

In The Philosophy of Artificial Life, ed. M. Boden
Oxford 1996.

1

ARTIFICIAL LIFE

CHRISTOPHER G. LANGTON

1. THE BIOLOGY OF POSSIBLE LIFE

Biology is the scientific study of life—in principle anyway. In practice, biology is the scientific study of life on Earth based on carbon-chain chemistry. There is nothing in its charter that restricts biology to the study of carbon-based life; it is simply that this is the only kind of life that has been available to study. Thus, theoretical biology has long faced the fundamental obstacle that it is impossible to derive general principles from single examples.

Without other examples it is extremely difficult to distinguish essential properties of life—properties that must be shared by any living system *in principle*—from properties that may be incidental to life, but which happen to be universal to life on Earth due *solely* to a combination of local historical accident and common genetic descent. Since it is quite unlikely that organisms based on different physical chemistries will present themselves to us for study in the foreseeable future, our only alternative is to try to synthesize alternative life-forms ourselves—*Artificial Life*: life made by man rather than by nature.

1.1. Artificial Life

Biology has traditionally started at the top, viewing a living organism as a complex biochemical machine, and has worked *analytically* down from there through the hierarchy of biological organization—decomposing a living organism into organs, tissues, cells, organelles, and finally molecules—in its pursuit of the mechanisms of life. Analysis means ‘the separation of an intellectual or substantial whole into constituents for individual study’. By composing our individual understandings of the dissected component parts of living organisms,

traditional biology has provided us with a broad picture of the mechanics of life on Earth.

But there is more to life than mechanics—there is also dynamics. Life depends critically on principles of dynamical self-organization that have remained largely untouched by traditional analytic methods. There is a simple explanation for this—these self-organizing dynamics are fundamentally non-linear phenomena, and non-linear phenomena in general depend critically on the interactions *between* parts: they necessarily disappear when parts are treated in isolation from one another, which is the basis for the analytic method.

Rather, non-linear phenomena are most appropriately treated by a *synthetic* approach. Synthesis means 'the combining of separate elements or substances to form a coherent whole'. In non-linear systems, the parts must be treated in each other's presence, rather than independently from one another, because they behave very differently in each other's presence than we would expect from a study of the parts in isolation.

Artificial Life is simply the synthetic approach to biology: rather than take living things apart, Artificial Life attempts to put living things together.

But Artificial Life is more than this. To understand the overall aims of the Artificial Life enterprise, one needs to do the following: (1) Broaden the scope of the attempts, beyond simply recreating 'the living state', to the synthesis of any and all biological phenomena, from viral self-assembly to the evolution of the entire biosphere. (2) Couple this with the observation that there is no reason, in principle, why the parts we use in our attempts to synthesize these biological phenomena need be restricted to carbon-chain chemistry. (3) Note that we expect the synthetic approach to lead us not only to, but quite often *beyond*, known biological phenomena: beyond *life-as-we-know-it* into the realm of *life-as-it-could-be*.

Thus, for example, Artificial Life involves attempts to (1) synthesize the process of evolution (2) in computers, and (3) will be interested in whatever emerges from the process, even if the results have no analogues in the 'natural' world. It is certainly of scientific interest to know what kinds of things *can* evolve in principle, whether or not they happened to do so here on Earth.

1.2. AI and the Behaviour Generation Problem

Artificial Life is concerned with generating *lifelike* behaviour. Thus it focuses on the problem of creating *behaviour generators*. A good place to start is to identify the mechanisms by which behaviour is generated and controlled in natural systems, and to recreate these mechanisms in artificial systems. This is the course we shall take later in this paper.

The related field of Artificial Intelligence is concerned with generating *intelligent* behaviour. It, too, focuses on the problem of creating behaviour generators. However, although it initially looked to natural intelligence to identify its underlying mechanisms, these mechanisms were not known, nor are they today. Therefore, following an initial flirt with neural nets, AI became wedded to the only other known vehicle for the generation of complex behaviour: the technology of serial computer programming. As a consequence, from the very beginning artificial intelligence embraced an underlying methodology for the generation of intelligent behaviour that bore no demonstrable relationship to the method by which intelligence is generated in natural systems. In fact, AI has focused primarily on the production of intelligent *solutions* rather than on the production of intelligent *behaviour*. There is a world of difference between these two possible foci.

By contrast, Artificial Life has the great good fortune that many of the mechanisms by which behaviour arises in natural living systems are known. There are still many holes in our knowledge, but the general picture is in place. Therefore, Artificial Life can start by recapturing natural life and has no need to resort to the sort of initial infidelity that is now coming back to haunt AI.

The key insight into the natural method of behaviour generation is gained by noting that *nature is fundamentally parallel*. This is reflected in the 'architecture' of natural living organisms, which consist of many millions of parts, each one of which has its own behavioural repertoire. Living systems are highly distributed and quite massively parallel. If our models are to be true to life, they must also be highly distributed and quite massively parallel. Indeed, it is unlikely that any other approach will prove viable.

2. HISTORICAL ROOTS OF ARTIFICIAL LIFE

Mankind has a long history of attempting to map the mechanics of his contemporary technology on to the workings of nature, trying to understand the latter in terms of the former.

It is not surprising, therefore, that early models of life reflected the principal technology of their era. The earliest models were simple statuettes and paintings—works of art which captured the static form of living things. These statues were provided with articulated arms and legs in the attempt to capture the dynamic form of living things. These simple statues incorporated no internal dynamics, requiring human operators to make them behave.

The earliest mechanical devices that were capable of generating their own behaviour were based on the technology of water transport. These were the early Egyptian water-clocks called *Clepsydra*. These devices made use of a

rate-limited process—in this case the dripping of water through a fixed orifice—to indicate the progression of another process—the position of the sun. Ctesibius of Alexandria developed a water-powered mechanical clock around 135 BC which employed a great deal of the available hydraulic technology—including floats, a siphon, and a water-wheel-driven train of gears.

In the first century AD, Hero of Alexandria produced a treatise on *Pneumatics*, which described, among other things, various simple gadgets in the shape of animals and humans that utilized pneumatic principles to generate simple movements.

However, it was really not until the age of mechanical clocks that artefacts exhibiting complicated internal dynamics became possible. Around AD 850 the *mechanical escapement* was invented, which could be used to regulate the power provided by falling weights. This invention ushered in the great age of clockwork technology. Throughout the Middle Ages and the Renaissance, the history of technology is largely bound up with the technology of clocks. Clocks often constituted the most complicated and advanced application of the technology of an era.

Perhaps the earliest clockwork simulations of life were the so-called 'Jacks', mechanical 'men' incorporated in early clocks who would swing a hammer to strike the hour on a bell. The word 'jack' is derived from 'jacomarchiadus', which means 'the man in the suit of armour'. These accessory figures retained their popularity even after the spread of clock dials and hands—to the extent that clocks were eventually developed in which the function of timekeeping was secondary to the control of large numbers of figures engaged in various activities, to the point of acting out entire plays.

