

Evolutionary Automata

Murray Shanahan

Imperial College
Department of Computing,
180 Queen's Gate,
London SW7 2BZ, England.
Email: mps@doc.ic.ac.uk

Abstract

The class of Evolutionary Automata is formally defined. An Evolutionary Automaton generates an evolving population of “organisms” within a computational medium. The organisms move randomly around a microworld, eat, reproduce asexually (with mutation), and die of “natural causes”. The results of some preliminary experiments conducted with an instance of this class are reported, in which a parameter was varied across a small range. The paper concludes with some speculation on the possible significance of these results for the study of diversity and complexity in Artificial Life.

Introduction

One of the most attractive features of the manifesto of Artificial Life research is that it legitimises the study of any life-like phenomenon that can be reproduced in a computer. Artificial Life research avoids the prolific detail present in real-life biological phenomena, but at the same time enjoys the freedom to explore a larger space of possible phenomena than actually arise in nature [Langton, 1987]. This wider perspective encourages the asking of deep scientific questions, such as “What preconditions are essential to the origin of all life-like phenomena?”

Unsurprisingly, evolution is one of the phenomena that has attracted most attention from Artificial Life researchers. Work on evolution currently falls into two main categories. First, some impressive results have been obtained through the simulation of evolutionary processes in microworlds, some of which are deliberate attempts to model aspects of biological reality (for example, [Yaeger, 1992]), whilst others are unashamedly artificial (for example, [Ray, 1990]). Second, a number of computer experiments have been conducted which are based on abstract automata, in particular cellular automata [Packard, 1988], [Langton, 1990], and boolean networks [Kaufmann, 1992, Chapters 5 & 6]. These experiments are not so much an attempt to model particular evolutionary phenomena in detail. Rather, their emphasis is on deepening our theoretical understanding of such processes in general.

This paper falls between these two categories. A class of Evolutionary Automata is described, in which a population of entities evolves in a microworld. This class is large — in the sense that different instances of it can reproduce a wide variety of effects — yet abstract enough for it to be susceptible to simple mathematical description and tractable formal analysis. (Holland's Echo model [1992, Chapter 10]

is similarly motivated.) The point of studying such a class is that it enables us to investigate deep scientific questions, such as “Under what conditions does evolution produce complexity and diversity?” using a tool that models evolution with a fine enough grain of detail for complexity and diversity actually to arise.

The first section of the paper presents the basic idea of an Evolutionary Automaton and a formal definition. The next two sections present some experiments with a particular automaton, in which the effect on evolution of varying one of the automaton's parameters is studied. It turns out that this parameter must be tuned to be close to a critical value for the automaton to produce interesting effects. The paper concludes with some discussion of the question of the conditions that are required for complexity and diversity to emerge from an evolutionary process. The potential relevance of the edge of chaos to this question is discussed, but in the light of the reported experiments its applicability seems minimal.

1. The Class of Evolutionary Automata

After some informal characterisation, this section formally defines the class of Evolutionary Automata (EA's). Essentially, an Evolutionary Automaton non-deterministically specifies a sequence of states of a microworld. The microworld is a square grid of locations. Both top and bottom edges and left and right edges are joined so that the microworld is, in effect, shaped like a torus. An evolving population of “organisms” inhabits the microworld. These organisms can move, eat, reproduce, and die of “natural causes.”

Each organism occupies a number of connected locations in the microworld. An organism's whereabouts in the microworld is constantly changing, but its shape (or phenotype) remains unchanged throughout its lifetime. Each organism has a genotype, which can be thought of as defining the “species” to which it belongs. Its shape is a function of its genotype. Organisms can be either alive or dead. Dead organisms don't move, eat or reproduce, and can be eaten by any other organism except members of their own species.¹ Each organism has a certain energy level. Eating increases its energy level and reproduction depletes it. There are four kinds of microworld event.

¹ Forbidding cannibalism prevents the establishment of species which are completely independent from the rest of their environment, and so encourages the setting up of food webs. I have built this restriction into the definition of an EA, but it could be relegated to the descriptions of particular automata.

- **Movements:** At each increment of the microworld clock, every living organism moves a random distance in a random direction, up to a certain maximum distance away.
- **Meals:** Under certain circumstances one organism can eat another, thus gaining a number of units of energy. Organisms have to be within a certain distance from each other for one of them to eat the other.
- **Births:** When an organism has sufficient units of energy, it will reproduce. Reproduction is asexual, but random mutations to the organism's genotype can occur.
- **Natural deaths:** As an organism increases in age, it becomes ever more likely that it will die of "natural causes." When this happens it becomes frozen in its current location, and doesn't eat or reproduce.

