If the Milieu is Reasonable: Lessons from Nature on Creating Life

Moshe Sipper
Logic Systems Laboratory, Swiss Federal Institute of Technology,
CH-1015 Lausanne, Switzerland. E-mail: Moshe.Sipper@di.epfl.ch.

Prologue

Anybody who looks at living organisms knows perfectly well that they can produce other organisms like themselves. This is their normal function, they wouldn't exist if they didn't do this, and it's plausible that this is the reason why they abound in the world. In other words, living organisms are very complicated aggregations of elementary parts, and by any reasonable theory of probability or thermodynamics highly improbable. That they should occur in the world at all is a miracle of the first magnitude; the only thing which removes, or mitigates, this miracle is that they reproduce themselves. Therefore, if by any peculiar accident there should ever be one of them, from there on the rules of probability do not apply, and there will be many of them, at least if the milieu is reasonable. (von Neumann, 1966)

If during the long course of ages and under varying conditions of life, organic beings vary at all in the several parts of their organisation, and I think this cannot be disputed; if there be, owing to the high geometrical powers of increase of each species, at some age, season, or year, a severe struggle for life, and this certainly cannot be disputed; then, considering the infinite complexity of the relations of all organic beings to each other and to their conditions of existence, causing an infinite diversity in structure, constitution, and habits, to be advantageous to them, I think it would be a most extraordinary fact if no variation ever had occurred useful to each being's own welfare, in the same way as so many variations have occurred useful to man. But if variations useful to any organic being do occur, assuredly individuals thus characterised will have the best chance of being preserved in the struggle for life; and from the strong principle of inheritance they will tend to produce offspring similarly characterised. This principle of preservation, I have called, for the sake of brevity, Natural Selection. (Darwin, 1859)

Genesis

Some three and a half billion years ago the first self-replicating molecules appeared on Earth, a humble beginning that marked the onset of the evolutionary avalanche that gave rise to Life, in all its glory, with the numerous species in existence today. Among the major scientific quests of this century, many are intimately tied to this process; these strike at the heart of our human essence, aiming to understand what life is, how it originated, and whence cometh the complex ecosystem which we daily witness with awe.

During the past decade, the traditional methods used to investigate such issues have been supplemented by novel ones in a field known as artificial life. This field of study is devoted to understanding life by attempting to abstract the fundamental dynamical principles underlying biological phenomena, and recreating these dynamics in other physical media – such as computers – making them accessible to new kinds of experimental manipulation and testing (Langton, 1992). While biological research is essentially *analytic*, trying to break down complex phenomena into their basic components, artificial life is *synthetic*, attempting to construct phenomena from their elemental units (Levy, 1992; Sipper, 1995). Adding powerful new tools to the scientific toolkit is, however, only part of the field's mission. As put forward by Langton (1992), in addition to providing new ways for studying biological phenomena associated with life here on Earth, *life-as-we-know-it*, artificial life lets us extend our studies to the larger domain of "bio-logic" of possible life, *life-as-it-could-be*.

These two parallel lines of research are by no means independent, indeed, the hope is that they will cross-fertilize each other. By studying organic life, we can perhaps gain insight on how to attain useful artificial systems, ranging from automated space explorers to autonomous vacuum cleaners, and, conversely, artificial life may enable the investigation of fundamental questions hitherto either to difficult, or downright impossible, to tackle.

One of these fundamental questions concerns the conditions under which life arises. If one examines our sole example, organic life, the minimal, *necessary* conditions seem to be: (1) the existence of self-replicating entities that are (2) subject to evolution (Kauffman, 1990; Ray, 1992). If these conditions are implemented in some artificial medium, will life arise? Though it is probably premature to provide a definite answer at this point, research carried out over the years has lent insight into the issues involved. This paper provides a glimpse into the questions that arise and some of their possible answers. We set out by describing the implementation of self-replication and evolution processes in artificial media. Then, we go on to observe that while these two conditions may be *necessary*, they are usually not *sufficient*. In other words, one must seek to understand the extra parameters needed to induce *evolvability*. And even if evolvability is attained, one must bear in mind the huge amount of resources that nature has been expending. Finally, we explore an issue that is related to the weak versus strong artificial life debate, the former position contending that observed phenomena are mere simulations, the latter advocating bona fide life. One possible approach to attaining strong artificial life is by building actual machines that can function in *our* environment, an idea that has recently emerged in the form of the nascent field of bio-inspired systems and evolvable hardware.

It is important to distinguish between two different questions: (1) what is life? and (2) what are the conditions under which life arises? While the two are obviously interlaced, our interest here lies with the latter. The existence of self-replicating entities subject to evolution, the minimal conditions discussed herein, can give rise to several characteristics cited by researchers in their definitions of life, including metabolism, teleonomic (purposeful) behavior, stability under perturbations, and more. This issue is beyond the scope of this paper, and the interested reader is referred to Bedau (1996), who summarizes different conceptions of life.

Self-replication

In the late 1940s eminent mathematician and physicist John von Neumann had become interested in the question of whether a machine can self-replicate, that is, produce copies of itself. Von Neumann wished to investigate the *logic* necessary for replication; he was not interested, nor did he have the tools, in building a working machine at the biochemical or genetic level. Remember that at the time DNA had not yet been discovered as the genetic material in nature.

To conduct a formal mathematical investigation of the issue, von Neumann used a model conceived by his colleague, mathematician Stanislaw Ulam. The model, known as a *cellular automaton*, consists of a large grid of cells (similar to a checkerboard), each possessing a certain state at a given moment. The number of possible states per cell is finite and is usually small (one can imagine each state being represented by a different color). All cells change state simultaneously such that the state of a cell at the next time step depends only on its state at the current time step and the states of its neighboring cells. The principle that guides state transformations is applied identically to all cells and is referred to as the *rule*. For example, a simple rule for a two-state (black/white) cellular automaton sets the state of a cell at the next time step to black if it has an even number of black neighbors, and to white if it has an odd number of black neighbors (Figure 1). Though very simply defined, cellular automata give rise to complex behavior.

