd simulation. In: Proceedings of the 2007 ACM SIGGRAPH/Eurographics symposium on computer animation. San Diego, CA, August 2007, pp 109–118
- Lenton T (2002) Gaia as a complex adaptive system. *Philos Trans R Soc B Biol Sci* 357(1421):683–695
- Levy S (1993) Artificial life: a report from the frontier where computers meet biology. Vintage, New York
- Lindenmayer A (1968) Mathematical models for cellular interaction in development, parts I and II. *J Theor Biol* 18:280–315

- Lindenmayer A, Rozenberg G (eds) (1976) Automata, languages, development. Elsevier, Amsterdam, The Netherlands
- Lindgren K, Nordahl M (1994) Cooperation and community structure in artificial ecosystems. *Artif Life* 1:15–37
- Lovelock J (1989) Geophysiology, the science of Gaia. *Rev Geophys* 27(2):215–222
- Lucic P, Teodorovic D (2002) Transportation modeling: an artificial life approach. In: ICTAI 2002: Proceedings of the 14th IEEE international conference on tools with artificial intelligence. Washington, DC, November 2002, pp 216–223
- Luisi PL, Varela F (1989) Self replicating micelles: a minimal version of a chemical autopoietic system. *Orig Life Evol Biosph* 19:633–643
- Maes P (1995) Artificial life meets entertainment: lifelike autonomous agents. *Commun ACM* 38(11):108–114
- Magenat-Thalmann N, Thalmann D (eds) (1994) Artificial life and virtual reality. Wiley, Chichester, UK
- Magnoli G, Bonanni L, Khalaf R (2002) Designing a DNA for adaptive architecture: a new built environment for social sustainability. In: Brebbia C, Sucharov L, Pascolo P (eds) Design and nature: comparing design in nature with science and engineering. WIT, Southampton, UK, pp 203–213
- Maturana HR, Varela FJ (1973) Autopoiesis: the organization of the living. In: Autopoiesis and cognition: the realization of the living. Boston studies in the philosophy of science, vol 42. D. Reidel Publishing Company, Dordrecht, Holland, pp 59–138. First published 1972 in Chile under the title *De Maquinas y Seres Vivos*, Editorial Universitaria S.A
- Maynard Smith J, Szathmáry E (1997) The major transitions in evolution. Oxford University Press, New York
- McMullin B (1992) The Holland α -universes revisited. In: Varela FJ, Bourgine P (eds) Toward a practice of autonomous systems: proceedings of the first European conference on artificial life. Complex Adaptive Systems. MIT Press, Cambridge, MA, pp 317–326. series Advisors: John H. Holland, Christopher Langton and Stewart W. Wilson
- McMullin B (2000a) John von Neumann and the evolutionary growth of complexity: looking backwards, looking forwards. *Artif Life* 6(4):347–361. <http://www.eeng.dcu.ie/~alife/bmcm-alj-2000/>
- McMullin B (2000b) Some remarks on autocatalysis and autopoiesis. *Ann N Y Acad Sci* 901:163–174. <http://www.eeng.dcu.ie/~alife/bmcm9901/>
- McMullin B (2004) 30 years of computational autopoiesis: a review. *Artif Life* 10(3):277–296. <http://www.eeng.dcu.ie/~alife/bmcm-alj-2004/>
- McMullin B, Varela FJ (1997) Rediscovering computational autopoiesis. In: Husbands P, Harvey I (eds) ECAL-97: Proceedings of the fourth European conference on artificial life, Brighton, UK, July 1997. Complex adaptive systems. MIT Press, Cambridge, MA. <http://www.eeng.dcu.ie/~alife/bmcm-ecal97/>
- McMullin B, Taylor T, von Kamp A (2001) Who needs genomes? In: Atlantic symposium on computational biology and genome information systems & technology, Regal University Center, Durham, NC, March 2001. <http://www.eeng.dcu.ie/~alife/bmcm-cbgi-2001/>
- Meinhardt H (2003) The algorithmic beauty of sea shells. Springer, Heidelberg