Finally, clockwork mechanisms appeared which had done away altogether with any pretence at timekeeping. These 'automata' were entirely devoted to imparting lifelike motion to a mechanical figure or animal. These mechanical automaton simulations of life included such things as elephants, peacocks, singing birds, musicians, and even fortune-tellers.

This line of development reached its peak in the famous duck of Vaucanson, described as 'an artificial duck made of gilded copper who drinks, eats, quacks, splashes about on the water, and digests his food like a living duck'.¹ Vaucanson's goal is captured neatly in the following description:

In 1735 Jacques de Vaucanson arrived in Paris at the age of 26. Under the influence of contemporary philosophic ideas, he had tried, it seems, to reproduce life artificially.

Unfortunately, neither the duck itself nor any technical descriptions or diagrams remain that would give the details of construction of this duck. The complexity

¹ All quotes concerning these mechanical ducks are from Chapuis and Droz (1958).

of the mechanism is attested to by the fact that one single wing contained over 400 articulated pieces.

One of those called upon to repair Vaucanson's duck in later years was a 'mechanician' named Reichsteiner, who was so impressed with it that he went on to build a duck of his own—also now lost—which was exhibited in 1847. Here is an account of this duck's operation from the newspaper *Das Freie Wort*:

After a light touch on a point on the base, the duck in the most natural way in the world begins to look around him, eyeing the audience with an intelligent air. His lord and master, however, apparently interprets this differently, for soon he goes off to look for something for the bird to eat. No sooner has he filled a dish with oatmeal porridge than our famished friend plunges his beak deep into it, showing his satisfaction by some characteristic movements of his tail. The way in which he takes the porridge and swallows it greedily is extraordinarily true to life. In next to no time the basin has been half emptied, although on several occasions the bird, as if alarmed by some unfamiliar noises, has raised his head and glanced curiously around him. After this, satisfied with his frugal meal, he stands up and begins to flap his wings and to stretch himself while expressing his gratitude by several contented quacks. But most astonishing of all are the contractions of the bird's body clearly showing that his stomach is a little upset by this rapid meal and the effects of a painful digestion become obvious. However, the brave little bird holds out, and after a few moments we are convinced in the most concrete manner that he has overcome his internal difficulties. The truth is that the smell which now spreads through the room becomes almost unbearable. We wish to express to the artist inventor the pleasure which his demonstration gave to us.

Fig. 1.1 shows two views of one of the ducks—there is some controversy as to whether it is Vaucanson's or Reichsteiner's. The mechanism inside the duck would have been completely covered with feathers and the controlling mechanism in the box below would have been covered up as well.

2.1. The Development of Control Mechanisms

Out of the technology of the clockwork regulation of automata came the more general—and perhaps ultimately more important—technology of *process control*. As attested to in the descriptions of the mechanical ducks, some of the clockwork mechanisms had to control remarkably complicated actions on the part of the automata, not only *powering* them but *sequencing* them as well.

Control mechanisms evolved from early, simple devices—such as a lever attached to a wheel which converted circular motion into linear motion—to later, more complicated devices—such as whole sets of cams upon which would ride many interlinked mechanical arms, giving rise to extremely complicated automaton behaviours.

Eventually *programmable controllers* appeared, which incorporated such devices as interchangeable cams, or drums with movable pegs, with which one could program arbitrary sequences of actions on the part of the automaton. The




FIG. 1.1. Two views of the mechanical duck attributed to Vaucanson. From Chapuis and Droz (1958). Reprinted by permission

writing and picture-drawing automata of Fig. 1.2, built by the Jaquet-Droz family, are examples of programmable automata. The introduction of such programmable controllers was one of the primary developments on the road to general-purpose computers.

2.2. Abstraction of the Logical 'Form' of Machines

During the early part of the twentieth century, the formal application of logic to the mechanical process of arithmetic led to the abstract formulation of a 'procedure'. The work of Church, Kleene, Gödel, Turing, and Post formalized the notion of a logical sequence of steps, leading to the realization that the essence of a mechanical process—the 'thing' responsible for its dynamic behaviour—is not a thing at all, but an abstract control structure, or 'program'—a sequence of simple actions selected from a finite repertoire. Furthermore, it was recognized that the essential features of this control structure could be captured within an abstract set of rules—a formal specification—without regard to the material out of which the machine was constructed. The 'logical form' of a machine was separated from its material basis of construction, and it was found that 'machiness' was a property of the former, not of the latter. Today, the




FIG. 1.2. Two views of a drawing automaton built by the Jaquet-Droz family. From Chapuis and Droz (1958). Reprinted by permission

formal equivalent of a 'machine' is an *algorithm*: the logic underlying the dynamics of an automaton, regardless of the details of its material construction. We now have many formal methods for the specification and operation of abstract machines: such as programming languages, formal language theory, automata theory, recursive function theory, etc. All these have been shown to be logically equivalent.

Once we have learned to think of machines in terms of their abstract, formal specifications, we can turn around and view abstract, formal specifications as potential machines. In mapping the machines of our common experience to formal specifications, we have by no means exhausted the space of *possible* specifications. Indeed, most of our individual machines map to a very small subset of the space of specifications—a subset largely characterized by methodical, boring, uninteresting dynamics.

2.3. General-Purpose Computers

Various threads of technological development—programmable controllers, calculating engines, and the formal theory of machines—have come together in the general-purpose, stored-program computer. Programmable computers are extremely general behaviour generators. They have no intrinsic behaviour of

their own. Without programs, they are like formless matter. They must be told how to behave. By submitting a program to a computer—that is: by giving it a formal specification for a machine—we are telling it to behave as if it were the machine specified by the program. The computer then ‘emulates’ that more specific machine in the performance of the desired task. Its great power lies in its plasticity of behaviour. If we can provide a step-by-step specification for a specific kind of behaviour, the chameleon-like computer will exhibit that behaviour. Computers should be viewed as *second-order* machines—given the formal specification of a first-order machine, they will ‘become’ that machine. Thus the space of possible machines is directly available for study, at the cost of a mere formal description: computers ‘realize’ abstract machines.

2.4. *Formal Limits of Machine Behaviours*

3. THE ROLE OF COMPUTERS IN STUDYING LIFE AND OTHER COMPLEX SYSTEMS

Artificial Intelligence and Artificial Life are each concerned with the application of computers to the study of complex, natural phenomena. Both are concerned with generating complex behaviour. However, the manner in which each field employs the technology of computation in the pursuit of its respective goals is strikingly different.

AI has based its underlying methodology for generating intelligent behaviour on the computational paradigm. That is, AI uses the technology of computation as a model of intelligence. AL, on the other hand, is attempting to develop a new computational paradigm based on the natural processes that support living organisms. That is, AL uses insights from biology to explore the dynamics of interacting information structures. AL has not adopted the computational paradigm as its underlying methodology of behaviour generation, nor does it attempt to 'explain' life as a kind of computer program.

One way to pursue the study of artificial life would be to attempt to create life *in vitro*, using the same kinds of organic chemicals out of which we are constituted. Indeed, there are numerous exciting efforts in this direction. This would certainly teach us a lot about the possibilities for alternative life-forms *within* the carbon-chain chemistry domain that could have (but didn't) evolve here.