A *machine* in the cellular automaton model is a collection of cells that can be regarded as operating in unison. Thus, one can observe simple "creatures" that are able to move within this austere universe, as demonstrated in Figure 2 for another well-known cellular automaton rule, the "game of life."

Von Neumann used this simple model to describe a universal constructing machine, which can read *assembly instructions* of *any* given machine, and construct that machine accordingly. These instructions are a collection of cells of various states, as is the new machine after being assembled – indeed, any compound element on the grid is simply a collection of cells. The universal constructor can build *any* machine given the appropriate "genome" or assembly instructions; thus, given its own description, it is capable of constructing a copy of itself, i.e., self-replicate. Should we want the offspring to self-replicate as well, we must *copy* the assembly instructions and attach them to it (Figure 3). In this manner, von Neumann showed that a replicative process is possible in artificial machines. The actual proof is quite elaborate and detailed in a book completed posthumously by von Neumann's colleague, Arthur Burks (von Neumann, 1966).

One of von Neumann's main conclusions was that the replicative process uses the assembly instructions in two distinct manners: as interpreted code (during actual assembly), and as uninterpreted data (copying of assembly instructions to offspring). During the following decade, when the basic genetic mechanisms began to unfold, it became clear that na-

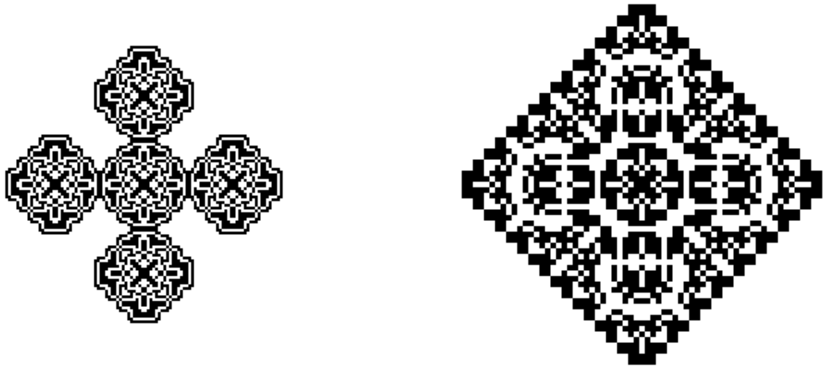

Figure 1: Patterns produced by a cellular automaton, where each cell can take on one of two states, represented by black and white squares. Starting from an arbitrary initial pattern, all cells change states simultaneously such that the state of a cell at the next time step depends only on its state at the current time step and the states of its four immediate neighbors (north, south, east, and west). Each cell in the grid follows the same simple rule which dictates that it becomes black if it has an even number of black neighbors, and white if it has an odd number of black neighbors. Shown above are the patterns generated after 90 time steps, that is, after all grid cells have undergone 90 state transformations (left), and after 120 time steps (right).

ture had "adopted" von Neumann's conclusions. The process by which assembly instructions (that is, DNA) are used to create a working machine (that is, proteins), indeed makes dual use of information: as interpreted code and as uninterpreted data. The former is referred to in biology as *translation*, the latter as *transcription*.

A major problem with von Neumann's approach (as well as his successors Banks (1970), Burks (1970), Codd (1968), and Pesavento (1995)) is the complexity of the constructor, which requires hundreds of thousands of cells. In addition, each cell can be in one of 29 states rather than just two. In 1984 Christopher Langton observed that although the capacity for universal construction is a *sufficient* condition for self-replication, it is not a *necessary* one; furthermore, natural systems are not capable of universal construction. Langton (1984) and his successors Byl (1989), Reggia et al. (1993), and Morita and Imai (1997) developed self-replicating automata which are much simpler than the universal constructor. These machines, however, lack any computing and constructing capabilities, their sole functionality being that of self-replication. Building upon these results, a number of researchers have recently demonstrated that one can attain self-replicating structures that can execute useful programs (Perrier et al., 1996; Tempesti, 1995). These latter machines are simpler than von Neumann's universal constructor, yet are able to self-replicate, with each "daughter" organism able to continue replicating in its turn, as well as to carry out other functions written into its genetic code.

Cellular automata exhibit what is known as *emergent behavior*. This term refers to the appearance of global information-processing capabilities that

