FROM COMPLEX ENVIRONMENTS TO COMPLEX BEHAVIORS

Filippo Menczer and Richard K. Belew

Computer Science and Engineering Department

University of California, San Diego

La Jolla, CA 92093-0114 USA

{fil,rik}@cs.ucsd.edu

Abstract

Adaptation of ecological systems to their environments is commonly viewed through some explicit fitness function defined *a priori* by the experimenter, or measured *a posteriori* by estimations based on population size and/or reproductive rates. These methods do not capture the role of environmental complexity in shaping the selective pressures that control the adaptive process. Ecological simulations enabled by computational tools such as the Latent Energy Environments (LEE) model allow us to characterize more closely the effects of environmental complexity on the evolution of adaptive behaviors. LEE is described in this paper. Its motivation arises from the need to vary complexity in controlled and predictable ways, without assuming the relationship of these changes to the adaptive behaviors they engender. This goal is achieved through a careful characterization of environments in which different forms of "energy" are well-defined. A genetic algorithm using endogenous fitness and local selection is used to model the evolutionary process. Individuals in the population are modeled by neural networks with simple sensory-motor systems, and variations in their behaviors are related to interactions with varying environments. We outline the results of three experiments that analyze different sources of environmental complexity and their effects on the collective behaviors of evolving populations.

1 Introduction

Adaptation of ecological systems to their environments is often viewed through some explicit fitness function defined *a priori* by the experimenter, or measured *a posteriori* by estimations based on population size and/or reproductive rates. These approaches have been used in the field as well as in computational methods. In the former case, evolutionary ecologists must choose some measurable trait to monitor and characterize the individuals in a population, perhaps finding which traits are best correlated with established measures of fitness such as reproductive rate or success, or environmental carrying capacity. Analogously, the modeler of an artificial life (ALife) experiment has to define how to map some phenotypic trait, be it somatic or behavioral, to a mechanism of offspring allocation, through some fitness function and selection scheme.

These methods have the merits of having allowed for enormous advances in understanding the dynamics of ecological systems, and having enlarged the set of tools at the disposal of ecological theory from purely analytical mathematics (such as differential equations) to explorative idea and minimal system models (such as computer simulations). However, the relationship between adaptive population and environment continues to look like a "black box" function mapping phenotypic traits to reproductive fitness. It is difficult, in the field, to refine the black box to the point of characterizing the way in which the environment shapes the selective pressures acting simultaneously on each trait and interfering with each other at the same time. The analytical tools commonly used by the theorist rapidly become intractable as the spatio-temporal dynamics of the environment become even moderately complex [Kareiva 1994].

Are we then hopeless in the face of characterizing the interactions between the complexity of the environment and that of the evolving population? This paper suggests that modeling techniques currently under development in the ALife community may play an important role in refining the role of the environment in the adaptive process. Genetic algorithms (GAs) have been used as computational models of evolution in a great many simulations of adapting populations. Recently, a new class of "open-ended" GAs seems to have emerged as a model of evolution different from Holland's [1992] classic GA. Mitchell and Forrest [1994] have associated this new class with the term *endogenous fitness* models, and we will use the same terminology in this paper.

One difference often pointed to in the notion of fitness (by which we mean reproductive rate) has to do with whether the definition is imposed externally. In classic GA models, the *(exogenous)* fitness function is determined *a priori* by the experimenter. On the contrary, endogenous fitness does not explicitly map features of the phenotype onto reproductive rate: reproduction is instead determined by the interaction between individuals and environment. Another feature which in our opinion well discriminates between the two classes can be reduced to the dimension of *optimization vs. adaptation.* While in both cases some fitness function is used to allocate offspring to alive individuals, in classic GAs this function is *normalized* over the whole population, because there is some comparison between an individual and the population (e.g., fitness average). On the contrary, fitness is computed with respect to a single individual in endogenous fitness models (e.g., no averaging). Therefore we suggest that endogenous fitness is a *local* measure, while exogenous fitness is global, with respect to the population. As a consequence, in a classic GA, an individual's survival performance in its environment is always scaled on a relative curve, by ranking and/or normalizing the population. This is true for any selection scheme (e.g., proportional, threshold, or tournament) and generation gap (e.g., steady-state or lock-step generations). Optimization of the genotype with respect to phenotypic fitness follows. On the other hand, an individual's endogenous fitness is taken at face value, without comparisons with other individuals. This measure results exclusively from interactions between an individual and its spatially local environment. Other members of the population (as well as members of other species) affect this measure only indirectly, for example by consuming part of a limited resource. Of course, if these conspecific interactions are themselves global, the distinction largely disappears: an individual's reproductive rate depends on environmental variables affected by the whole population. For spatially structured environments, however, interactions and selection retain their local flavor. Thus an individual needs to adapt to its local environment, and not necessarily to be better than the rest of the population, in order to be selected for reproduction.