- Meyer J, Wilson S (eds) (1991) From animals to animats: Proceedings of the first international conference on simulation of adaptive behavior, London, 1990. MIT Press Cambridge, MA
- Miranda F, Kögler J, Del Moral Hernandez E, Lobo Netto M (2001) An artificial life approach for the animation of cognitive characters. *Comput Graph* 25(6):955–964
- Mondada F, Pettinaro G, Guignard A, Kwee I, Floreano D, Deneubourg J, Nolfi S, Gambardella L, Dorigo M (2004) SWARM-BOT: a new distributed robotic concept. *Autonomous Robot* 17(2):193–221
- Moura L (2004) Man + robots: symbiotic art. Institut d'Art Contemporain, Villeurbanne, France
- Okuhara K, Domoto E, Ueno N, Fujita H (2003) Recycling design using the artificial life technology to optimize evaluation function. In: ECO Design'03: 2003 third international symposium on environmentally conscious design and inverse manufacturing. Tokyo, Japan, December 2003, pp 662–665
- Pargellis AN (2001) Digital life behavior in the amoeba world. *Artif Life* 7(1):63–75. doi: 10.1162/106454601300328025, <http://www.mitpressjournals.org/doi/abs/10.1162/106454601300328025>, <http://www.mitpressjournals.org/doi/pdf/10.1162/106454601300328025>
- Pattee H (ed) (1973) Hierarchy theory. The challenge of complex systems. Georges Braziller, New York
- Paun G (1998) Computing with membranes. Technical Report, Turku Center for Computer Science, TUCS Report 208
- Paun G, Rozenberg G (2002) A guide to membrane computing. *Theor Comput Sci* 287:73–100
- Pesavento U (1995) An implementation of von Neumann's self-reproducing machine. *Artif Life* 2(4):337–354
- Pfeifer R, Iida F, Bongard J (2005) New robotics: design principles for intelligent systems. *Artif Life* 11(1–2):99–120
- Plsek P, Greenhalgh T (2001) The challenge of complexity in health care. *Br Med J* 323(7313):625–628
- Pollack J, Blair A, Land M (1997) Coevolution of a backgammon player. In: Artificial life V: Proceedings of the fifth international workshop on the synthesis and simulation of living systems. Nara, Japan, May 1996, pp 92–98

- Popper KR, Eccles JC (1977) *The self and its brain: an argument for interactionism*. Routledge & Kegan Paul plc, London. First published 1977, Springer-Verlag, Berlin. This edition first published 1983
- Prokopenko M, Wang P, Valencia P, Price D, Foreman M, Farmer A (2005) Self-organizing hierarchies in sensor and communication networks. *Artif Life* 11 (4):407–426
- Prusinkiewicz P, Lindenmayer A (1990) *The algorithmic beauty of plants*. Springer, New York
- Rabin S (2002) *AI game programming wisdom*. Charles River Media, Hingham, MA
- Rashevsky N (1940) *Advances and applications of mathematical biology*. University of Chicago Press, Chicago, IL
- Rasmussen S (2008) *Protocells*. MIT Press, Cambridge, MA
- Rasmussen S, Knudsen C, Feldberg R, Hindsholm M (1990) *The Coreworld: emergence and evolution of cooperative structures in a computational chemistry*. *Physica* 42D:111–134
- Rasmussen S, Baas NA, Mayer B, Nilsson M (2001a) Defense of the ansatz for dynamical hierarchies. *Artif Life* 7(4):367–373
- Rasmussen S, Baas NA, Mayer B, Nilsson M, Olesen MW (2001b) Ansatz for dynamical hierarchies. *Artif Life* 7(4):329–353
- Rasmussen S, Bedau MA, Chen L, Deamer D, Krakauer DC, Packard NH, Stadler PF (eds) (2008) *Protocells: bridging nonliving and living matter*. MIT Press, Cambridge, MA
- Ratze C, Gillet F, Mueller J, Stoffel K (2007) Simulation modelling of ecological hierarchies in constructive dynamical systems. *Ecol Complexity* 4:13–25