However, biomolecules are extremely small and difficult to work with, requiring rooms full of special equipment, replete with dozens of 'postdocs' and graduate students willing to devote the larger part of their professional careers to the perfection of electrophoretic gel techniques. Besides, although the creation of life *in vitro* would certainly be a scientific feat worthy of note—and probably even a Nobel prize—it would not, in the long run, tell us much more about the space of *possible* life than we already know.

Computers provide an alternative medium within which to attempt to synthesize life. Modern computer technology has resulted in machinery with tremendous potential for the creation of life *in silico*.

Computers should be thought of as an important laboratory tool for the study of life, substituting for the array of incubators, culture dishes, microscopes, electrophoretic gels, pipettes, centrifuges, and other assorted wet-lab paraphernalia, one simple-to-master piece of experimental equipment devoted exclusively to the incubation of information structures.

The advantage of working with information structures is that information has no intrinsic size. The computer is *the* tool for the manipulation of information, whether that manipulation is a consequence of our actions or a consequence of the actions of the information structures themselves. Computers themselves will not be alive, rather they will support informational universes within which

dynamic populations of informational 'molecules' engage in informational 'biochemistry'.

This view of computers as workstations for performing scientific experiments within artificial universes is fairly new, but it is rapidly becoming accepted as a legitimate, even necessary, way of pursuing science. In the days before computers, scientists worked primarily with systems whose defining equations could be solved analytically, and ignored those whose defining equations could *not* be so solved. This was largely the case because, in the absence of analytic solutions, the equations would have to be integrated over and over again, essentially simulating the time behaviour of the system. Without computers to handle the mundane details of these calculations, such an undertaking was unthinkable except in the simplest cases.

However, with the advent of computers, the necessary mundane calculations can be relegated to these idiot-savants, and the realm of numerical simulation is opened up for exploration. 'Exploration' is an appropriate term for the process, because the numerical simulation of systems allows one to 'explore' the system's behaviour under a wide range of parameter settings and initial conditions. The heuristic value of this kind of experimentation cannot be overestimated. One often gains tremendous insight for the essential dynamics of a system by observing its behaviour under a wide range of initial conditions. Most importantly, however, computers are beginning to provide scientists with a new paradigm for modelling the world. When dealing with essentially unsolvable governing equations, the primary reason for producing a formal mathematical model—the hope of reaching an analytic solution by symbolic manipulation—is lost. Systems of ordinary and partial differential equations are not very well suited for implementation as computer algorithms. One might expect that other modelling technologies would be more appropriate when the goal is the *synthesis*, rather than the *analysis*, of behaviour.⁶

This expectation is easily borne out. With the precipitous drop in the cost of raw computing power, computers are now available that are capable of simulating physical systems from first principles. This means that it has become possible, for example, to model turbulent flow in a fluid by simulating the motions of its constituent particles—not just approximating *changes* in concentrations of particles at particular points, but actually computing their motions exactly (Frisch *et al.* 1986; Wolfram 1986; Toffoli and Margolus 1987).

What does all this have to do with the study of life? The most surprising lesson we have learned from simulating complex physical systems on computers is that *complex behaviour need not have complex roots*. Indeed, tremendously interesting and beguilingly complex behaviour can emerge from collections of

⁶ See Toffoli (1984) for a good exposition.

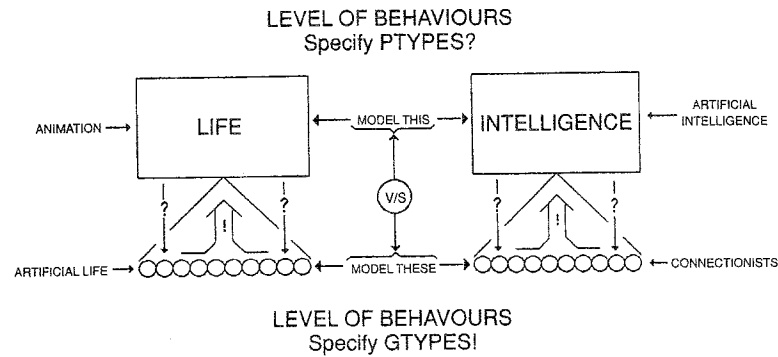


FIG. 1.4. The bottom-up versus the top-down approach to modelling complex systems. From Langton (1989a)

extremely simple components. This leads directly to the exciting possibility that much of the complex behaviour exhibited by nature—especially the complex behaviour that we call life—*also* has simple generators. Since it is very hard to work backwards from a complex behaviour to its generator, but very simple to create generators and synthesize complex behaviour, a promising approach to the study of complex natural systems is to undertake the general study of the kinds of behaviour that can emerge from distributed systems consisting of simple components (Fig. 1.4).

4. NON-LINEARITY AND LOCAL DETERMINATION OF BEHAVIOUR

4.1. Linear vs. Non-linear Systems

As mentioned briefly above, the distinction between linear and non-linear systems is fundamental, and provides excellent insight into why the principles underlying the dynamics of life should be so hard to find. The simplest way to state the distinction is to say that *linear systems* are those for which the behaviour of the whole is just the sum of the behaviour of its parts, while for *non-linear systems*, the behaviour of the whole is *more* than the sum of its parts.

Linear systems are those which obey the *principle of superposition*. We can break up complicated linear systems into simpler constituent parts, and analyse these parts *independently*. Once we have reached an understanding of the parts in isolation, we can achieve a full understanding of the whole system by *composing*

our understandings of the isolated parts. This is the key feature of linear systems: by studying the parts in isolation, we can learn everything we need to know about the complete system.

This is not possible for non-linear systems, which do *not* obey the principle of superposition. Even if we could break such systems up into simpler constituent parts, and even if we could reach a complete understanding of the parts in isolation, we would not be able to compose our understandings of the individual parts into an understanding of the whole system. The key feature of non-linear systems is that their primary behaviours of interest are properties of the *interactions between parts*, rather than being properties of the parts themselves, and these interaction-based properties necessarily disappear when the parts are studied independently.

Thus analysis is most fruitfully applied to linear systems. Analysis has *not* proved anywhere near as effective when applied to non-linear systems: the non-linear system must be treated as a whole.

A different approach to the study of non-linear systems involves the inverse of analysis: *synthesis*. Rather than start with the behaviour of interest and attempting to analyse it into its constituent parts, we start with constituent parts and put them together in the attempt to *synthesize* the behaviour of interest.

Life is a property of *form*, not *matter*, a result of the organization of matter rather than something that inheres in the matter itself. Neither nucleotides nor amino acids nor any other carbon-chain molecule is alive—yet put them together in the right way, and the dynamic behaviour that emerges out of their interactions is what we call life. It is effects, not things, upon which life is based—life is a kind of behaviour, not a kind of stuff—and as such, it is constituted of simpler behaviours, not simpler stuff. *Behaviours themselves* can constitute the fundamental parts of non-linear systems—*virtual parts*, which depend on non-linear interactions between physical parts for their very existence. Isolate the physical parts and the virtual parts cease to exist. It is the *virtual parts* of living systems that Artificial Life is after, and synthesis is its primary methodological tool.