It will be immediately obvious that this description includes many features we would expect to find in a population undergoing Darwinian natural selection. To obtain an instance of the class of Evolutionary Automaton, a number of parameters have to be supplied. These include,

- The mapping from genotype to phenotype,
- The conditions under which one organism can eat another,
- The mutation rate,
- The maximum distance an organism can move per time step, and
- The probability that an organism will reproduce given its energy level.

In addition, a particular run of an EA must be parameterised by grid-size, and an initial state. By varying all these parameters, a large number of effects can be produced. Now comes the formal definition.

Preliminary definitions. A *location* is a pair of integers. The *grid* of a microworld of size N is the set of locations $\langle x,y \rangle$ such that $\text{abs}(x) \leq N$ and $\text{abs}(y) \leq N$, where $\text{abs}(z)$ is the absolute value of z . A *cell* is a pair comprising a type and a location (the type allows for different kinds of cell). Two cells are said to be *connected* if their corresponding locations are adjacent or if there exists a third cell which is connected to both of them. A *phenotype* is a set of connected cells including a distinguished cell $\langle t, \langle 0,0 \rangle \rangle$ called the *centre*. Anything can be a *genotype*.

Definition 1. Let Org be an infinite set (the carrier set of potential organisms) and Gen be a set of genotypes. Then a *state* of a microworld of size N is a 6-tuple $\langle O,G,L,S,A,E \rangle$, where

- O is a subset of Org (all the extant organisms, both alive and dead),
- G is a mapping from O to Gen (the organism's genotype),
- L is a mapping from O to the microworld's grid (the organism's whereabouts — actually the location of its centre),
- S is a mapping from O to True or False (whether the organism is alive or dead),
- A is a mapping from O to the naturals $\{ \}$ (the organism's age), and

- E is a mapping from O to $\{ \}$ (the organism's energy level).

Definition 2. Let Gen be a set of genotypes and Phe be a set of phenotypes. An *Evolutionary Automaton* is an 8-tuple $\langle P,M,U,R,X,D,B,T \rangle$, where

- P is a mapping from Gen to Phe ,
- M is a mapping from $\text{Phe} \times \text{Phe}$ to True or False (whether an organism with a given phenotype can eat one with another phenotype) such that if $M(X,Y)=\text{True}$ then $M(Y,X)=\text{False}$,
- U is a mapping from Gen to 2^{Gen} (the range of genotypes a given genotype can mutate into),
- R is a real number in the range $0..1$ (the probability of mutation),
- X is a natural number (the maximum distance an organism can move in one time step),
- D is a function from $\{ \}$ to reals in the range $0..1$ (the probability that an organism will die of natural causes given its age),
- B is a function from $\text{Phe} \times \{ \}$ to reals in the range $0..1$ (the probability that an organism of a given phenotype will reproduce when it has a given energy level), and
- T is a natural number (how close two organisms have to get before one can eat the other).

Let Org be an infinite set of organisms, Gen be a set of genotypes and Phe be a set of phenotypes. Then given the initial state $\langle O,G,L,S,A,E \rangle$ of a microworld of size N , an Evolutionary Automaton $\langle P,M,U,R,X,D,B,T \rangle$ generates a sequence of microworld states by repeatedly executing four procedures — deaths, moves, meals and births — which are specified below. Each procedure updates one or more of the variables O, G, L, S, A, E representing the state of microworld.

The Deaths procedure alters the status, denoted by S , of any number of living organisms depending on their age. It also updates the age of each living organism. The death rate can be controlled by varying the function D . Obviously, the longer an organism is allowed to live, the higher the chances of the survival of its species. A low death rate favours organisms that eat only occasionally, but eat a lot. A high death rate favours organisms that eat little and often.

Procedure Deaths(O,S,A)

For all x in O such that $S(x)$ **Do**

$A(x) := A(x)+1$

With probability $D(A(x))$ **Do**

$S(x) := \text{False}$

The Moves procedure alters the location, denoted by L , of each living organism. Two or more organisms can have the same location. Note that the organisms have no control over their own movement. They just float about randomly. In fact, they have no "behaviour" at all. The only feature they have, which determines their capacity to survive, is their shape, which is a function of their genotype. The maximum distance an organism can move in one time step is controlled by the X parameter. A small value for X will encourage the formation of localised colonies of the same species. The greater the value of X , the faster a species will spread around the microworld.