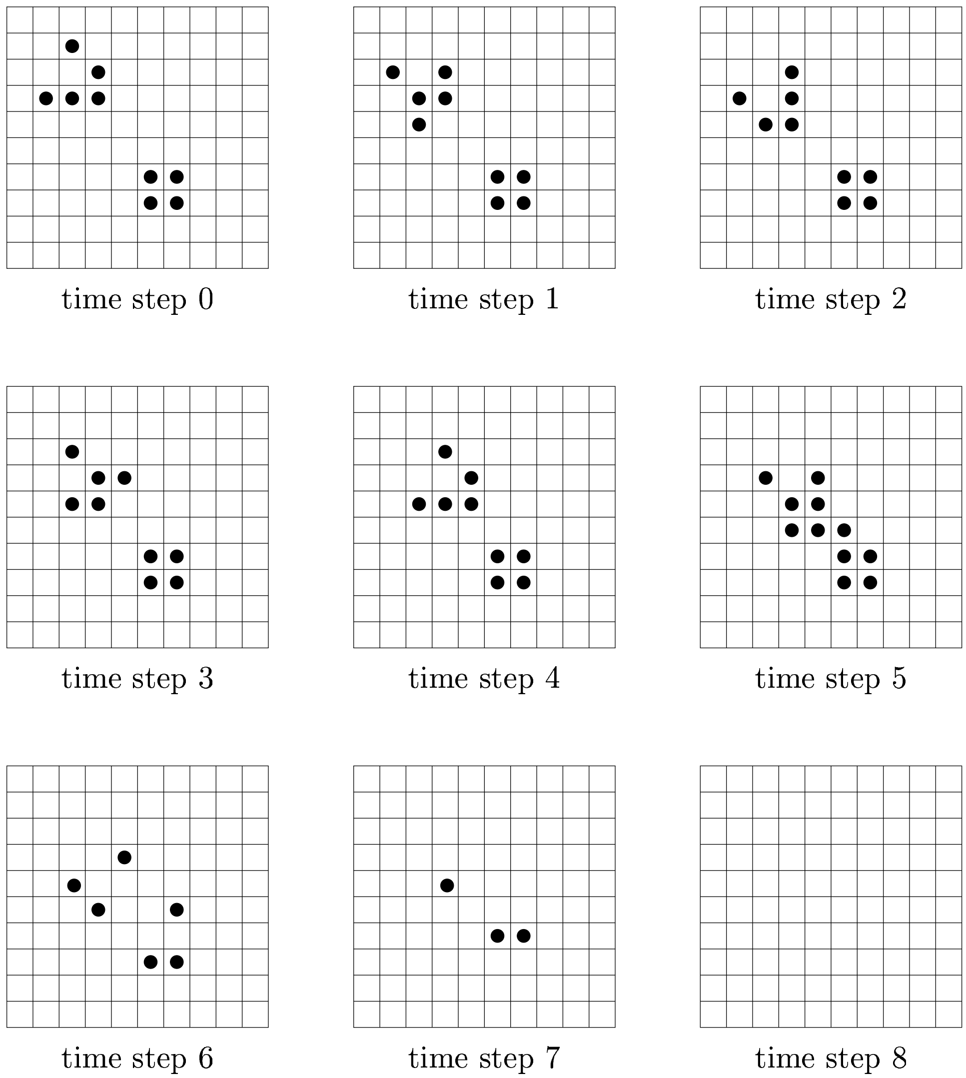

Figure 2: The "game of life" cellular automaton was defined by John H. Conway in 1968: "Life occurs on a virtual checkerboard. The squares are called cells. They are in one of two states: alive or dead. Each cell has eight possible neighbors, the cells which touch its sides or its corners. If a cell on the checkerboard is alive, it will survive in the next time step if there are either two or three neighbors also alive. It will die of overcrowding if there are more than three live neighbors, and it will die of exposure if there are fewer than two. If a cell on the checkerboard is dead, it will remain dead in the next time step unless exactly three of its eight neighbors are alive. In that case, the cell will be "born" in the next time step" (Berlekamp et al., 1982; Levy, 1992). Shown above at time step 0 are two patterns, a stationary square, known as a "block," and a moving creature, known as a "glider." The latter displaces itself one square diagonally every four time steps. Upon meeting other creatures, such as the block, both are subject to this simple universe's basic rule, or "laws of physics," which cause their mutual annihilation in this case. Note that actions, such as "movement," are purely in the eyes of the beholder – the most basic level consists of simple state transformations, with no movement at all. However, by considering higher-level views of the system, one can introduce new terms to describe the perceived phenomena.

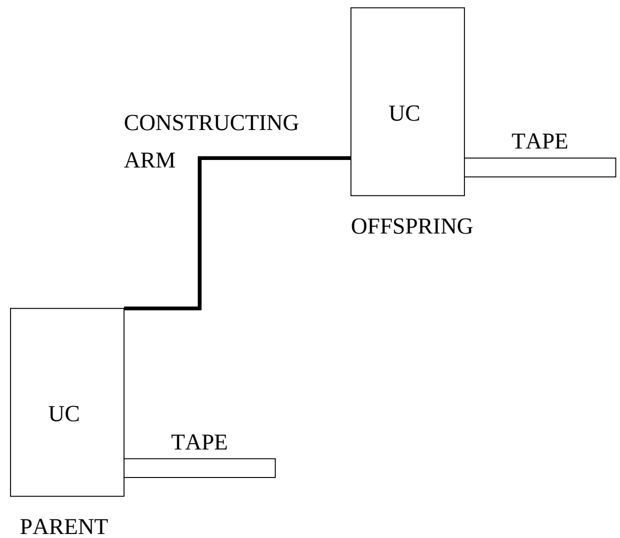

Figure 3: A schematic diagram of von Neumann's self-replicating cellular automaton. The machine is a universal constructor (UC) capable of constructing, through the use of a "constructing arm," any configuration whose description (genome) can be stored on its input tape. Thus, given its own description, the machine is capable of constructing a copy of itself, i.e., self-replicate.

are not explicitly represented in the system's elementary components or in their interconnections (Crutchfield and Mitchell, 1995; Forrest, 1991; Sipper, 1997a). As put forward by Steels (1994), a system's behavior is emergent if it can only be defined using descriptive categories that are not necessary to describe the behavior of the constituent components. In physical systems, temperature and pressure are examples of emergent phenomena – they occur in large ensembles of molecules and are due to interactions at the molecular level, though an individual molecule by itself possesses neither temperature nor pressure. Nature abounds in such systems in which the actions of simple, locally-interacting components give rise to coordinated global behavior. Examples include insect colonies, cellular assemblies, the retina, and the immune system. Cellular automata are one of the oft-used models for studying emergence, since observed phenomena arise as a result of applying the local "laws of physics," that is, the local rule by which state transformations take place. Interestingly, it has been suggested by Ed Fredkin that our own universe is a cellular automaton (Levy, 1992).

An interesting point concerning the cellular automata defined above is their deterministic mode of operation – a given initial pattern of states always leads to the same development in time, that is, to the same sequence of patterns. However, probabilistic cellular automata have also been treated, where cellular state updates take place in a probabilistic manner (Toffoli and Margolus, 1987). This leads to non-determinism and to what might be considered the equivalent of quantum-mechanical effects.

Interlude: Permutation City

The Autoverse was a "toy" universe, a computer model which obeyed its own simplified "laws of physics" – laws far easier to deal with mathematically than the equations of real-world quantum mechanics.

... Autoverse was a vast array of cubic cells... the laws governing individual cells drove everything that happened at higher levels... The cellular automaton which was the Autoverse did nothing whatsoever but apply these rules uniformly to every cell; these were its fundamental "laws of physics." Here, there were no daunting quantum-mechanical equations to struggle with – just a handful of trivial arithmetic operations, performed on integers. And yet the impossibly crude laws of the Autoverse still managed to give rise to "atoms" and "molecules" with a "chemistry" rich enough to sustain "life."