4.2. The Parsimony of Local Determination of Behaviour

It is easier to generate complex behaviour from the application of simple, *local* rules than it is to generate complex behaviour from the application of complex, *global* rules. This is because complex global behaviour is usually due to non-linear interactions occurring at the local level. With bottom-up specifications, the system computes the local, non-linear interactions explicitly and the global behaviour, which was implicit in the local rules, emerges spontaneously without being treated explicitly.

With top-down specifications, however, local behaviour must be implicit in global rules! This is really putting the cart before the horse! The global rules must 'predict' the effects on global structure of many local, non-linear interactions—something which we have seen is intractable, even impossible, in the general case. Thus top-down systems must take computational short cuts and explicitly deal with special cases, which results in inflexible, brittle, and unnatural behaviour.

Furthermore, in a system of any complexity, the number of possible global states is astronomically enormous, and grows exponentially with the size of the system. Systems that attempt to supply *global* rules for *global* behaviour simply *cannot* provide a different rule for every global state. Thus the global states must be classified in some manner, and categorized using a coarse-grained scheme according to which the global states within a category are indistinguishable. The rules of the system can only be applied at the level of resolution of these categories. There are many possible ways to implement a classification scheme, most of which will yield different partitionings of the global-state space. Any rule-based system must necessarily *assume* that finer-grained differences do not matter, or must include a finite set of tests for 'special cases', and then must assume that no *other* special cases are relevant.

For most complex systems, however, fine differences in the global state can result in enormous differences in global behaviour, and there may be no way in principle to partition the space of global states in such a way that specific fine differences have the appropriate global impact.

On the other hand, systems that supply *local* rules for *local* behaviours *can* provide a different rule for each and every possible local state. Furthermore, the size of the local-state space can be completely independent of the size of the system. In local rule-governed systems, each local state, and consequently the global state, can be determined exactly and precisely. Fine differences in the global state will result in very specific differences in the local state and, consequently, will affect the invocation of local rules. As fine differences affect local behaviour, the difference will be felt in an expanding patch of local states, and in this manner—propagating from local neighbourhood to local neighbourhood—fine differences in global state can result in large differences in global behaviour. The only 'special cases' explicitly dealt with in locally determined systems are exactly the set of all possible local states, and the rules for these are just exactly the set of all local rules governing the system.

5. BIOLOGICAL AUTOMATA

Organisms have been compared to extremely complicated and finely tuned biochemical machines. Since we know that it is possible to abstract the logical

form of a machine from its physical hardware, it is natural to ask whether it is possible to abstract the logical form of an organism from its biochemical wetware. The field of Artificial Life is devoted to the investigation of this question.

In the following sections we shall look at the manner in which behaviour is generated in a bottom-up fashion in living systems. We then generalize the mechanisms by which this behaviour generation is accomplished, so that we may apply them to the task of generating behaviour in artificial systems.

We shall find that the essential machinery of living organisms is quite a bit different from the machinery of our own invention, and we would be quite mistaken to attempt to force our preconceived notions of abstract machines on to the machinery of life. The difference, once again, lies in the exceedingly parallel and distributed nature of the operation of the machinery of life, as contrasted with the singularly serial and centralized control structures associated with the machines of our invention.

5.1. Genotypes and Phenotypes

The most salient characteristic of living systems, from the behaviour generation point of view, is the *genotype/phenotype* distinction. The distinction is essentially one between a specification of machinery—the genotype—and the behaviour of that machinery—the phenotype.

The *genotype* is the complete set of genetic instructions encoded in the linear sequence of nucleotide bases that makes up an organism's DNA. The *phenotype* is the physical organism itself—the structures that emerge in space and time as the result of the interpretation of the genotype in the context of a particular environment. The process by which the phenotype develops through time under the direction of the genotype is called *morphogenesis*. The individual genetic instructions are called *genes* and consist of short stretches of DNA. These instructions are 'executed'—or *expressed*—when their DNA sequence is used as a template for transcription. In the case of protein synthesis, transcription results in a duplicate nucleotide strand known as a *messenger RNA*—or *mRNA*—constructed by the process of base-pairing. This mRNA strand may then be modified in various ways before it makes its way out to the cytoplasm where, at bodies known as *ribosomes*, it serves as a template for the construction of a linear chain of *amino acids*. The resulting *polypeptide* chain will fold up on itself in a complex manner, forming a tightly packed molecule known as a *protein*. The finished protein detaches from the ribosome and may go on to serve as a passive structural element in the cell, or may have a more active role as an *enzyme*. Enzymes are the functional molecular 'operators' in the logic of life.

One may consider the genotype as a largely unordered 'bag' of instructions, each one of which is essentially the specification for a 'machine' of some sort—

passive or active. When instantiated, each such 'machine' will enter into the ongoing logical 'fray' in the cytoplasm, consisting largely of local interactions between other such machines. Each such instruction will be 'executed' when its own triggering conditions are met and will have specific, local effects on structures in the cell. Furthermore, each such instruction will operate within the context of all the other instructions that have been—or are being—executed.

The phenotype, then, consists of the structures and dynamics that emerge through time in the course of the execution of the parallel, distributed 'computation' controlled by this genetic 'bag' of instructions. Since genes' interactions with one another are highly non-linear, the phenotype is a non-linear function of the genotype.

5.2. Generalized Genotypes and Phenotypes

In the context of Artificial Life, we need to generalize the notions of *genotype* and *phenotype*, so that we may apply them in non-biological situations. We shall use the term *generalized genotype*—or *GTYPE*—to refer to any largely unordered set of low-level rules, and we shall use the term *generalized phenotype*—or *PTYPE*—to refer to the behaviours and/or structures that emerge out of the interactions among these low-level rules when they are activated within the context of a specific environment. The *GTYPE*, essentially, is the specification for a set of machines, while the *PTYPE* is the behaviour that results as the machines are run and interact with one another.

This is the bottom-up approach to the generation of behaviour. A set of entities is defined, and each entity is endowed with a specification for a simple behavioural repertoire—a *GTYPE*—that contains instructions which detail its reactions to a wide range of *local* encounters with other such entities or with specific features of the environment. Nowhere is the behaviour of the set of entities as a whole specified. The global behaviour of the aggregate—the *PTYPE*—emerges out of the collective interactions among individual entities.

It should be noted that the *PTYPE* is a multi-level phenomenon. First, there is the *PTYPE* associated with each particular instruction—the effect which that instruction has on an entity's behaviour when it is expressed. Second, there is the *PTYPE* associated with each individual entity—its individual behaviour within the aggregate. Third, there is the *PTYPE* associated with the behaviour of the aggregate as a whole.