Procedure Moves(O,S,L)

For all x in O such that $S(x)$ **Do**

Let θ be a random direction

Let δ be a random distance $\leq X$

$L(x) := L(x)$ translated by δ units in direction θ

The Meals procedure, which updates the O and E variables, is where the two vital mappings M and P come into play. Alteration of these parameters can give rise to completely different behaviours. This is where selection takes place. It is important to note that the M mapping does not supply a global measure of fitness. It does not even have to be a transitive relation, thus allowing a paper/scissors/stone relationship between species of the form A eats B , B eats C but C eats A . The absence of a global fitness function makes open-ended evolution possible.

Procedure Meals(O,S,E)

Let s be the set of all x,y in O such that $S(x)$ and $S(y)$ and the distance between x and $y < T$

While s not empty **Do**

Let $\langle a,b \rangle$ be a randomly chosen member of s

$s := s - \langle a,b \rangle$

If $M(P(G(a),P(G(b))))$ or $[S(b)=\text{False and } G(a)\neq G(b)]$ **Then**

$O := O - b$

$E(a) := E(a) + E(b) + \text{size of } b$

Finally, we have the Births procedure. This procedure, whenever a new organism is born, will update all the variables comprising the state. The genotype of a newborn organism is liable to mutation, which will of course alter its phenotype. Mutation can be constrained by the U parameter to ensure that evolutionary change is gradual. The mutation rate is controlled through the parameter R . A high mutation rate will lead to runaway evolution, in which numerous new species are created, only to be superseded by even newer species before they have time to establish themselves.

Procedure Births(O,G,L,S,A,E)

For all x in O such that $S(x)$ **Do**

With probability $B(P(G(x)),E(x))$ **Do**

Pick any y in Org not in O

$O := O + y$; $L(y) := L(x)$

$S(y) := \text{True}$; $E(y) := 0$

$A(y) := 0$; $G(y) := G(x)$

With probability R **Do**

Let $G(y)$ be any member of $U(G(y))$

Before moving on to a specific example, I would like to point out a few features of this class of automata. First, unlike the microworlds used in many previous computer simulations of evolution, that manipulated by an EA contains just organisms. There are no obstacles, no landscape features, in fact no inanimate objects at all. An organism's environment consists entirely of other organisms (dead or alive). Second, as already mentioned, organisms have no behavioural characteristics. In much previous work of this kind, species evolve with different behavioural tendencies. In contrast, all evolution in an EA is evolution of shape. However, the M mapping, which determines who eats who based on their shapes, can be any

(computable) function.² Suitable instantiations of M can lead to very interesting evolutionary effects, which can then be studied without having to grapple with the complexities of behaviour.

Evolutionary Automata bear some resemblance to genetic algorithms [Holland, 1992]. However, there are several very important differences which it is worth emphasising.

- There is no global measure of fitness in an EA, whereas there is in a genetic algorithm. Rather, the M function allows for open-ended evolution.
- Geographical effects are possible in an EA. This limits the influence of each organism to its neighbours. In a genetic algorithm, a single global ranking of all members of the population is used to decide which will survive to the next generation.
- In a genetic algorithm, unlike an EA, organisms do not persist from one time step to the next, and there is no concept of natural death. In a genetic algorithm, each time step is a new generation.
- Unlike genetic algorithms, EA's only feature asexual reproduction. EA's with crossover are an obvious topic for further study.

One way to generalise the class of EA's is to endow each organism with an internal state which can be changed by encounters with other organisms. The M function would have to allow for forms of interaction other than eating. This would permit the evolution of organisms that cooperate in various ways through the trade of information, with both members of their own species and members of other species. As a special case of this, genetic material could be stored and exchanged via the internal state, and reproduction could be made a function of this state. Experiments with such generalisations are a topic for future research.

2. A Simple Evolutionary Automaton

This section and the next describe some computer experiments carried out with a simple but non-trivial EA. These experiments are on a very small scale, and should be thought of as a first foray into the space of possible EA's. In each of these experiments, the microworld is a grid of 50 by 50 locations.