Max Lambert's original reason for designing the Autoverse had been the hope of observing self-replicating molecular systems – primitive life – arising from simple chemical mixtures.

They'd invented a new physics... Everything was driven from the bottom up, by the lowest level of physical laws, just as it was in the real world. The price of this simplicity was that an Autoverse bacterium didn't necessarily behave like its real-world counterparts.
(Egan, 1994).

Evolution

The only process currently known to have produced an ecosystem of living creatures, and in particular, of intelligent beings, is that of natural evolution. Darwin (1859) laid out the core of the currently accepted theory of evolution, its major elements being (Ray, 1994a):

- Individuals vary in their viability in the environments that they occupy.
- This variation is heritable.
- Self-replicating individuals tend to produce more offspring than can survive on the limited resources available in the environment.
- In the ensuing struggle for survival, the individuals best adapted to the environment are the ones that will survive to reproduce.

The continual workings of this process over the millenia causes populations of organisms to change, generally becoming better adapted to their environment.

The issue of self replication was discussed in the previous section, and we now turn our attention to the second issue, namely, evolution. The idea of applying the biological principle of natural evolution to artificial systems was introduced in the 1950s and the 1960s, when several researchers studied evolutionary systems with the idea that evolution could be used as an optimization tool for engineering problems. Central to all the different methodologies is the notion of solving problems by evolving an initially random population of candidate solutions, through the application of operators inspired by natural genetics and natural selection, such that in time "fitter" (i.e., better) solutions emerge (Holland, 1975; Koza, 1992; Michalewicz, 1996; Mitchell, 1996). Nowadays, these so-called evolutionary algorithms are ubiquitous, having been successfully applied to numerous problems from different domains, including optimization, automatic programming, machine learning, economics, operations research, immune systems, ecology, population genetics, studies of evolution and learning, and social systems (Mitchell, 1996). An interesting line of research that has recently emerged upon the scene is that of evolving cellular automata (Mitchell et al., 1994; Mitchell, 1996; Sipper, 1996; Sipper, 1997a; Sipper, 1997b).

In order to execute an evolutionary algorithm one creates, at random, a *population* of individuals. Each individual represents a possible solution to an a priori given problem, and is often represented by a simple string of characters. This initial population is referred to as the first generation. The following generations are then formed by evolution so that in time the population comes to consist of better (fitter) individuals. Each new generation is created by selecting good parents from the previous generation that are then subjected to "genetic" operators, such as crossover and mutation. A parent that is better at solving the problem at hand stands a better chance of being selected, thereby having its genes remain in the gene pool. Crossover creates a new individual from two parents by mixing their genetic material, while mutation introduces a small amount of copying errors (Figure 4). The selection-crossover-mutation process continues until the next generation is formed, in (abstract) analogy to nature: a given generation consists of different individuals whose chances of survival stand in relation to their fitness – the better (fitter) an individual, the higher its probability of survival, and in time the population's overall fitness increases.

A fundamental issue that must be addressed is what comprises a good (fit) individual within the population of artificial organisms. This question, while highly complex in nature, has a simpler answer in the context of evolutionary computation – fitness is imposed externally by an outside observer, in accordance with the particular problem at hand. For example, if the goal is to design an airplane wing, then the population might consist of individual designs, initially created at random. The fitness in this case will reflect how well an individual wing design performs in some kind of test environment, for example, a wind tunnel (or a computer simulation of such a tunnel). Obviously, the first-generation individuals will exhibit poor performance, however, there will be some variability among them. Even such minute differences can be picked up by evolution and enhanced to ultimately produce a good solution (in our example, a wing design that can be used in an actual aircraft). Note that evolution proceeds without human intervention – after the environment is set, an initial population is generated at random and evolution treads along until a satisfying solution is found. A prime advantage of evolutionary methods from an engineering standpoint is the potential for adaptability – when an unforeseen event occurs, the system can evolve, that is, adapt to the new situation, in analogy to nature.

A major difference between natural and artificial evolution pertains to the issue of open-endedness. When the fitness criterion is imposed by the user in accordance with the task to be solved (currently the rule with artificial evolution techniques), one attains a form of *guided*, or *directed* evolution. This is to be contrasted with *open-ended* evolution occurring in nature, which admits no externally imposed fitness criterion, but rather an implicit, emergent, dynamical one (that could arguably be summed up as survivability). A significant aspect of the natural environment to which any given organism must adapt is all the other organisms with which it interacts; this is often referred to as coevolution. Darwin (1859) himself wrote "of the coadaptations of organic beings to each