- Ray T (1994) An evolutionary approach to synthetic biology. *Artif Life* 1(1/2):179–209
- Ray T (1998) Selecting naturally for differentiation: preliminary evolutionary results. *Complexity* 3 (5):25–33
- Ray TS (1992) An approach to the synthesis of life. In: Langton CG, Taylor C, Farmer JD, Rasmussen S (eds) *Artificial life II*. Santa Fe Institute studies in the sciences of complexity, vol X. Addison-Wesley, Redwood City, CA, pp 371–408
- Ray TS (1995) A proposal to create two biodiversity reserves: one digital and one organic. <http://life.ou.edu/pubs/reserves/>, Project proposal for Network Tierra
- Rechenberg I (1975) *Evolution strategies* (in German). Holzmann-Froboog, Stuttgart, Germany
- Reeke GN Jr, Edelman GM (1988) *Daedalus*, Winter 1998: Real brains and artificial intelligence. *Proc Am Acad Arts Sci* 117(1):143–173
- Reggia J, Armentrout S, Chou H, Peng Y (1993) Simple systems that exhibit self-directed replication. *Science* 259:1282–1287
- Regis E (2008) *What is life? Investigating the nature of life in the age of synthetic biology*. Farrar, Straus & Giroux, New York
- Reynolds C (1987) *Flocks, herds and schools: a distributed behavioral model*. In: *Proceedings of the 14th annual conference on computer graphics and interactive techniques*, July 1987. ACM, New York, Anaheim, CA, July 1987, pp 25–34
- Reynolds C (1999) Steering behaviors for autonomous characters. In: *Game developers conference*, vol 1999, San Jose, CA, 1999, pp 763–782
- Rinaldo K (1998) *Artificial life art*, introduction to the special section of Leonardo. *Leonardo* 31:370
- Rocha LM (ed) (2000) *The physics and evolution of symbols and codes: reflections on the work of Howard Pattee*. *BioSystems* (special issue) 60(1–3) 149–157. Elsevier
- Romero J, Machado P (2007) *The art of artificial evolution: a handbook of evolutionary art and music*. Springer, Berlin
- Ronald E, Sipper M (2000) *Engineering, emergent engineering, and artificial life: unsurprising, unsurprising surprise, and surprising surprise*. In: *Artificial life VII: Proceedings of the seventh international conference on artificial life*. The MIT Press, Cambridge, MA, pp 523–528
- Rosen R (1959) On a logical paradox implicit in the notion of a self-reproducing automaton. *Bull Math Biophys* 21:387–394
- Rosen R (1972) Some relational cell models: the metabolism-repair systems. In: Rosen R (ed) *Foundations of mathematical biology*, vol II. Academic, New York
- Rosen R (1985) *Organisms as causal systems which are not mechanisms: an essay into the nature of complexity*. In: Rosen R (ed) *Theoretical biology and complexity*. Academic, Orlando, FL, Chap 3, pp 165–203
- Rosen R (1991) *Life itself*. Columbia University Press, New York
- Rothblum U (1975) Algebraic eigenspace of nonnegative matrices. *Linear Algebra Appl* 12:281–292
- Rumelhart DE, McClelland JL, the PDP Research Group (1986) *Parallel distributed processing*, vol 1: Foundations. MIT Press, Cambridge, MA
- Salomaa A (1973) *Formal languages*. Academic Press, New York
- Samuel A (July 1959) Some studies in machine learning using the game of checkers. *IBM J* 3(3): 210–229, http://ieeexplore.ieee.org/xpl/freeabs_all.jsp?arnumber=5392560
- Searle JR (1980) *Minds, brains, and programs*. *Behav Brain Sci* 3:417–457, includes peer commentaries. Also reprinted (without commentaries) as Hauge-land J (ed) (1981) *Mild design*. MIT Press, Cambridge, MA, Chap 10, pp 282–305 and Boden MA (ed) (1990) *The philosophy of artificial intelligence*.