Gould and Lewontin [1979] were among the first to criticize the view equating adaptation with optimization from a biological standpoint. Ecologists and population geneticists speak of different fitness measures as appropriate in different environmental conditions, for example density-dependent vs. frequency-dependent selection, or reproductive rate vs. expected reproductive success [Stearns 1992]. Recent computational models of adaptive populations have reduced the gap with biological models and increasingly made use of fitness measures defined in biology to characterize the role of the environment in their simulations. Well known models focusing on adaptation to the environment rather than optimization have been proposed by Holland [1992] and Ackley and Littman [1992] among others. The SPO model [Todd and Wilson 1993] considers models of reproductive behavior which make it more a matter of individual "choice" than a response to metabolic/energy considerations. In short, ALife's growing interest in simulation of evolution's creative, open-ended aspects can draw on a growing body of literature for mathematical and computational analyses of the behavior [Brooks and Maes 1994].

In this paper we use a class of models, called *Latent Energy Environments* (LEE), to analyze a number of experiments where the endogenous fitness paradigm is applied to populations adapting to simulated environments. Each environment is designed to explore a different dimension along which to vary the complexity of the survival task. We can then monitor the corresponding complexity of the collective behaviors evolved in response to such variations.

In the next section, we overview the LEE model and its differences with classic GAs. In Section 3 we outline issues relative to individual representations. In Section 4 we analyze the general behavior of LEE populations to draw a mathematical connection between environment and population. In Section 5 we report on three experiments simulating the evolution of different adaptive behaviors and highlight the role of environmental complexity in shaping the selective pressures that allow such behaviors to evolve. In the last section we draw some conclusions, outlining the relevance of our approach for natural and artificial systems. The Appendix introduces the LEE software package that is available for public use.

2 Overview of Latent Energy Environments

It has generally been too easy for GA modelers to assume a fairly direct correspondence between the genotypic data-structure manipulated by genetic operators (mutation, crossover, etc.) and the phenotype's fitness (i.e., number of offspring) ultimately evaluated by the environment. The appropriateness of phenotypic behaviors is obviously conditional upon environmental context, and the definition of "adaptive" behavioral features must therefore depend on a complex interplay of genetic and environmental circumstances.

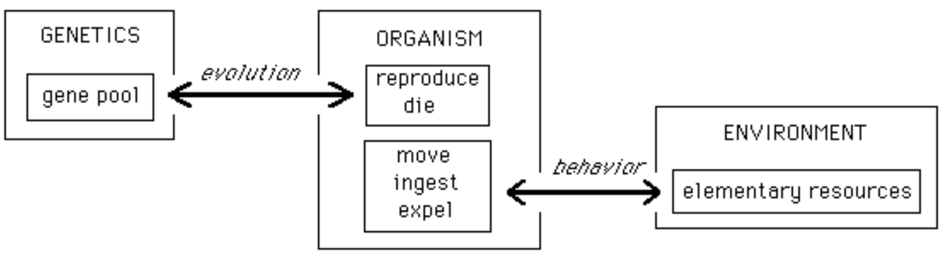

Figure 1: General view of the LEE model.

A general view of the LEE model is offered in Figure 1. The arrow labeled *behavior* represents all types of interactions between organisms and environments. The arrow labeled *evolution* represents, roughly, all genetics-mediated interactions among organisms. In the next subsections, the essential features of such interactions are illustrated in greater depth.

2.1 Environmental complexity

The ALife coupling of organisms with environments brings with it a major methodological problem within ALife research: results reporting behaviors of different organisms in different environments are incommensurate. It is therefore difficult to assess whether an apparently superior behavior is the consequence of more sophisticated adaptive techniques or due to relative *complexity* of the environments. It would be very desirable to be able to define artificial environments of controlled complexity, within which a wide range of ALife techniques might be directly compared.