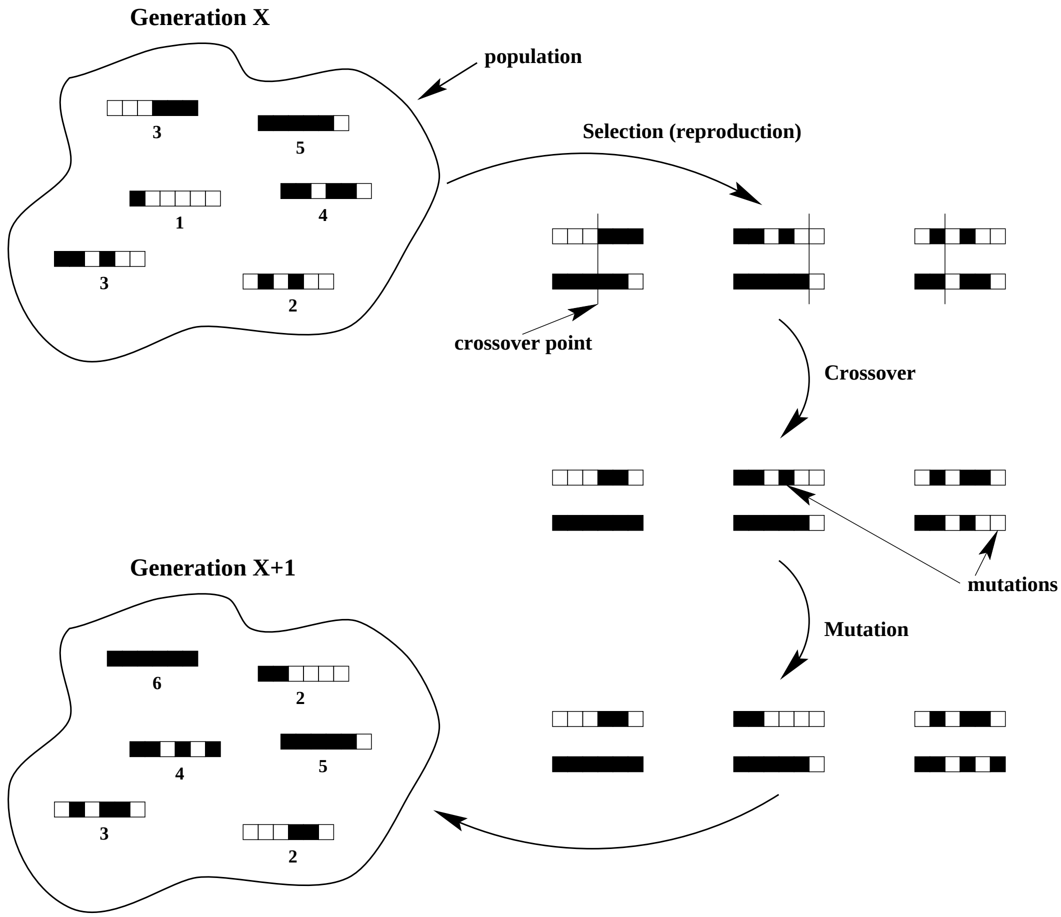

Figure 4: Demonstration of an evolutionary algorithm over one generation. The population consists of six individuals, each one represented by an artificial genome containing six genes. A gene can take on one of two values (marked by black and white boxes). In this simple example, the fitness of an individual equals the number of black boxes (genes) in its genome (fitness values are displayed below the genomes). Selection (reproduction) is performed probabilistically: the higher an individual's fitness, the better its chance of being selected. Thus, some parents get selected more than once while others not at all. Each selected pair of parents is recombined to produce two offspring, an operation known as crossover. This is done by exchanging all genes to the right of a randomly selected crossover point. Mutation is then applied with low probability by simply flipping the gene's value. Note that application of the genetic operators on the population of generation X has yielded a perfect individual, with a fitness value of 6, at generation $X+1$. Furthermore, the average fitness of the population, computed over all individuals, has increased.

other." Open-ended, undirected evolution is the only form of evolution known to produce such devices as eyes, wings, and nervous systems, and to give rise to the formation of species. Undirectedness may have to be applied to artificial evolution if we want to observe the emergence of completely novel systems.

The question of whether open-ended evolution can be embedded within a computer was posed by Thomas Ray who devised a virtual world called *Tierra*, consisting of computer programs that can undergo evolution (Ray, 1992). In contrast to evolutionary computation where fitness is defined by users, the Tierra "creatures" (programs) receive no such direction. Rather, they compete for the natural resources of their computerized environment, namely CPU time and memory. Since only a finite amount of these are available, the virtual world's natural resources are limited, as in nature, serving as the basis for competition between creatures. Ray inoculated his system with a single, self-replicating organism, called the "Ancestor," which is the only engineered (man-made) creature in Tierra. He then set his system loose and witnessed the emergence of an ecosystem within the Tierra world, including organisms of various sizes, parasites, hyper-parasites, and so on. The evolved parasites, for example, are small creatures that use the replication code of larger organisms (such as the Ancestor) to self-replicate. In this manner they proliferate rapidly without the need for the excess replication code. Ray has recently extended his Tierra environment to run on the Internet, rather than on a single computer, hoping that by increasing the scale of the system new phenomena may arise that have not been observed on a single computer (Ray, 1994b; Ray, 1996).

Interlude: Permutation City

"I want you to construct a seed for a biosphere... I want you to design a prebiotic environment – a planetary surface, if you'd like to think of it that way – and one simple organism which you believe would be capable, in time, of evolving into a multitude of species and filling all the potential ecological niches."

"Think of this whole project as... A sketch of a proof."
"A proof of what?"
"That Autoverse life could – in theory – be as rich and complex as life on Earth."
(Egan, 1994).

Self-replication + Evolution = Life ?

We have established so far that self-replication and evolution can be attained in artificial systems. Though there is much yet to be explored along these lines, an even more basic question arises: while these two factors seem to comprise a *necessary* condition for the creation of life, and in particular of biodiversity, are they *sufficient*? In other words, what does it take to induce *evolvability*? Consider, for example, computer programs written in a standard programming language such as FORTRAN or C. A brute-force attempt to apply an evolutionary process that randomly reshuffles such programs, would, in most cases, lead nowhere, churning out a plethora of non-functional systems (Kauffman, 1990). (In fact, it is possible to evolve computer programs by an evolutionary methodology known as genetic programming (Koza, 1992). However, this requires meticulously setting the evolutionary scenario, including the genomic encoding of programs, and the genetic operators applied to them.)

Ray (1994a) studied the evolvability issue within the Tierra framework by creating four versions of his artificial world. The four differed in the basic machine language which comprises the underlying genetic system. Though the differences were subtle they nonetheless gave rise to a wide variability in terms of rates, degrees, and patterns of evolution. Dawkins (1989) looked into this question by employing a simple environment of evolving organisms, denoted biomorphs. He too noted that tweaking this simple world's genetic system has a large influence on evolvability, that is, on the possible evolutionary outcomes. Finally, Kauffman (1990; 1995) used a model known as random boolean networks to study the requirements for evolvability in complex systems. For now, it seems that we are still far from possessing a definite answer to this question, which could provide us with clear guidelines for constructing artificial worlds.