- Oxford readings in philosophy. Oxford University Press, Oxford, Chap 3, pp 67–88
- Simple J, Woolridge N, Lumsden C (2005) Review: in vitro, in vivo, in silico: computational systems in tissue engineering and regenerative medicine. *Tissue Eng* 11(3–4):341–356
- Shao W, Terzopoulos D (2007) Autonomous pedestrians. *Graph Model* 69(5–6):246–274
- Sims K (1991) Artificial evolution for computer graphics. *Comput Graph* 25(4):319–328
- Sims K (1994) Evolving 3D morphology and behavior by competition. In: *Artificial life IV: Proceedings of the fourth international workshop on the synthesis and simulation of living systems*. MIT Press, Cambridge, MA, pp 28–39
- Sipper M (1998) Fifty years of research on self-replication: an overview. *Artif Life* 4(3):237–257
- Smith HO, Hutchison CA, Pfannkoch C, Venter JC (2003) Generating a synthetic genome by whole genome assembly: ϕ x174 bacteriophage from synthetic oligonucleotides. *Proc Natl Acad Sci* 100(26):15440–15445
- Snare A (1999) Typogenetics. <http://www.csse.monash.edu.au/hons/projects/1999/Andrew.Snare/thesis.pdf>, Thesis for Bachelor of Science (Computer Science) Honours, School of Computer Science and Software Engineering, Monash University
- Spafford E (1994) Computer viruses as artif life. *Artif Life* 1(3):249–265
- Spector L, Klein J, Perry C, Feinstein M (2005) Emergence of collective behavior in evolving populations of flying agents. *Genet Programming Evolvable Mach* 6(1):111–125
- Speroni di Fenizio P, Dittrich P, Zeigler J, Banzhaf W (2000) Towards a theory of organizations. In: *German workshop on artificial life*, vol 5, Bayreuth, Germany, April 2000. <http://www.cs.mun.ca/~banzhaf>
- Steels L, Brooks R (1995) *The artificial life route to artificial intelligence*. Lawrence Erlbaum Associates, New Haven, CT
- Straatman B, White R, Banzhaf W (2008) An artificial chemistry-based model of economies. In: *Artificial Life XI: Proceedings of the eleventh international conference on simulation and synthesis of living systems*. MIT Press, Cambridge, MA, pp 592–602
- Strand E, Huse G, Giske J (2002) Artificial evolution of life history and behaviour. *Am Nat* 159:624–644
- Strogatz S (2001) Exploring complex networks. *Nature* 410:268–276
- Sugiura K, Suzuki H, Shiose T, Kawakami H, Katai O (2003) Evolution of rewriting rule sets using string-based tierra. In: *Banzhaf W et al. (eds) Advances in artificial life - Proceedings ECAL 2003*, Budapest, Hungary, April 2003. Lecture notes in computer science, vol 2801. Springer, Berlin, pp 69–77
- Tanner H, Jadbabaie A, Pappas G (2005) Flocking in fixed and switching networks. In: *IEEE conference on decision and control*, vol 1, Seville, Spain, December 2005. p 2
- Taylor C, Jefferson D (1994) Artificial life as a tool for biological inquiry. *Artif Life* 1(1–2):1–13
- Terzopoulos D (1999) Artificial life for computer graphics. *Comm of the ACM* 42(8):32–42
- Terzopoulos D, Tu X, Grzeszczuk R (1994) Artificial fishes with autonomous locomotion, perception, behavior, and learning in a simulated physical world. In: *Artificial life IV: Proceedings of the fourth international workshop on the synthesis and simulation of living systems*. MIT Press, Cambridge, MA, pp 17–27
- Thatcher JW (1970) Universality in the von Neumann cellular model. In: *Burks AW (ed) Essays on cellular automata*. University of Illinois Press, Urbana, IL, pp 132–186 (Essay Five)
- Theraulaz G, Bonabeau E (1999) A brief history of stigmergy. *Artif Life* 5(2):97–116
- Todd S, Latham W (1992) Artificial life or surreal art. In: *Toward a practice of autonomous systems: Proceedings of the first European conference on artificial life*. MIT Press, Cambridge, MA, p 504
- Todd S, Latham W (1994) *Evolutionary art and computers*. Academic London
- Toffoli T, Margolus N (1987) *Cellular automata machines*. MIT Press, Cambridge, MA
- Turing A (1936) On computable numbers, with an application to the Entscheidungsproblem. *Proc London Math Soc Ser 2* 42:230–265
- Turing AM (1950) Computing machinery and intelligence. *Mind* LIX(236):433–460, also reprinted as, *Boden MA (ed) (1990) The philosophy of artificial intelligence*. Oxford readings in philosophy. Oxford University Press, Oxford, Chap 2, pp 40–66
- Turing AM (1952) The chemical basis of morphogenesis. *Philos Trans R Soc B* 237:37–72
- Varela FJ, Maturana HR, Uribe R (1974) Autopoiesis: the organization of living systems, its characterization and a model. *BioSystems* 5:187–196
- Varela FJ, Thompson ET, Rosch E (1992) *The embodied mind: cognitive science and human experience*. MIT Press, Cambridge, MA
- von Neumann J (1945) First draft of a report on the EDVAC. (A corrected version was formally published in the *IEEE Ann Hist Comput* 15(4):27–75 1993.)