Godfrey-Smith [1996] provides us with a very good starting point for the definition of both environmental and behavioral complexity, equating complexity to *heterogeneity.* This useful simplification allows environments to be characterized in terms of the number of distinct states they present to an organism, the frequency at which these change, etc. But while Godfrey-Smith abstracts away from specific properties of organization, we want to ground our notion of complexity on observable measures of environmental organization.

There have been several attempts to define generic conditions on environmental complexity that are analytically tractable without constraining evolution's creative potential. Rössler [1974] proposed food density as a simple complexity metric for the environment facing an organism. Even a random walk may be an adequate foraging technique if food is abundant, but as environments become more scarce, more coherent movement is required and the foraging organism may need to depend on landmarks, cognitive maps, etc. [Wilson 1991]. Food density is clearly a useful dimension for ALife simulation, but this single dimension of environmental variability must be extended to include other factors if we are to be able to test the full repertoire of ALife models. Holland's ECHO model describes an extraordinarily rich domain in which many intra- and inter-specic evolutionary questions might be explored [Holland 1992; Jones and Mitchell 1993]. In fact the range of possible experiments, or outcomes to a single experiment in ECHO is so broad that, like real biological systems, its behavior is often difficult to analyze.

Latent Energy Environments are a modest step from the single dimension of food density towards the richness of the ECHO model. To accomplish a useful measure of environmental complexity without compromising analytical tractability, we extend the standard model of "food" as a spatially localized element of the environment, required for survival. Food is replaced in LEE with a series of inert environmental "atomic elements" that must be combined by organisms in order to realize the "energy" they require for survival. Consider a simple discrete world with cells placed on a two-dimensional grid. Let the only source of (positive or negative) energy be via binary reactions; i.e., combining two elements results in an energy gain or loss, and possibly other by-product element(s). Let the base elements belong to an artificial set of types $\{A_1, ..., A_t\}$. Furthermore, let elements of each of these types be generated according to some spatio-temporal distribution. All the possible reactions can be represented by a symmetric reaction matrix indexed by element types, as shown in the example matrix below:

$$\begin{array}{c|ccc} & A_1 & A_j & A_t \\ \hline A_1 & 10+2A_2 & & \text{non reactive} \\ A_i & & E_{ij}+P_{ij} & \\ A_t & & & -3+A_t \end{array} \qquad (1)$$

A reaction occurs when two reactive elements come into contact. This can only happen when an individual catalyzes such a reaction, since reactive elements never occupy the same cell. In the example matrix (1), two generic elements A_i and A_j yield energy E_{ij} ($E_{ij}>0$ for exothermic reactions, $E_{ij}<0$ for endothermic ones) and one or more by-products indicated by P_{ij} . The reaction matrix and the spatio-temporal distributions of elements together represent the laws of physics and chemistry regulating the interactions between an individual and its environment — and among individuals, since the environment of an individual includes the rest of the population. In short, matrix and distributions are the parts of the LEE model which characterize the environment.

We now need some definitions. For any world configuration, there may be many possible combinations of existing elements according to the possible reactions in the matrix. After one of these reactions takes place, two elements are consumed and some new elements (by-products) may appear, giving rise to a new configuration. We can think of this process as a decision tree, where each reaction leads to a new node down the tree, and a path along the tree represents a reaction chain. Eventually, the chain terminates if no possible reactions exist among the remaining elements: a leaf is reached in the tree. Every reaction chain has a corresponding potential energy, given by the arithmetic sum of energies released and/or absorbed by the reactions along the path. For each world configuration, we call *latent energy* the set of potential energies corresponding to all reaction chains starting from that node. The energy remains latent until it is realized as potential energy by choosing a particular reaction chain. By specifying the table's interaction energies and controlling the rate at which elements of each type are introduced, the amount of energy available in its latent form can be precisely regulated. Recall now that a collection of the various elements is distributed across the two-dimensional space of the LEE world, and define the *work* required to release latent energy associated with a pair of elements to be the distance one element must be moved in order to occupy the cell of the other. Then, by controlling the spatial distributions of elements in the world, we can dynamically regulate the amount of work required to combine them.

The next step is to introduce organisms into this abiotic environment, by considering their ability to *move* through the environment and mix its elements as their most fundamental behaviors. We also endow the basic LEE individual with an internal body cavity (gut) and assume ingestion of an element into this gut occurs any time an organism occupies a cell containing one. An element consumed in this fashion is then carried with the organism as it moves. Organisms therefore become the agents actively mixing the otherwise static elements of the LEE environment.