Even if evolvability is attained, one must bear in mind the huge amount of resources that nature has been expending, in terms of time, area size, and population size. Darwin (1859) himself had noted the importance of these elements, writing, on the time factor: "The mind cannot possibly grasp the full meaning of the term of a hundred million years; it cannot add up and perceive the full effects of many slight variations, accumulated during an almost infinite number of generations." On area size: "... on the whole I am inclined to believe that largeness of area is of more importance..." And, on population size: "A large number of individuals... is, I believe, an extremely important element of success." Darwin remarked that the evolutionary process goes on "... for millions on millions of years; and during each year on millions of individuals of many kinds..." As beautifully put by him "nature is prodigal in variety, though niggard in innovation" – the huge amount of resources expended thereby enabling, nonetheless, progression.

Interlude: Permutation City

Then, changing the length scale by a factor of a million, she started up twenty-one tiny cultures of Autobacterium lamberti...

Arranging for A. lamberti *to mutate was easy; like a real-world bacterium, it made frequent errors every time it duplicated its analogue of DNA. Persuading it to mutate "usefully" was something else. Max Lambert himself – inventor of the Autoverse, creator of* A. lamberti... – *had spent much of the last fifteen years of his life trying to discover why the subtle differences between real-world and Autoverse biochemistry made natural selection so common in one system, and so elusive in the other. Exposed to the kind of stressful opportunities which* E. coli *would have exploited within a few dozen generations, strain after strain of* A. lamberti *had simply died out.*

"What exactly do you mean by a 'planetary environment'?"
"Whatever you think is reasonable. Say – thirty million square kilometers?"

"We've given the Autoverse a lot of resources; seven thousand years, for most of us, has been about three billion for Planet Lambert." (Egan, 1994).

The Real McCoy

So far we have discussed two major factors in our planet's natural history, self-replication and evolution, considering their implementation in artificial media. Up until now we did not discuss the actual media used, and in point of fact all the works discussed above were implemented as simulations on a general-purpose computer, such as a PC or a workstation. This raises a fundamental issue, namely, the distinction between so-called weak artificial life, where observed phenomena are considered to be mere simulations, and strong artificial life, involving bona fide life. The latter's goal was stated by Langton (1986): "We would like to build models that are so life-like that they cease to be *models* of life and become *examples* of life themselves." One of the basic tenets of strong artificial life is that veritable phenomena are observed, the underlying belief being that life is not necessarily carbon-based. The sole implication of the term 'artificial' is that the systems in question are man-made, that is, the basic components were not created by nature through evolution. However, the higher-level phenomena are completely genuine. Thus, it is argued that the replicative process exhibited by self-replicating automata is as real as that carried out in nature, the difference resting solely in the basic components: artificial cells versus live ones. The reader wishing to learn more on the issue of strong versus weak artificial life is referred to Boden (1996) and Langton et al. (1992).

One possible approach to attaining strong artificial life is by building actual machines that can function in *our* environment. This idea, whose origins can be traced to the cybernetics movement of the 1940s and the 1950s, has recently resurged in the form of the nascent field of bio-inspired systems and evolvable hardware. The field draws on ideas from the evolutionary computation domain as well as on novel hardware innovations (Sanchez and Tomassini, 1996). Sipper et al. (1997) have recently introduced the POE model for classifying bio-inspired hardware along three axes, inspired by three levels of organization observed in nature: **P**hylogeny, **O**ntogeny, and **E**pigenesis (Figure 5a):

Phylogeny: The first level concerns the temporal evolution of the genetic program, the hallmark of which is the evolution of species, or *phylogeny*.

Ontogeny: Upon the appearance of multicellular organisms, a second level of biological organization manifests itself. The successive divisions of the mother cell, the zygote, with each newly formed cell possessing a copy of the original genome, is followed by a specialization of the daughter cells in accordance with their surroundings, i.e., their position within the ensemble. This latter phase is known as cellular differentiation. *Ontogeny* is thus the developmental process of a multicellular organism.

Epigenesis: The ontogenetic program is limited in the amount of information that can be stored, thereby rendering the complete specification of the organism impossible. (A well-known example is that of the human brain with some 10^{10} neurons and 10^{14} connections, far too large a number to be completely specified in the four-character genome of length approximately 3×10^9.) Therefore, upon reaching a certain level of complexity, there must emerge a different process that permits the individual to integrate the vast quantity of interactions with the outside world. This process is known as *epigenesis*, and primarily includes the nervous system, the immune system, and the endocrine system.

The distinction between the axes cannot be easily drawn where nature is con-

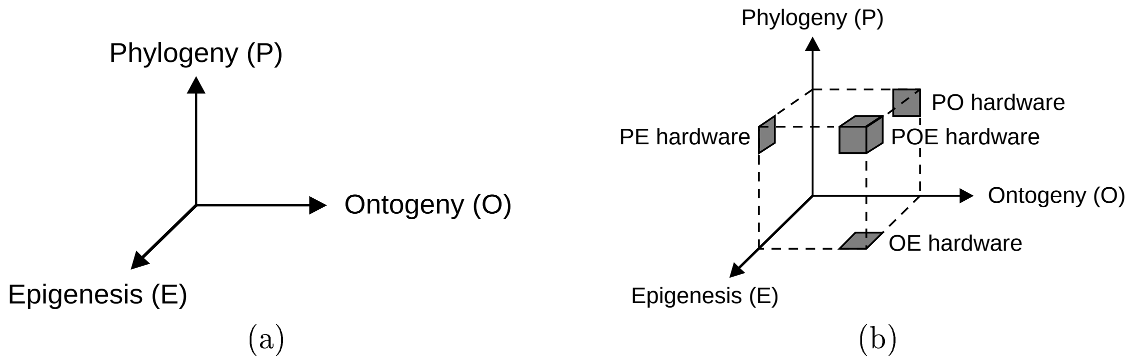

Figure 5: The POE model. (a) Partitioning the space of bio-inspired hardware systems along three axes: phylogeny, ontogeny, and epigenesis. (b) In the future we may witness the oncoming of systems that are situated along two, and ultimately all three axes. For example, a system that exhibits self-replication and evolution would be situated within the PO plane. Given time, it could possibly "infiltrate" the epigenetic axis, through the evolution of, say, some form of learning mechanism.