A variety of P 's and M 's were tried, apart from those I'm about to describe. These included a mapping from genotype to phenotype which was very similar to the recursively defined mapping used by Dawkins [1987] in his "biomorphs" program. Using this mapping, with a suitable instantiation of M , evolution can produce spectacular results in a few hundred time steps. But much of the apparent complexity in the resulting organisms is inherent in the recursive nature of the mapping. Because I wanted to study the production of complexity and diversity by

² It's a straightforward exercise, for example, to construct an EA that evolves an improving population of noughts and crosses (tick-tack-toe) playing organisms. An organism's body is a look-up table, mapping board positions to moves. The M function plays two organisms against each other, reading off each organism's moves from this look-up table. The winner eats the loser.

selection alone, the main experiment described in this section uses a simpler mapping.

Three cell types are used: red, white and blue. A genotype has an infinitely extendable tree structure, recursively defined as follows.

- A *colour* is either red, white or blue.
- A *gene* is either,
 - A pair $\langle c,n \rangle$, where c is a colour and n is a natural number, or
 - A list $[\langle c,n \rangle, g_1, g_2, g_3]$, where c is a colour, n is a natural number, and g_1 to g_3 are genes.
- A *direction* is either north, south, east or west.
- A *genotype* is a pair $\langle d,g \rangle$, where d is a direction and g is a gene.

For example, $\langle \text{east}, [[\langle \text{white}, 2 \rangle, \langle \text{red}, 1 \rangle], [\langle \text{blue}, 2 \rangle, \langle \text{red}, 1 \rangle], \langle \text{blue}, 0 \rangle, \langle \text{red}, 0 \rangle], \langle \text{red}, 1 \rangle] \rangle$ is a genotype. The corresponding phenotype is shown in Figure 1. Informally, to form the phenotype for the genotype $\langle d, [\langle c,n \rangle, g_1, g_2, g_3] \rangle$, you begin at $\langle 0,0 \rangle$ and add n cells of colour c in direction d , then rotate 90° anticlockwise and add the cells for g_1 , move back to direction d and add the cells for g_2 , then finally rotate 90° clockwise and add the cells for g_3 . More precisely,

- $P(\langle d,g \rangle) = \langle 0,0 \rangle \cup \text{Grow}(g,d,\langle 0,0 \rangle)$, where
 $\text{Grow}(\langle c,n \rangle, d, \langle x,y \rangle) =$ the next n cells (with colour c) in direction d starting from $\langle x,y \rangle$
 $\text{Grow}([\langle c,n \rangle, g_1, g_2, g_3], d, \langle x,y \rangle) = \text{Grow}(g_1, d, \langle x,y \rangle) \cup \text{Grow}(g_2, d, \langle x,y \rangle) \cup \text{Grow}(g_3, d, \langle x,y \rangle) \cup \text{Grow}(\langle c,n \rangle, d, \langle x,y \rangle)$,
where $\langle x',y' \rangle$ is n units from $\langle x,y \rangle$ in direction d



Figure 1: An Example Phenotype

The rule that determines whether one organism eats another (the mapping M) works by awarding points to the two organisms. An organism gets two points for every one of its cells that covers and beats one of its opponent cells. The “beats” relation is not transitive. A red cell beats a white cell, a white cell beats a blue cell, and a blue cell beats a red cell. In addition, an organism gets a point for every one of its cells that is not covered by any of its opponent’s cells. If one organism is awarded more points than the other, then it eats the other. Let $\text{Points}(p_1, p_2)$ be the number of points awarded to p_1 against p_2 , according to the above description. Then we have,

- $M(p_1, p_2) = \text{True}$ if $\text{Points}(p_1, p_2) > \text{Points}(p_2, p_1)$
 $M(p_1, p_2) = \text{False}$ if $\text{Points}(p_1, p_2) \leq \text{Points}(p_2, p_1)$

The rest of the parameters of the EA are instantiated as follows. The crucial parameter for the suite of experiments I will discuss is the penultimate one, B . This will be varied, by varying the value of θ , to give different effects, as we’ll see shortly.