This is true for natural systems as well. We can talk about the phenotypic trait associated with a particular gene, we can identify the phenotype of an individual cell, and we can identify the phenotype of an entire multi-cellular organism—its body, in effect. *PTYPES* should be complex and multi-level. If we want

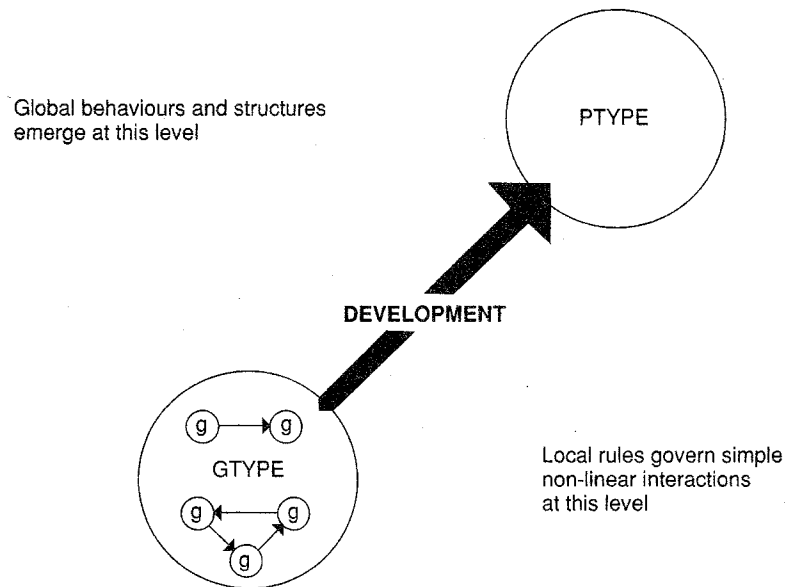


FIG. 1.5. The relationship between *GTYPE* and *PTYPE*. From Langton (1989a)

to simulate life, we should expect to see hierarchical structures emerge in our simulations. In general, phenotypic traits at the level of the whole organism will be the result of many non-linear interactions between genes, and there will be no single gene to which one can assign responsibility for the vast majority of phenotypic traits.

In summary, *GTYPES* are low-level rules for behaviors—i.e. abstract specifications for 'machines'—which will engage in local interactions within a large aggregate of other such behaviors. *PTYPES* are the behaviors—the structures in time and space—that *develop* out of these non-linear, local interactions (Fig. 1.5).

5.3. Unpredictability of *PTYPE* from *GTYPE*

Non-linear interactions between the objects specified by the *GTYPE* provide the basis for an extremely rich variety of possible *PTYPES*. *PTYPES* draw on the full combinatorial potential implicit in the set of possible interactions between

low-level rules. The other side of the coin, however, is that we cannot predict the PTYPES that will emerge from specific GTYPES, due to the general unpredictability of non-linear systems. If we wish to maintain the property of predictability, then we must restrict severely the non-linear dependence of PTYPE on GTYPE, but this forces us to give up the combinatorial richness of possible PTYPES. Therefore, a trade-off exists between behavioural richness and predictability (or 'programmability'). We shall see in the section on evolution that the lack of programmability is adequately compensated for by the increased capacity for adaptiveness provided by a rich behavioural repertoire.

As discussed previously, we know that it is impossible in the general case to determine *any* non-trivial property of the future behaviour of a sufficiently powerful computer from a mere inspection of its program and its initial state alone (Hopcroft and Ullman 1979). A Turing machine—the formal equivalent of a general-purpose computer—can be captured within the scheme of GTYPE/PTYPE systems by identifying the machine's transition table as the GTYPE and the resulting computation as the PTYPE. From this we can deduce that in the general case it will not be possible to determine, by inspection alone, any non-trivial feature of the PTYPE that will emerge from a given GTYPE in the context of a particular initial configuration. In general, the only way to find out anything about the PTYPE is to start the system up and watch what happens as the PTYPE develops under control of the GTYPE and the environment.

Similarly, it is not possible in the general case to determine what specific alterations must be made to a GTYPE to effect a desired change in the PTYPE. The problem is that any specific PTYPE trait is, in general, an effect of many, many non-linear interactions between the behavioural primitives of the system (an 'epistatic trait' in biological terms). Consequently, given an arbitrary proposed change to the PTYPE, it may be impossible to determine by any formal procedure exactly what changes would have to be made to the GTYPE to effect that—and *only* that—change in the PTYPE. It is not a practically computable problem. There is no way to calculate the answer—short of exhaustive search—even though there may be an answer!⁷

The only way to proceed in the face of such an unpredictability result is by a process of trial and error. However, some processes of trial and error are more efficient than others. In natural systems, trial and error are interlinked in such a way that error guides the choice of trials under the process of evolution by natural selection. It is quite likely that this is the *only* efficient, *general* procedure that could find GTYPES with specific PTYPE traits when non-linear functions are involved.

⁷ An example in biology would be: What changes would have to be made to the genome in order to produce six fingers on each hand rather than five?

6. RECURSIVELY GENERATED OBJECTS

In the previous section, we described the distinction between genotype and phenotype, and we introduced their generalizations in the form of GTYPES and PTYPES. In this section, we shall review a general approach to building GTYPE/PTYPE systems based on the methodology of *recursively generated objects*.

A major appeal of this approach is that it arises naturally from the GTYPE/PTYPE distinction: the local developmental rules—the recursive description itself—constitute the *GTYPE*, and the developing structure—the recursively generated object or behaviour itself—constitutes the *PTYPE*.

Under the methodology of recursively generated objects, the 'object' is a structure that has sub-parts. The rules of the system specify how to modify the most elementary, 'atomic' sub-parts, and are usually sensitive to the *context* in which these atomic sub-parts are embedded. That is, the state of the 'neighbourhood' of an atomic sub-part is taken into account in determining which rule to apply in order to modify that sub-part. It is usually the case that there are no rules in the system whose context is the entire structure; that is, there is no use made of *global* information. Each piece is modified solely on the basis of its own state and the state of the pieces 'nearby'.

Of course, if the initial structure consists of a single part—as might be the case with the initial seed—then the context for applying a rule is necessarily global. The usual situation is that a structure consists of *many* parts, only a local subset of which determine the rule that will be used to modify any one sub-part of the structure.

A recursively generated object, then, is a kind of PTYPE, and the recursive description that generates it is a kind of GTYPE. The PTYPE will emerge under the action of the GTYPE, developing through time via a process akin to morphogenesis.

We shall illustrate the notion of recursively generated objects with examples taken from the literature on L-systems, cellular automata, and computer animation.

6.1. Example 1: Lindenmayer Systems

Lindenmayer systems (L-systems) consist of sets of rules for rewriting strings of symbols, and bear strong relationships to the formal grammars treated by Chomsky. We shall give several examples of L-systems illustrating the methodology of recursively generated objects.⁸

In the following ' $X \rightarrow Y$ ' means that one replaces every occurrence of

⁸ For a more detailed review, see Prusinkiewicz (1991).

symbol X in the structure with string Y . Since the symbol X may appear on the right as well as the left sides of some rules, the set of rules can be applied 'recursively' to the newly rewritten structures. The process can be continued *ad infinitum* although some sets of rules will result in a 'final' configuration when no more changes occur.

Simple Linear Growth. Here is an example of the simplest kind of L-system. The rules are *context free*, meaning that the context in which a particular part is situated is *not* considered when altering it. There must be only one rule per part if the system is to be deterministic.