- von Neumann J (1949) Theory and organization of complicated automata. In: *Burks AW (ed) (1966) Theory of self-reproducing automata [by] John von Neumann*. University of Illinois Press, Urbana, IL, pp 27– (Part One), based on transcripts of lectures delivered at the University of Illinois, in December 1949. Edited for publication by A.W. Burks

- von Neumann J (1951) The general and logical theory of automata. In: Taub AH (ed) John von Neumann: collected works. Volume V: Design of computers, theory of automata and numerical analysis. Pergamon Press, Oxford, chap 9, pp 288–328. First published Jeffress LAA (ed) (1951) Cerebral mechanisms in behavior—the Hixon symposium. Wiley, New York, pp 1–41
- von Neumann J (1966) The theory of automata: Construction, reproduction, homogeneity. In: Burks AW (ed) (1966) Theory of self reproducing automata [by] John von Neumann. University of Illinois Press, Urbana, IL, pp 89–250. Based on an unfinished manuscript by von Neumann. Edited for publication by A.W. Burks
- Vyssotsky VA (1972) Darwin: a game of survival and (hopefully) evolution. Software Pract Experience 2: 91–96. <http://www.cs.dartmouth.edu/doug/darwin.pdf>, Attachment to letter appearing in Computer Recreations column, signed by M. D. McIlroy, R. Morris, and V. A. Vyssotsky
- Waddington CH (ed) (1969) Towards a theoretical biology. 2: Sketches. Edinburgh University Press, Edinburgh, UK
- Wang F, Tang S (2004) Artificial societies for integrated and sustainable development of metropolitan systems. IEEE Intell Syst 19(4):82–87
- Watts D (1999) Small worlds: the dynamics of networks between order and randomness. Princeton University Press, Princeton, NJ
- White R, Engelen G (1993) Cellular automata and fractal urban form: a cellular modelling approach to the evolution of urban land-use patterns. Environ Plan A 25:1175–1199
- Whitelaw M (2004) Metacreation: art and artificial life. MIT Press, Cambridge, MA
- Wiener N (1965) Cybernetics, 2nd edn: or the control and communication in the animal and the machine. MIT Press, Cambridge, MA
- Wilkinson D (2003) The fundamental processes in ecology: a thought experiment on extraterrestrial biospheres. Biol Rev 78:171–179
- Willis M, Hiden H, Marenbach P, McKay B, Montague GA (1997) Genetic programming: an introduction and survey of applications. In: Zalzala A (ed) GALE-SIA'97: Second international conference on genetic algorithms in engineering systems: innovations and applications, Glasgow, UK, September 1997. Institution of Electrical Engineers (IEE), Savoy Place, London, UK, pp 314–319, <http://ieeexplore.ieee.org/stamp/stamp.jsp?arnumber=00681044>
- Wolfram S (1984) Computation theory of cellular automata. Commun Math Phys 96(1):15–57
- Wolfram S (1986) Theory and applications of cellular automata. Advanced Series on Complex Systems. World Scientific Publication, Singapore
- Wolfram S (1994) Cellular automata and complexity: collected papers. Perseus Books Group, Reading, MA
- Wolkenhauer O (2007) Why systems biology is (not) systems biology. BIOforum Europe 2007, vol 4, pp 2–3
- Yovits MC, Cameron S (eds) (1960) Self-organizing systems. Pergamon Press, Oxford. Proceedings of an interdisciplinary conference, 5 and 6 May, 1959
- Zenobia B, Weber C, Daim T (2008) Artificial markets: a review and assessment of a new venue for innovation research. Technovation. doi:10.1016/j.technovation.2008.09.002