In these preliminary experiments, individuals incur constant metabolic costs simply for being alive. That is, there is no additional energy cost for moving (as opposed to remaining in one cell of the environment) or for carrying elements. The work associated with a pair of elements, based on the Euclidean distance between them in the environment, therefore becomes a lower-bound on the energy cost a most-efficient individual would incur in moving in order to realize their latent energy. Subtracting work from latent energy and summing over all pairs of elements in the environment, we can estimate optimal energy efficiencies, and from this the maximum *carrying capacity* of an environment. In short, by controlling how much latent energy is available and how much work must be done in order to realize this energy, we can obtain a quantitative measure of how difficult it is for an organism (or ecology of organisms) to survive.

The above definitions — latent energy, work, carrying capacity — allow us to control parameters of the environment without specifying just how latent energy is to be realized by one or more species. The monitoring of population dynamics makes it possible to compare ecologies in different environments, in terms of how efficiently each exploits the energy latent in its own environment. Conversely, in a shared environment, competition for the finite resources intrinsically defines the ground on which adaptive success can be gauged. The simple model of metabolism achieved by the range of available reactions provides a space in which to explore the effects of discoveries due to more complex and efficient exploitations of available resources. The principle of conservation of energy permits us to do so without sacrificing analytical tractability.

2.2 Steady-state and asynchrony

The evolutionary process is modeled in LEE by the basic algorithm shown in Figure 2. This type of GA is called *steady-state* because successive generations are interleaved through time rather than lock-step. While from an operational viewpoint decreasing the generation gap results in more diversity being preserved within the population and convergence being slowed down [De Jong and Sarma 1993], our modeling choice is primarily driven by biological plausibility. One important feature of the model is that the population size does *not* remain constant throughout an experiment, and *extinction* is possible. We will show that population size becomes stable spontaneously and robustly when the environmental conditions allow it, without this being imposed externally.

For each time step, every organism executes a basic cycle. This is an intrinsically parallel process; for sequential machines, however, it can be simulated via sequential calls to the organisms. This is the case in our implementation, where each organism is called in random order so as to minimize the spurious interferences of the sequential simulation [Huberman and Glance 1993].

2.3 Endogenous fitness

Both the terms *open-ended* and *endogenous fitness* only capture a partial aspect of the distinctive features that discriminate these models from classic GAs: namely, the fact that the fitness function is not explicitly defined in terms of behavior. The stress therefore is on the *creative* power of endogenous fitness models, which was in fact one of the main motivating factors for LEE [Menczer and Belew 1996]. A more careful analysis of the evolutionary dynamics in LEE, however, permits us to make the distinction more objective. To this end, consider the reproduction and death thresholds, α and ω, in the pseudo-code of Figure 2. With different functional forms for α and ω, the algorithm reduces to either a classic or an endogenous fitness GA.

```
for each time cycle {
    for each alive organism i {
        sense world;
        compute actions;
        perform actions;
        Energy(i) ← Energy(i) + Σ_{actions(i)} Benefits − Σ_{actions(i)} Costs;
        Energy_Envt ← Energy_Envt − Σ_{actions(i)} Benefits;
        if ( Energy(i) ≥ α ) {
            reproduce;      /* clone genotype         */
            mutate;         /* get new genotype       */
            develop;        /* get new phenotype j    */
            Energy(i) ← Energy(i) − E'; Energy(j) ← E';
        }
        else if ( Energy(i) ≤ ω ) {
            die;            /* delete organism i      */
            Energy_Envt ← Energy_Envt + Energy(i);
        }
    }
    Energy_Envt ← Energy_Envt + Energy_Replenish;
}
```

Figure 2: Pseudo-code for the main loop of the LEE GA. Energy is always conserved. The thresholds α and ω are key to the local selective process. Note that several aspects of the algorithm are left to be defined by particular implementations. For example, the computations of benefits and costs associated to actions. Furthermore, single or multiple offspring can be generated depending on what fraction of a parent's energy is allocated to each offspring. In the experiments reported in this paper, offspring are generated one at a time and receive half of the parent's energy.

Classic GA	**Endogenous fitness**
Global selection	Local selection
Rank-normalization	Shared environment
Constant selective pressure	Constant resource flow
Frequency dependence	Density dependence
Exploitation	Exploration
Premature convergence	Genetic drift
Sensitivity	Robustness
Optimization to global environment	Adaptation to local environment

Table 1: Main discriminating features of classic and endogenous fitness GAs.