- $U(\langle d,g \rangle) =$ the set of all genotypes obtainable either,
 - 1) By replacing the direction d by $d \pm 90^\circ$, or
 - 2) By replacing any gene $\langle c,n \rangle$ in g where $n \leq 1$ by $\langle c, n+1 \rangle$, or
 - 3) By replacing any gene $\langle c,n \rangle$ in g where $n > 1$ by $\langle c, n \pm 1 \rangle$, or

- 4) By replacing any gene $\langle c,n \rangle$ in g by $[\langle c,n \rangle, g_1, g_2, g_3]$, where each g_i has the form $\langle c_i, n_i \rangle$ where c_i is a colour and some $n_i = 1$ whilst the other two $n_i = 0$

- $R = 0.1667$ (1 in 6)
- $X = 5$
- $D(n) = \frac{n^3}{125}$ or 1, whichever is smaller
- $B(p,e) = \begin{cases} 0 & \text{if } e < s \\ \frac{e}{s^\theta} & \text{if } s \leq e < s^\theta \\ 1 & \text{if } e > s^\theta \end{cases}$ where $s =$ size of p
- $T = 5$

As should be clear, the mapping U is designed to ensure that (almost) any genotype can evolve from any other, but in small steps. But what is the rationale behind the apparently arbitrary choice of the other parameter values? Many of these values, although they are the result of experimentation to find values which give interesting behaviour, would tolerate quite a lot of variation. The mutation rate (R) must be neither so small as to arrest evolution altogether, nor so large as to prevent species from establishing themselves. The ratio of the amount of movement allowed per time step (X) to the maximum distance at which organisms can eat each other (T)³ must be neither so small as to prevent organisms from finding food before they die of natural causes, nor so large as to negate geographical effects altogether. The function D guarantees that no organism can live longer than five time steps, whilst ensuring that the probability of dying of natural causes is small for young organisms and grows rapidly towards age five.

The function B , which determines the probability that an organism will reproduce, is intended to give some advantage to smaller organisms to counteract the disadvantage they are given by the mapping M , thus encouraging the setting up of food webs. The larger an organism, the more it has to eat to have a reasonable chance of reproducing. The degree of this advantage is controlled by the internal parameter θ . A large θ accentuates the advantage of being small.

3. Experimental Results

Each of the parameters of the EA described above, in particular R , X , D , B and T , could be varied experimentally in order to gain an understanding of the workings of the family of EA’s with similar genotype to phenotype mappings and similar rules for eating, hopefully thus improving our understanding of evolutionary processes in general. In the small suite of experiments whose results I will now present, the value of θ (internal to the parameter B) was varied from 1 to 3. Each run involved an initial population of 150 organisms, randomly placed in the microworld. Each organism in the initial state was given a

³ For efficiency regions, in the actual implementation the microworld was divided into five by five regions, and two organisms were deemed to be within eating distance of each other if any parts of their bodies were in the same region.

randomly generated genotype of the form $\langle d, \langle c, n \rangle \rangle$, where $n \leq 3$.

Twenty experiments of 200 time steps were conducted for each of $\theta=1$, $\theta=2$ and $\theta=3$. The experiments were performed with an Apple Macintosh Quadra, and were implemented in LPA MacProlog. Particular attention was given to the behaviour of three variables throughout each run: the total number of different species, the mean size of organism, and the total population. Two organisms belong to different species if they have different genotypes. The number of species in the microworld at any given time is the number of different genotypes having two or more exemplars. In other words, unique mutations that never reproduce are not considered to have inaugurated a new species.




Figure 2: Graph Shapes

When the variables mentioned above are plotted against time, several distinct graph shapes are to be found, three of which are denoted A, B, and C in Figure 2. There was a strong correlation between population and number of

other words, an exploding population of single-celled organisms was generated. In general, the dominant species in the population was alternately red, white then blue, then red again, with increasing colonies of one colour often chasing declining colonies of another across the microworld.

The most interesting results were obtained with the intermediate value for θ . With $\theta=2$, there was considerable variation in the shapes of the graphs from one run to another. Examples of both the above extremes of behaviour were obtained. But in most of the runs, the population and number of species produced graphs of shape C, whilst mean organism size produced a graph of shape A. The upturn in mean organism size corresponded with the decline in population. This last kind of behaviour corresponds to an initial phase during which small organisms thoroughly establish themselves, supporting a small population of moderately successful predators, followed by the emergence of one or more super-predators which are so successful that they eat all available food, resulting ultimately in mass extinction. Graphs for an actual run with $\theta=2$ are shown in Figure 3.