cerned, indeed the definitions themselves may be subject to discussion. Sipper et al. (1997) therefore *defined* each of the above axes *within the framework* of the POE model as follows: the phylogenetic axis involves *evolution*, the ontogenetic axis involves the *development* of a single individual from its own genetic material, essentially without environmental interactions, and the epigenetic axis involves *learning* through environmental interactions that take place after formation of the individual. Sipper et al. (1997) presented an overview of the nascent field of bio-inspired systems, describing actual hardware implementations situated along each of these axes. In particular, hardware realizations of self-replicating (ontogenetic) and evolving (phylogenetic) systems, such as those described in this paper, have been demonstrated (Sipper et al., 1997; Sipper, 1997a). Sipper et al. (1997) also presented an outlook on the field's possible future development, concluding that we may witness the oncoming of systems that are situated along two, and ultimately all three axes. For example, a system that exhibits self-replication and evolution would be situated within the PO plane. Given time, it could possibly "infiltrate" the epigenetic axis, through the evolution of, say, some form of learning mechanism (Figure 5b).

Looking (and dreaming) toward the future, one can imagine nano-scale (bioware) systems becoming a reality, which will be endowed with evolutionary, reproductive, regenerative, and learning capabilities. Such systems could give rise to novel species which will coexist alongside carbon-based organisms.

Interlude: Permutation City

The catch was, if a molecule obeyed only Autoverse physics – the internal logic of the self-contained computer model – then how could she, outside the model, interact with it at all? By constructing little surrogate hands in the Autoverse, to act as remote manipulators?... Manipulating the contents of the Autoverse meant violating its laws. (Egan, 1994).

Elephants Don't Play Chess

Evolution on Earth has given rise not only to the plethora of species in existence today, but also to intelligent beings. According to Brooks (1990), an examination of the evolution of life on Earth reveals that most of the time was spent developing basic intelligence. The elemental faculties evolved enable mobility in a dynamic environment and sensing of the surroundings to a degree sufficient to achieve the necessary maintenance of life and reproduction. This point pertains to an essential difference between artificial life and artificial intelligence (AI). Whereas AI has traditionally concentrated on complex human functions such as chess playing, text comprehension, medical diagnosis, and so on, artificial life concentrates on basic natural behaviors, emphasizing survivability in complex environments. The issues dealt with by AI appeared only very recently on the evolutionary scene (a mere few thousand years) and mostly in humans. This suggests that problem-solving behavior, language, expert knowledge, and reason are all rather simple once the essence of being and reacting is available. The idea is expressed in the title of one of Brooks' papers, *Elephants Don't Play Chess* (Brooks, 1990), suggesting that these creatures are nonetheless highly intelligent, able to survive and reproduce in a complex, dynamic environment.

In conclusion, we began by observing two factors that played a key role in the formation of species in the organic world: self-replication and evolution. Following nature's example, we explored the possibility of implementing these two processes in artificial media, as computer simulations, or as actual hardware that interacts with the real world. The years to come will see what wonders shall sprout from our artificial seeds.

Epilogue: Permutation City

Maria was elated, and a little dazed. People had been trying to achieve a spontaneous adaptation like this for sixteen years. *She didn't even know why she'd finally succeeded...*

"There are six hundred and ninety million species currently living on Planet Lambert. All obeying the laws of the Autoverse. All demonstrably descended from a single organism which lived three billion years ago – and whose characteristics I expect you know by heart. Do you honestly believe that anyone could have designed *all that?" (Egan, 1994).*

Acknowledgments

I am grateful to Daniel Mange and Eduardo Sanchez for their valuable comments and suggestions.

⋆ Further information on the issues discussed in this paper can be obtained via:
`http://lslwww.epfl.ch/~moshes/caslinks.html`

References

Banks, E. R. 1970. Universality in cellular automata. In *IEEE 11th Annual Symposium on Switching and Automata Theory*, pp 194–215. Santa Monica, California.

Bedau, M. A. 1996. The nature of life. In M. A. Boden (ed.), *The Philosophy of Artificial Life*, Chapt. 13, pp 332–357. Oxford University Press, New York.

Berlekamp, E. R., Conway, J. H., and Guy, R. K. 1982. *Winning Ways for your Mathematical Plays*, vol. 2, Chapt. 25, pp 817–850. Academic Press, New York.

Boden, M. A. (ed.) 1996. *The Philosophy of Artificial Life.* Oxford University Press, New York.

Brooks, R. A. 1990. Elephants don't play chess. *Robotics and Autonomous Systems*, vol. 6, no. 1-2, pp 3–15.

Burks, A. (ed.) 1970. *Essays on Cellular Automata.* University of Illinois Press, Urbana, Illinois.

Byl, J. 1989. Self-reproduction in small cellular automata. *Physica D*, vol. 34, pp 295–299.

Codd, E. F. 1968. *Cellular Automata.* Academic Press, New York.

Crutchfield, J. P. and Mitchell, M. 1995. The evolution of emergent computation. *Proceedings of the National Academy of Sciences USA*, vol. 92, no. 23, pp 10742–10746.

Darwin, C. 1859. *The Origin of Species.* John Murray, London.

Dawkins, R. 1989. The evolution of evolvability. In C. G. Langton (ed.), *Artificial Life*, vol. VI of *SFI Studies in the Sciences of Complexity*, pp 201–220. Addison-Wesley.

Egan, G. 1994. *Permutation City.* Harper, New York.

Forrest, S. (ed.) 1991. *Emergent Computation: Self-organizing, Collective, and Cooperative Phenomena in Natural and Artificial Computing Networks.* The MIT Press, Cambridge, MA.

Holland, J. H. 1975. *Adaptation in Natural and Artificial Systems.* The University of Michigan Press, Ann Arbor, Michigan.

Kauffman, S. A. 1990. Requirements for evolvability in complex systems: Orderly dynamics and frozen components. *Physica D*, vol. 42, pp 135–152.

Kauffman, S. A. 1995. *At Home in the Universe: The Search for Laws of Self-Organization and Complexity.* Oxford University Press, New York.

Koza, J. R. 1992. *Genetic Programming.* The MIT Press, Cambridge, Massachusetts.