The classic case occurs when α and ω depend on the energy distribution over the population, causing an interaction among different (possibly all) organisms at the fitness level. This selection scheme is *global* in that it requires an external normalization of the population that no single individual can compute. Such normalization is therefore a supervised process. The simplest example is $\alpha = \omega = \langle E \rangle_p$, where $\langle \cdot \rangle_p$ indicates population average. Then the global aspect of selection is just averaging. In other global selection schemes, such as proportional selection, the normalization may be based on other forms of supervision, for example fitness ranking. In all global selection instances, however, a constant selective pressure is imposed: when the population improves on average, selection becomes accordingly more stringent. The optimizing potential of the GA results from the more or less greedy exploitation of the information contained in the current population, and premature convergence is a typical drawback when the adaptive landscape has local optima.

The endogenous fitness case results when α and ω are constants, for example in LEE we use $\alpha = const > \omega = 0$. Selective pressure results indirectly from the *local* competition for the finite resources in the shared environment. The currency for such resources is the energy latent in the environment flowing from sources outside the system to the work sink. If the organisms do not interact in other ways, the fitness is said to be *density dependent* [Roughgarden 1971]. The lack of a constant selective pressure allows individuals to explore the adaptive landscape without competing with others who are not part of their own local environment. Genetic drift is the typical obstacle to optimal convergence. We have summarized the main discriminating features between the two GA classes in Table 1.

3 Individual representation

One of the fundamental features of LEE is that energy in an environment can be released only through the behaviors of organisms. Behaviors induce reaction sequences, thus catalyzing the transformation of latent energy into usable energy. The notion of behavior — actions taken by the organism that change the world and/or the organism’s relationship to the world — immediately binds an organism to its environment in an intrinsic way.

In any adaptive model, an individual’s behavior is strongly dependent upon its representation. Each individual in LEE is represented by a *body* and a *brain*. The body comprises sensory and motor apparati, and internal reservoirs for energy and environmental resources. The brain consists of a neural network that models the behavior of the individual given its environmental conditions. The fact that the individual is embodied makes its interactions with the environment crucial in determining the degree to which behaviors are well adapted.

3.1 Body

The brain interacts with the environment external to the body through a sensory-motor system. Figure 3 provides an illustration of an organism's body structure with a typical sensory-motor system. The body provides a strong coupling between environmental complexity and difficulty of the survival task [Menczer and Parisi 1992b]. Note that the body also determines the placement and orientation of an organism's sensors and motors.