Considerable variation was present even within the set of runs displaying the last kind of behaviour described. In one of the twenty runs, although conforming to the same overall pattern, the population went through a sequence of

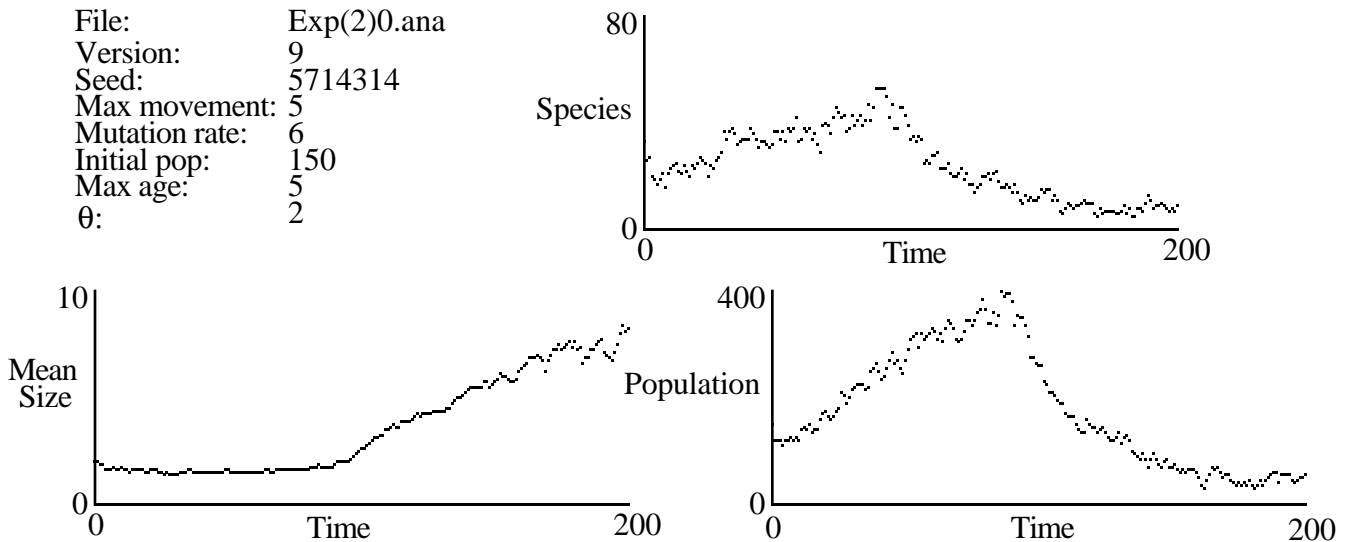


Figure 3: $\theta = 2$

species. In all twenty experiments with $\theta=1$, population and number of species produced graphs with approximately shape B, whilst mean size produced a graph with approximately shape A. In other words, with such a low θ , successful predators emerge very quickly. These predators rapidly eat all available food, and soon become extinct.

With $\theta=3$, all twenty experiments generated graphs for population and number of species with approximately shape A. The mean size of organism, however, remained constant and low (close to one), reflecting the fact that predators find it very hard to establish themselves with such a high θ . In

peaks and troughs coinciding with troughs and peaks in mean organism size. This is just the sort of cyclical pattern often taken to indicate the presence of a homeostatic relationship between predator and prey. Had evolution been arrested before the advent of the first super-predator, this homeostasis might have persisted indefinitely. In addition, not all of the runs ended with total extinction, and it would have been interesting to continue beyond 200 time steps to see when, if ever, the whole population would die out. Clearly, much further experimentation is required with $\theta=2$ to see just what is going on.

Discussion

Langton [1990] asks, “Under what conditions can we expect a dynamics of information to emerge spontaneously and come to dominate the behaviour of a physical system?” The answer, he suggests, is that “living systems . . . must have learned to steer a delicate course between too much order and too much chaos,” and “evolution can be viewed as the process of gaining control over more and more ‘parameters’ affecting a system’s relationship to the vital phase transition [between order and chaos].”

So, according to Langton, life exists at the edge of chaos, in other words in the “complex regime” to be found in the vicinity of the phase transition between order and chaos. But even if we accept this as the basis of the answer to Langton’s question, we cannot stop there. With the complex regime already established, how is it that evolution has been creative enough to produce the diversity and complexity⁴ we find in life on Earth? Without this diversity and complexity, we surely wouldn’t exist. So, within the complex regime, what happens? Let’s zoom in on the complex regime, and ask ourselves, “Within a physical system whose behaviour has come to be dominated by a dynamics of information, under what conditions can we expect diversity and complexity to evolve?”