The rules (the 'recursive description' or GTYPE):

- 1) $A \rightarrow CB$
- 2) $B \rightarrow A$
- 3) $C \rightarrow DA$
- 4) $D \rightarrow C$

When applied to the initial seed structure 'A', the following structural history develops (each successive line is a successive time-step):


time	structure	rules applied (L to R)
0	A	(initial 'seed')
1	C B	(rule 1 replaces A with CB)
2	D A A	(rule 3 replaces C with DA and rule 2 replaces B with A)
3	C C B C B	(rule 4 replaces D with C and rule 1 replaces the two As with CBs)
4	... (etc.) ...	

And so forth.

The 'PTYPE' that emerges from this kind of recursive application of a simple, local rewriting rule can get extremely complex. These kinds of grammars (whose rules replace single symbols) have been shown to be equivalent to the operation of finite-state machines. With appropriate restrictions, they are also equivalent to the 'regular languages' defined by Chomsky.

Branching Growth. L-systems incorporate meta-symbols to represent branching points, allowing a new line of symbols to branch off from the main 'stem' (see Fig. 1.6).

The following grammar produces branching structures. The '()' and '[]' notations indicate left and right branches, respectively, and the strings within them indicate the structure of the branches themselves.



$n=5, \delta=18^\circ$

ω : plant
 p_1 : plant \rightarrow internode + [plant + flower] - - //
 [- - leaf] internode [+ + leaf] -
 [plant flower] + + plant flower
 p_2 : internode \rightarrow F seg [// & leaf] [// ^ ^ leaf] F seg
 p_3 : seg \rightarrow seg F seg
 p_4 : leaf \rightarrow [' { +f-ff-f+ | +f-ff-f }]
 p_5 : flower \rightarrow [& & pedicel ' / wedge // // wedge // //
 wedge // // wedge // // wedge]
 p_6 : pedicel \rightarrow FF
 p_7 : wedge \rightarrow [' ^ F] [{ & & & -f+f | -f+f }]

FIG. 1.6. An L-system plant grown from rules incorporating graphical rendering information. From Prusinkiewicz (1991)

is the entire universe: a lattice of finite automata. The local rule set—the GTYPE—in this case is the transition function obeyed homogeneously by every automaton in the lattice. The local context taken into account in updating the state of each automaton is the state of the automata in its immediate neighbourhood. The transition function for the automata constitutes a *local* physics for a simple, discrete space/time universe. The universe is updated by applying the local physics to each local ‘cell’ of its structure over and over again. Thus, although the physical structure itself does not develop over time, its *state* does.

Within such universes, one can embed all manner of processes, relying on the context sensitivity of the rules to local neighbourhood conditions to propagate information around within the universe ‘meaningfully’. In particular, one can embed general-purpose computers. Since these computers are simply particular configurations of states within the lattice of automata, *they can compute over the very set of symbols out of which they are constructed*. Thus, structures in this universe can compute and construct other structures, which also may compute and construct.

For example, here is the simplest known structure that can reproduce itself:

```

  2 2 2 2 2 2 2 2
  2 1 7 0 1 4 0 1 4 2
  2 0 2 2 2 2 2 2 0 2
  2 7 2      2 1 2
  2 1 2      2 1 2
  2 0 2      2 1 2
  2 7 2      2 1 2
  2 1 2 2 2 2 2 2 1 2 2 2 2 2
  2 0 7 1 0 7 1 0 7 1 1 1 1 1 2
  2 2 2 2 2 2 2 2 2 2 2 2

```

Each number is the state of one automaton in the lattice. Blank space is presumed to be in state ‘0’. The ‘2’-states form a sheath around the ‘1’-state data path. The ‘7 0’ and ‘4 0’ state pairs constitute signals embedded within the data path. They will propagate counter-clockwise around the loop, cloning off copies which propagate down the extended tail as they pass the T-junction between loop and tail. When the signals reach the end of the tail, they have the following effects: each ‘7 0’ signal extends the tail by one unit, and the two ‘4 0’ signals construct a left-hand corner at the end of the tail. Thus for each full cycle of the instructions around the loop, another side and corner of an ‘offspring-loop’ will be constructed. When the tail finally runs into itself after four cycles, the collision of signals results in the disconnection of the two loops as well as the construction of a tail on each of the loops.

After 151 time-steps, this system will evolve to the following configuration:

```

  2
  2 1 2
  2 7 2
  2 0 2
  2 1 2
  2 2 2 2 2 2 2 7 2      2 2 2 2 2 2 2 2
  2 1 1 1 7 0 1 7 0 2    2 1 7 0 1 4 0 1 4 2
  2 1 2 2 2 2 2 2 1 2    2 0 2 2 2 2 2 2 0 2
  2 1 2      2 7 2 2 7 2      2 1 2
  2 1 2      2 0 2 2 1 2      2 1 2
  2 4 2      2 1 2 2 0 2      2 1 2
  2 1 2      2 7 2 2 7 2      2 1 2
  2 0 2 2 2 2 2 2 0 2    2 1 2 2 2 2 2 2 1 2 2 2 2 2
  2 4 1 0 7 1 0 7 1 2    2 0 7 1 0 7 1 0 7 1 1 1 1 1 2
  2 2 2 2 2 2 2 2      2 2 2 2 2 2 2 2 2 2 2 2

```

Thus, the initial configuration has succeeded in reproducing itself.

Each of these loops will go on to reproduce itself in a similar manner, giving rise to an expanding *colony* of loops, growing out into the array.

These embedded self-reproducing loops are the result of the recursive application of a rule to a seed structure. In this case, the primary rule that is being recursively applied constitutes the ‘physics’ of the universe. The initial state of the loop itself constitutes a little ‘computer’ under the recursively applied physics of the universe: a computer whose program causes it to construct a copy of itself. The ‘program’ within the loop computer is also applied recursively to the growing structure. Thus, this system really involves a double level of recursively applied rules. The mechanics of applying one recursive rule within a universe whose physics is governed by another recursive rule had to be worked out by trial and error. This system makes use of the signal propagation capacity to embed a structure that itself *computes* the resulting structure, rather than having the ‘physics’ directly responsible for developing the final structure from a passive seed.

This captures the flavour of what goes on in natural biological development: the genotype codes for the constituents of a dynamic process in the cell, and it is this dynamic process that is primarily responsible for mediating—or ‘computing’—the expression of the genotype in the course of development.

6.3. Example 3: Flocking ‘Boids’

The previous examples were largely concerned with the growth and development of *structural* PTYPES. Here, we give an example of the development of a *behavioural* PTYPE.

Craig Reynolds (1987) has implemented a simulation of flocking behaviour. In this model—which is meant to be a general platform for studying the qualitatively similar phenomena of flocking, herding, and schooling—one has a large collection of autonomous but interacting objects (which Reynolds refers to as ‘Boids’), inhabiting a common simulated environment.

The modeller can specify the manner in which the individual Boids will respond to *local* events or conditions. The global behaviour of the aggregate of Boids is strictly an emergent phenomenon, none of the rules for the individual Boids depends on global information, and the only updating of the global state is done on the basis of individual Boids responding to local conditions.