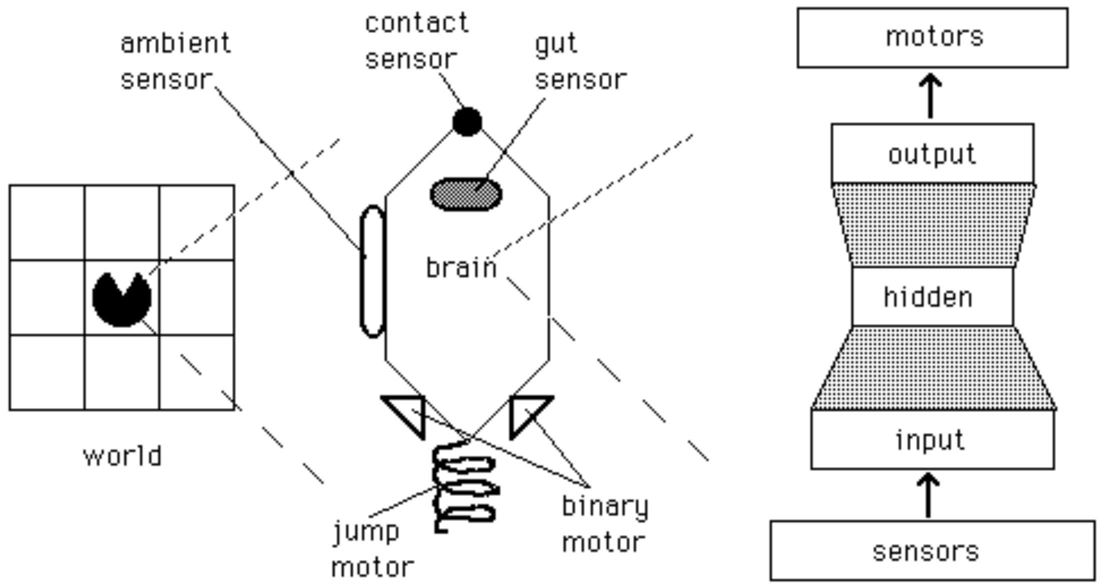

Figure 3: A LEE organism. On the left, the organism is displayed situated and oriented in the environment; in the middle, the body of the organism is expanded and different possible sensors and motors are shown; on the right, the architecture of a neural net modeling the organism's brain is illustrated. No learning occurs during the life of the organisms in the experiments reported in this paper. However, the simulator includes mechanisms of non-Lamarckian learning by reinforcement or sensory prediction.

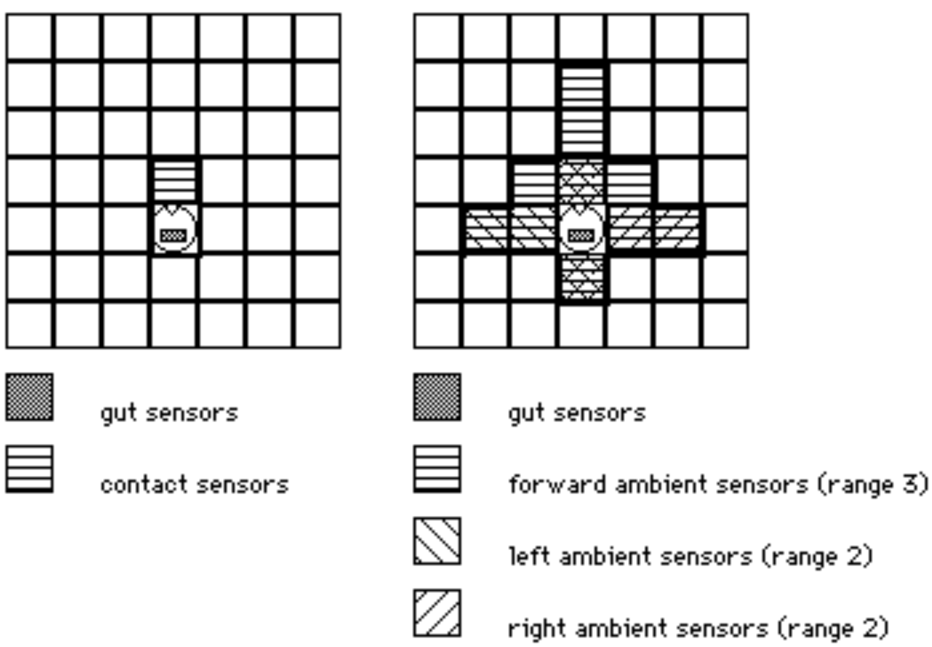

Figure 4: Two examples of sensory configurations. Both systems use internal gut sensors. One uses only contact external sensors, while the other uses three ambient sensors with different ranges and orientations. In LEE, each sensor can be designed to signal the presence of any complex of different elements; however, the experiments reported in this paper only use sensors that signal single elements.

The sensory system is composed of a set of sensors with different characteristics. These collect information from either the external world or the internal body of the organism, and map it onto the brain input. Sensors may differ in range, directionality, sensitivity, resolution, specificity, accuracy, etc. Examples of external sensor types used in the experiments reported here are *contact* and *ambient* sensors. Contact sensors provide a binary presence/absence indication of some element (or complex of elements) in the space directly adjacent to the sensor. They are sufficient to support avoidance behaviors, but not approach. Ambient sensors signal the presence of some element (or complex), summed over cells in a neighborhood of the sensor and weighted inversely according to their distance[1]. Ambient sensors can underlie approaching behavior only if the brain possesses either a memory to compare *temporal* differences in a single sensor, or multiple ambient sensors placed and oriented differently to compare *spatial* differences. Two examples of sensory configurations are shown in Figure 4.