Langton, C. G. 1984. Self-reproduction in cellular automata. *Physica D*, vol. 10, pp 135–144.

Langton, C. G. 1986. Studying artificial life with cellular automata. *Physica D*, vol. 22, pp 120–149.

Langton, C. G. 1992. Preface. In C. G. Langton, C. Taylor, J. D. Farmer, and S. Rasmussen (eds.), *Artificial Life II*, vol. X of *SFI Studies in the Sciences of Complexity*, pp xiii–xviii. Addison-Wesley, Redwood City, CA.

Langton, C. G., Taylor, C., Farmer, J. D., and Rasmussen, S. (eds.) 1992. *Artificial Life II*, vol. X of *SFI Studies in the Sciences of Complexity.* Addison-Wesley, Redwood City, CA.

Levy, S. 1992. *Artificial Life: The Quest for a New Creation.* Random House.

Michalewicz, Z. 1996. *Genetic Algorithms + Data Structures = Evolution Programs.* Springer-Verlag, Heidelberg, third edition.

Mitchell, M. 1996. *An Introduction to Genetic Algorithms.* MIT Press, Cambridge, MA.

Mitchell, M., Crutchfield, J. P., and Hraber, P. T. 1994. Evolving cellular automata to perform computations: Mechanisms and impediments. *Physica D*, vol. 75, pp 361–391.

Morita, K. and Imai, K. 1997. Logical universality and self-reproduction in reversible cellular automata. In *Proceedings of The First International Conference on Evolvable Systems: From Biology to Hardware (ICES96)*, Lecture Notes in Computer Science. Springer-Verlag, Heidelberg. (to appear).

Perrier, J.-Y., Sipper, M., and Zahnd, J. 1996. Toward a viable, self-reproducing universal computer. *Physica D*, vol. 97, pp 335–352.

Pesavento, U. 1995. An implementation of von Neumann's self-reproducing machine. *Artificial Life*, vol. 2, no. 4, pp 337–354.

Ray, T. S. 1992. An approach to the synthesis of life. In C. G. Langton, C. Taylor, J. D. Farmer, and S. Rasmussen (eds.), *Artificial Life II*, vol. X of *SFI Studies in the Sciences of Complexity*, pp 371–408. Addison-Wesley, Redwood City, CA.

Ray, T. S. 1994a. Evolution, complexity, entropy and artificial reality. *Physica D*, vol. 75, pp 239–263.

Ray, T. S. 1994b. *A Proposal to Create a Network-Wide Biodiversity Reserve for Digital Organisms.* unpublished. See also: *Science*, vol. 264, May, 1994, page 1085.

Ray, T. S. 1996. A network-wide biodiversity reserve for digital organisms. In *Imagina 96 Proceedings*, pp 15–26. Institut National de l'Audiovisuel, Bry-sur-Marne, France. See also: Continuing report on the network Tierra experiment at: http://www.hip.atr.co.jp/~ray /tierra/netreport/netreport.html.

Reggia, J. A., Armentrout, S. L., Chou, H.-H., and Peng, Y. 1993. Simple systems that exhibit self-directed replication. *Science*, vol. 259, pp 1282–1287.

Sanchez, E. and Tomassini, M. (eds.) 1996. *Towards Evolvable Hardware*, vol. 1062 of *Lecture Notes in Computer Science*. Springer-Verlag, Heidelberg.

Sipper, M. 1995. An introduction to artificial life. *Explorations in Artificial Life (special issue of AI Expert)* pp 4–8. Miller Freeman, San Francisco, CA.

Sipper, M. 1996. Co-evolving non-uniform cellular automata to perform computations. *Physica D*, vol. 92, pp 193–208.

Sipper, M. 1997a. *Evolution of Parallel Cellular Machines: The Cellular Programming Approach*. Springer-Verlag, Heidelberg.

Sipper, M. 1997b. The evolution of parallel cellular machines: Toward evolware. *BioSystems*, vol. 42, pp 29–43.

Sipper, M., Sanchez, E., Mange, D., Tomassini, M., Pérez-Uribe, A., and Stauffer, A. 1997. A phylogenetic, ontogenetic, and epigenetic view of bio-inspired hardware systems. *IEEE Transactions on Evolutionary Computation*, vol. 1, no. 1.

Steels, L. 1994. The artificial life roots of artificial intelligence. *Artificial Life Journal*, vol. 1, no. 1/2, pp 75–110. The MIT Press, Cambridge, MA.

Tempesti, G. 1995. A new self-reproducing cellular automaton capable of construction and computation. In F. Morán, A. Moreno, J. J. Merelo, and P. Chacón (eds.), *ECAL'95: Third European Conference on Artificial Life*, vol. 929 of *Lecture Notes in Computer Science*, pp 555–563. Springer-Verlag, Heidelberg.

Toffoli, T. and Margolus, N. 1987. *Cellular Automata Machines*. The MIT Press, Cambridge, Massachusetts.

von Neumann, J. 1966. *Theory of Self-Reproducing Automata*. University of Illinois Press, Illinois. Edited and completed by A. W. Burks.

Moshe Sipper is a senior researcher in the Logic Systems Laboratory at the Swiss Federal Institute of Technology, Lausanne, Switzerland. He received a B.A. in Computer Science from the Technion – Israel Institute of Technology, an M.Sc. and a Ph.D. from Tel Aviv University. His chief interests involve the application of biological principles to artificial systems, including evolutionary computation, cellular automata (with an emphasis on evolving cellular machines), bio-inspired systems, evolving hardware, complex adaptive systems, artificial life, and neural networks. Dr. Sipper has authored and co-authored several scientific papers in these areas, as well as his recent book *Evolution of Parallel Cellular Machines: The Cellular Programming Approach* (Heidelberg: Springer-Verlag, 1997). He was co-organizer of a special session entitled "Toward Evolware," held as part of the IEEE International Conference on Evolutionary Computation (ICEC'97), and is Program Chairman of the Second International Conference on Evolvable Systems: From Biology to Hardware (ICES98), to be held in Lausanne in September 1998.