We might expect the idea of the edge of chaos to apply to this question too. As well as its products, the evolutionary process itself, we might speculate, has to steer a course between too much order and too much chaos. For evolution to be frozen in the ordered regime is for it to lack the creative power to produce enough variation to support diversity or complexity. On the other hand, for evolution to be in the chaotic regime is for it to produce too much variation, enough variation to overwhelm the forces of cumulative selection, preventing the establishment of the sort of stable ecosystem in which we find ourselves.

However, even the preliminary exploration of the world of EA’s presented in this paper suggests that the truth is much more complicated than this. Section 2 defined a simple EA, and Section 3 presented some experimental results based on this automaton. In these experiments, the number of species is a crude measure of diversity, whilst the mean organism size is a crude measure of complexity. Although based on a small sample, these experiments have shown that, with this EA, the conditions necessary for complexity and diversity to arise are extremely fragile. The parameter θ must be close to 2 for there to be a realistic chance of diversity and complexity arising. If θ is too high, we get a degree of diversity but without complexity. If it is too low, we get complexity at the expense of stability. Experiments with a much larger population would probably reveal a more consistent pattern of behaviour with θ set at this critical value.⁵

It would be nice to be able to report the discovery of something at the edge of something else, especially

⁴ Two different notions of complexity are being used here. The first applies to organisms and the second applies to a kind of computation. Precise definitions of both concepts are still pending.

⁵ Imperial College has recently taken delivery of a Fujitsu AP1000 parallel machine, on which it is hoped to test this hypothesis.

something important at the edge of something poetic. However, the results in Section 3 don’t support any grand conclusions. The region in which θ is just right isn’t the edge of chaos. But it is crucially important. And we might expect similar results with many other parameters. Reading the Artificial Life literature, it’s easy to be misled into thinking that a whole range of spectacular evolutionary effects arise inevitably when evolution is simulated in an artificial microworld. Researchers rarely dwell on the hours of tinkering (or sheer good fortune) required to achieve these effects. But it’s just this tinkering that’s most interesting when it comes to answering the big questions. How did nature achieve the tuning of all these parameters without the aid of a scientist’s intervention?

On a more specific level, the results in Section 3 emphasise the fact that it is possible for a species to be too successful. In this artificial microworld, if a species evolves which is part of a complicated food web but which is highly successful at the expense of that food web, then it places the whole ecosystem in danger. (To what extent, if any, this phenomenon mirrors a possibility present in the evolution of life on Earth, I cannot say.) If the θ parameter itself were made part of an organism’s genotype then, with a large population, we might see this parameter being tuned by evolution. It might then be appropriate to speak of the “second-order” selection of whole ecosystems or food webs, in which parameters are spontaneously tuned to support a degree of stability. (See [Kaufmann, 1993, Chapter 6] for ideas along similar lines.)

This suggests a promising project for the future, which is to answer the following question. What is the simplest instance of the smallest generalisation of the class of EA’s, which generates stable ecosystems supporting diversity and complexity, and which spontaneously tunes its parameters to the critical values required to do this?

Acknowledgments

The author is supported by a SERC research fellowship.

References

- [Dawkins, 1987] R.Dawkins, The Evolution of Evolvability, *Proceedings Artificial Life*, pages 201-220.
- [Holland, 1992] J.H.Holland, *Adaptation in Natural and Artificial Systems* (2nd edition), MIT Press (1992).
- [Kaufmann, 1993] S.A.Kaufmann, *The Origins of Order*, Oxford University Press (1993).
- [Langton, 1987] C.G.Langton, Artificial Life, *Proceedings Artificial Life*, pages 1-47.
- [Langton, 1990] C.G.Langton, Life at the Edge of Chaos, *Proceedings Artificial Life II*, pages 41-91.
- [Packard, 1988] N.H.Packard, Adaptation Toward the Edge of Chaos, in *Dynamic Patterns in Complex Systems*, ed. S.Kelso, A.J.Mandell & M.Shlesinger, World Scientific (1988), pages 293-301.
- [Ray, 1990] T.S.Ray, An Approach to the Synthesis of Life, *Proceedings Artificial Life II*, pages 371-408.
- [Yaeger, 1992] L.Yaeger, Computational Genetics, Physiology, Metabolism, Neural Systems, Learning, Vision, and Behavior or PolyWorld: Life in a New Context, *Proceedings Artificial Life III*, pages 263-298.