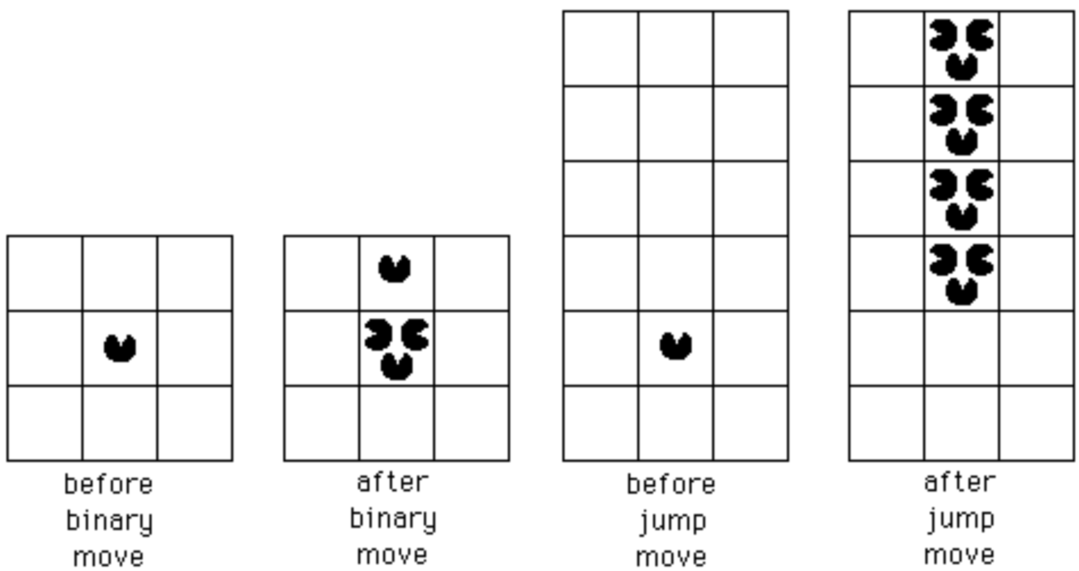

Figure 5: The functioning of two motor systems. The positions and orientations that are possible after a move are shown for a binary and a jump motor (with power 1 and 4, respectively). Note that different positions and orientations are not equally probable. It is possible for LEE organisms to have more than one motor, in which case the moves are additive; however, the experiments reported in this paper only use motor systems made of single motors.

The sensory information is elaborated by the brain to produce an output, interpreted as an action (movement) in the world. Motors function as the output effectors of the organism, changing its location and/or orientation with respect to a fixed environment referential frame. Motors may differ in energy efficiency, power, orientation, accuracy, etc. Examples of motor types used in the experiments of this paper are *binary* and *jump* motors. A binary motor is very simple and has been used in many other ALife experiments (e.g., by Nolfi, Parisi and Elman [1994]). It allows an organism to move to the cell immediately ahead, turn 90 degrees right or left, or stand still. A jump motor

[1] This metric is defined behaviorally, in terms of the number of moves required of the organism to reach this location. Distance is therefore well-specified with respect to some motor system.

moves the organism ahead by some distance (determined by the output) and randomly changes its facing direction. In Figure 5 we illustrate the operation of these motor systems.

Each organism also has an internal cavity we call its *gut*. A gut is capable of holding a number of atoms specified by the organism's genotype. An organism automatically ingests atoms over which it travels, causing these atoms to be placed in its gut (irrespective of gut capacity). If the gut already contains some atoms that are reactive with the new atoms, one or more reactions take place and as a result the organism either accrues (exothermic reaction) or loses (endothermic reaction) energy. If at the end of these reactions the gut is still filled with non reactive atoms beyond its capacity, some atoms are randomly expelled before the next ingestion, leaving a "trace" of the organism's presence in the environment. This can be used as a primitive form of communication among individuals. Furthermore, the gut can function as an important form of "memory" by the use of internal *gut* sensors (cf. Figure 4), as has been shown elsewhere [Menczer and Belew 1994]. Some, all, or none of these somatic characteristics may be allowed to evolve in LEE experiments.

3.2 Brain

The brain receives input corresponding to the organism's stimuli and produces output controlling its behaviors. We use a well-studied type of neural network as our computational model of an organism's brain [Rumelhart and McClelland 1986]. The signals from each sensor are used as the inputs to a feed-forward neural network with one or more hidden layer of units,[2] and the neural net's output units are used to control the motors (cf. Figure 3). The network mapping sensory states to motor actions is part of an organism's genome and thus evolves by selection and mutations of the neural net's connection weights, represented as real numbers. Mutations are obtained by adding a random deviate (uniformly distributed in some user-specified interval) to a fraction of the connection weights.