Each Boid in the aggregate shares the same behavioural ‘tendencies’:

- to maintain a minimum distance from other objects in the environment, including other Boids,
- to match velocities with Boids in its neighbourhood, and
- to move towards the perceived centre of mass of the Boids in its neighbourhood.

These are the only rules governing the behaviour of the aggregate.

These rules, then, constitute the generalized genotype (GTYPE) of the Boids system. They say nothing about structure, or growth and development, but they determine the behaviour of a set of interacting objects, out of which very natural motion emerges.

With the right settings for the parameters of the system, a collection of Boids released at random positions within a volume will collect into a dynamic flock, which flies around environmental obstacles in a very fluid and natural manner, occasionally breaking up into sub-flocks as the flock flows around both sides of an obstacle. Once broken up into sub-flocks, the sub-flocks reorganize around their own, now distinct and isolated centres of mass, only to re-merge into a single flock again when both sub-flocks emerge at the far side of the obstacle and each sub-flock feels anew the ‘mass’ of the other sub-flock (Fig. 1.7).

The flocking behaviour itself constitutes the generalized phenotype (PTYPE) of the Boids system. It bears the same relation to the GTYPE as an organism’s morphological phenotype bears to its molecular genotype. The same distinction between the *specification* of machinery and the *behaviour* of machinery is evident.

6.4. Discussion of Examples

In all the above examples, the recursive rules apply to *local structures* only, and the PTYPE—structural or behavioural—that results at the global level emerges out of all local activity taken collectively. Nowhere in the system are there rules




FIG. 1.7. A flock of ‘Boids’ negotiating a field of columns. Sequence generated by Craig Reynolds. From Langton (1989a)

for the behaviour of the system at the global level. This is a much more powerful and simple approach to the generation of complex behaviour than that typically taken in AI, for instance, where ‘expert systems’ attempt to provide global rules for global behaviour. Recursive, ‘bottom up’ specifications yield much more natural, fluid, and flexible behaviour at the global level than typical ‘top down’ specifications, and they do so *much* more parsimoniously.

Importance of Context Sensitivity. It is worth while to note that *context-sensitive* rules in GTYPE/PTYPE systems provide the possibility for non-linear interactions among the parts. Without context sensitivity, the systems would be linearly

decomposable, information could not 'flow' throughout the system in any meaningful manner, and complex long-range dependencies between remote parts of the structures could not develop.

Feedback between the Local and the Global Levels. There is also a very important feedback mechanism *between* levels in such systems: the interactions among the low-level entities give rise to the global-level dynamics which, in turn, affects the lower levels by *setting the local context* within which each entity's rules are invoked. Thus, local behaviour supports global dynamics, which shapes local context, which affects local behaviour, which supports global dynamics, and so forth.

6.5. Genuine Life in Artificial Systems

It is important to distinguish the ontological status of the various levels of behaviour in such systems. At the level of the individual behaviours, we have a clear difference in kind: Boids are *not* birds, they are not even remotely like birds, they have no cohesive physical structure, but rather they exist as information structures—processes—within a computer. But—and this is *the* critical 'But'—at the level of behaviours, *flocking Boids and flocking birds are two instances of the same phenomenon: flocking.*

The behaviour of a flock as a whole does not depend critically on the internal details of the entities of which it is constituted, only on the details of the way in which these entities behave in each other's presence. Thus, flocking in Boids is true flocking, and may be counted as another empirical data point in the study of flocking behaviour in general, right up there with flocks of geese and flocks of starlings.

This is *not* to say that flocking Boids capture *all* the nuances upon which flocking behaviour depends, or that the Boids' behavioural repertoire is sufficient to exhibit all the different modes of flocking that have been observed—such as the classic 'V' formation of flocking geese. The crucial point is that we have captured, within an aggregate of artificial entities, a *bona fide* lifelike behaviour, and that the behaviour emerges within the artificial system in the same way that it emerges in the natural system.

The same is true for L-systems and the self-reproducing loops. The constituent parts of the artificial systems are different kinds of things from their natural counterparts, but the emergent behaviours that they support are the same kinds of thing as their natural counterparts: genuine morphogenesis and differentiation for L-systems, and genuine self-reproduction in the case of the loops.

The claim is the following. The 'artificial' in Artificial Life refers to the component parts, not the emergent processes. If the component parts are implemented

correctly, the processes they support are *genuine*—every bit as genuine as the natural processes they imitate.

The *big* claim is that a properly organized set of artificial primitives carrying out the same functional roles as the biomolecules in natural living systems will support a process that will be 'alive' in the same way that natural organisms are alive. Artificial Life will therefore be *genuine* life—it will simply be made of different stuff than the life that has evolved here on Earth.

7. EVOLUTION



7.1. Evolution: From Artificial Selection to Natural Selection

Modern organisms owe their structure to the complex process of biological evolution, and it is very difficult to discern which of their properties are due to chance and which to necessity. If biologists could 'rewind the tape' of evolution and start it over, again and again, from different initial conditions, or under different regimes of external perturbations along the way, they would have a full *ensemble* of evolutionary pathways to generalize over. Such an ensemble would allow them to distinguish universal, necessary properties (those which were observed in all the pathways in the ensemble) from accidental, chance properties (those which were unique to individual pathways). However, biologists cannot rewind the tape of evolution, and are stuck with a single, actual evolutionary trace out of a vast, intuited ensemble of possible traces.

Although studying computer models of evolution is not the same as studying the 'real thing', the ability to freely manipulate computer experiments—to 'rewind the tape', perturb the initial conditions, and so forth—can more than make up for their 'lack' of reality.

It has been known for some time that one can evolve computer programs by the process of natural selection among a population of variant programs. Each individual program in a population of programs is evaluated for its performance on some task. The programs that perform best are allowed to 'breed' with one another via *Genetic Algorithms* (Holland 1975; Goldberg 1989). The offspring of these better-performing parent programs replace the worst-performing programs in the population, and the cycle is iterated. Such evolutionary approaches to program improvement have been applied primarily to the tasks of function optimization and machine learning.

However, such evolutionary models have rarely been used to study evolution itself (Wilson 1989). Researchers have primarily concentrated on the *results*, rather than on the *process*, of evolution. In the spirit of von Neumann's research on self-reproduction via the study of self-reproducing *automata*, the following

8. CONCLUSION

This article is intended to provide a broad overview of the field of Artificial Life, its motivations, history, theory, and practice. In such a short space, it cannot hope to go into depth in any one of these areas. Rather, it attempts to convey the 'spirit' of the Artificial Life enterprise via several illustrative examples coupled with a good deal of motivating explanation and discussion.

The field of Artificial Life is in its infancy, and is currently engaged in a period of extremely rapid growth, which is producing many new converts to the principles detailed here. However, it is also raising a significant amount of controversy, and is not without its critics. The notion of studying biology via the study of patently non-biological things is an idea that is hard for the traditional biological community to accept. The acceptance of Artificial Life techniques within the biological community will be directly proportional to the contributions it makes to our understanding of biological phenomena.

That these contributions are forthcoming, I have no doubt. However, high-quality research in Artificial Life is difficult, because it requires that its practitioners be experts in both the computational sciences and the biological sciences. Either of these alone is a full-time career, and so the danger lurks of doing either masterful biology but trivial computing, or doing masterful computing but trivial biology.

Therefore, I strongly suggest incorporating a trick from nature: cooperate! As is amply illustrated in many of the examples discussed in this article, nature often discovers that collections of individuals easily solve problems that would be extremely difficult or even impossible for individuals to solve on their own. Collaborations between biologists and computer scientists are quite likely to be the most appropriate vehicles for making significant contributions to our understanding of biology via the pursuit of Artificial Life.

So, if you are a computer expert dying to hack together an evolution program, go find yourself a top-notch evolutionary biologist to collaborate with, one who will bring to the enterprise an in-depth understanding of the subtleties of the evolutionary process plus a proper set of open questions about evolution towards which your evolution program might be addressed.

On the other hand, if you are a field biologist interested in doing some numerical simulations in order to understand the ecological dynamics you are observing in the field, hook up with a top-notch parallel-computing expert, who will bring to the enterprise a thorough knowledge of the subtleties involved in multi-agent interactions, and will be in possession of an equally open set of questions, which you very well might find to be strikingly related to your own.

Above all, when in doubt, turn to Mother Nature. After all, she is smarter than you!¹⁰

REFERENCES

- Axelrod, R. (1984), *The Evolution of Cooperation* (New York: Basic Books).
 — and Hamilton, W. D. (1981), 'The Evolution of Cooperation', *Science*, 211: 1390–6.
 Burks, A. W. (1970) (ed.), *Essays on Cellular Automata* (Urbana, Ill.: University of Illinois Press).
 Buss, L. (1987), *The Evolution of Individuality* (Princeton: Princeton University Press).
 Chapuis, A., and Droz, E. (1958), *Automata: A Historical and Technological Study*, trans. A. Reid (London: Batsford Ltd.).
 Dawkins, R. (1989), 'The Evolution of Evolvability', in Langton (1989b), 201–20.
 Frisch, U., Hasslacher, B., and Pomeau, Y. (1986), 'Lattice Gas Automata for the Navier-Stokes Equation', *Physical Reviews and Letters*, 56: 1505–8.
 Goldberg, D. E. (1989), *Genetic Algorithms in Search, Optimization, and Machine Learning* (Reading, Mass.: Addison-Wesley).
 Hamilton, W. D. (1980), 'Sex Versus Non-Sex Versus Parasite', *OIKOS*, 35: 282–90.
 — (1982), 'Pathogens as Causes of Genetic Diversity in their Host Populations', in R. M. Anderson and R. M. May (eds.), *Population Biology of Infectious Diseases* (Berlin: Springer-Verlag, 1982), 269–96.
 Hillis, W. D. (1991), 'Co-evolving Parasites Improve Simulated Evolution as an Optimization Procedure', in Langton *et al.* (1991), 313–24.
 Holland, J. H. (1975), *Adaptation in Natural and Artificial Systems: An Introductory Analysis with Applications to Biology, Control, and Artificial Intelligence* (Ann Arbor: University of Michigan Press).
 — (1986), 'Escaping Brittleness: The Possibilities of General Purpose Learning Algorithms Applied to Parallel Rule-Based Systems', in R. S. Michalski, J. G. Carbonell, and T. M. Mitchell (eds.), *Machine Learning II* (New York: Kaufman), 593–623.
 Hopcroft, J. E., and Ullman, J. D. (1979), *Introduction to Automata Theory, Languages, and Computation* (Menlo Park, Calif.: Addison-Wesley).
 Koza, J. R. (1991), 'Genetic Evolution and Co-evolution of Computer Programs', in Langton *et al.* (1991), 603–30.
 Langton, C. G. (1984), 'Self-Reproduction in Cellular Automata', *Physica D*, 10/1–2: 135–44.
 — (1986), 'Studying Artificial Life with Cellular Automata', *Physica D*, 22: 120–49.

¹⁰ A large number of people have assisted in writing this paper, which is based on my lecture notes for the Complex Systems Summer School and on the overview of Artificial Life that served as an introduction to the proceedings of the first Artificial Life workshop. Besides the people credited in the latter paper, I should like to thank the following people for their help with this version: Tom Ray, Kristian Lindgren, Danny Hillis, and John Koza. I should also like to thank Ronda Butler-Villa and Della Ulbarri for their patience with me and for their skill in preparing the figures and text for publication.

- Langton, C. G. (1989a), 'Artificial Life', in Langton (1989b).
- (1989b), *Artificial Life: Proceedings of an Interdisciplinary Workshop on the Synthesis and Simulation of Living Systems* (Santa Fe Institute Studies in the Sciences of Complexity, Proceedings, 6; Redwood City, Calif.: Addison-Wesley).
- Taylor, C., Farmer, J. D. and Rasmussen, S. (1991) (eds.), *Artificial Life II* (Santa Fe Institute Studies in the Sciences of Complexity, Proceedings, 10; Redwood City, Calif.: Addison-Wesley).
- Lindgren, K. (1991), 'Evolutionary Phenomena in Simple Dynamics', in Langton *et al.* (1991), 295–312.
- Margolus, L. (1970), *Origin of Eucaryotic Cells* (New Haven: Yale University Press).
- Prusinkiewicz, P. (1991), *The Algorithmic Beauty of Plants* (Berlin: Springer-Verlag).
- Ray, T. S. (1991), 'An Approach to the Synthesis of Life', in Langton *et al.* (1991), 371–408, and Chapter 3 below.
- Reynolds, C. W. (1987), 'Flocks, Herds, and Schools: A Distributed Behavioral Model', Proceedings of SIGGRAPH '87, *Computer Graphics* V 21/4: 25–34.
- Toffoli, T. (1984), 'Cellular Automata as an Alternative to (Rather than an Approximation of) Differential Equations in Modeling Physics', in J. D. Farmer, T. Toffoli, and S. Wolfram (eds.), *Cellular Automata: Proceedings of an Interdisciplinary Workshop* (Los Alamos, New Mexico, March 7–11, 1983) = *Physica D* (special issue), 10/1–2.
- and Margolus, N. (1987), *Cellular Automata Machines* (Cambridge, Mass.: MIT Press).
- Ulam, S. (1962), 'On Some Mathematical Problems Connected with Patterns of Growth of Figures', *Proceedings of Symposia in Applied Mathematics*, 14: 215–24. Repr. in Burks (1970).
- Von Neumann, J. (1966), *Theory of Self-Reproducing Automata*, ed. and completed by A. W. Burks (Urbana, Ill.: University of Illinois Press).
- Wilson, S. W. (1989), 'The Genetic Algorithm and Simulated Evolution', in Langton (1989), 157–65.
- Wolfram, S. (1986), 'Cellular Automaton Fluids 1: Basic Theory', *Journal of Statistical Physics*, 45: 471–526.