A different way in which behaviors can adapt to the environment, within the lifetime of an individual is by *learning.* The role of learning by prediction of sensory states and by reinforcement — two mechanisms of unsupervised learning that do not assume the presence of any external teachers — have been explored elsewhere [Menczer and Belew 1994]. In this paper we will only use genetic evolution, i.e., no adaptation can occur at time scales shorter than a lifetime.[3]

4 Population, carrying capacity, and fitness

By observing some emerging properties of LEE populations, we have shown that the task of analyzing endogenous fitness models is greatly facilitated by relying on well-founded constraints such as the conservation of energy [Menczer and Belew 1996]. In

[2] We will use a single hidden layer in the experiments of this paper.
[3] Note, however, that the duration of a lifetime is not fixed in LEE. We will show that it is, rather, an adaptive life-history trait.

this section we summarize those results, to illustrate the roles that different measures of fitness play in latent energy environments.

There is a connection between the size of an evolving population and the resources available from its environment. Such resources impose a limit on how many individuals can be sustained. However, different behaviors may result in different efficiencies in using the resources of the same environment. Thus it is the combination of population behavior and environment that characterizes the *carrying capacity* of that environment. For example, extinction indicates that the carrying capacity is insufficient to support a population large enough to withstand stochastic fluctuations.

LEE allows us to quantitatively estimate adaptedness of behaviors by monitoring the population size throughout an experiment. To see how, let us analyze the relationship between latent energy and population size in a particular example. The experiments reported in the rest of the paper will make use of analogous derivations to predict and evaluate evolved behaviors. Consider the following reaction energy matrix for a simple environment in which two types of atomic elements, A and B, are distributed uniformly:

	A	B
A	$-\beta E$	E
B		$-\beta E$

(2)

where E and β are environmental parameters subject to the constraints:

$$\begin{gathered} E > 0, \\ 0 < \beta < 1, \end{gathered} \tag{3}$$

and there are no by-products. Matrix (2), along with the spatio-temporal distributions of the elements, determines the environmental complexity. Conditions (3) make the environment non-zero-sum. The minimal set of elements that can be found in the environment and completely transformed into energy by alternative behavioral strategies is {A,A,B,B}. In fact, with these elements an organism can catalyze the following sets of reactions:

$$(A+A),(B+B) \to -2\beta E, \tag{4}$$

$$(A+B),(A+B) \to +2E, \tag{5}$$

where strategy (5) is clearly more advantageous than (4), given the conditions (3). If r is the expected rate of replenishment for both A and B elements, then a set {A,A,B,B} is produced every $2/r$ cycles. Thus strategies (4) and (5) produce energy changes

$$\Delta E_{(4)} = \frac{-2\beta E}{2/r} = -r\beta E, \tag{6}$$

$$\Delta E_{(5)} = \frac{2E}{2/r} = rE, \tag{7}$$

per unit time (cycle), respectively. Let us now consider the situation at *equilibrium.* On average, two conditions are verified: first, by definition, the population size remains constant; second, energy is consumed by organisms in the population at the same rate at which it is produced (otherwise the population size would change). Using brackets to indicate time averages, we can write:

$$p \equiv \langle \text{population size} \rangle = \text{const}, \tag{8}$$

$$\langle \Delta E \rangle = 0. \tag{9}$$

Since energy is always conserved, the only consumed energy is that lost in the form of work, that is, used for moving in the world. At any time step, all individuals get to make a move. Here for simplicity we set the cost of all moves equal to a constant that we call c. Then, using Eq. (8), the average energy used by the population per unit time is pc.

To calculate how much energy is produced, we must know the strategy used on average by the population for combining elements. In other words, we need to determine how efficiently the latent energy contained in the elements is transformed into usable energy. It is important to note that resources cannot build up forever in the environment, otherwise by statistical arguments the equilibrium would become less and less stable. Therefore, by choosing an appropriate time unit, we can make the additional assumption (verified in every simulation) that elements are transformed into energy at the same rate as they are replenished by the environment. Strategies (4) and (5) provide the connection between elements and produced energy. We then introduce a probability distribution over these two strategies: let η be the probability of strategy (5), so that $(1-\eta)$ is that of strategy (4). The average energy produced per unit time by the population is obtained summing the energy changes of Eqs. (6) and (7), weighted by the corresponding probabilities:

$$\eta \Delta E_{(5)} + (1-\eta)\Delta E_{(4)} = r\left[\eta E + (\eta - 1)\beta E\right]$$

so that we can finally rewrite Eq. (9) as

$$\langle \Delta E \rangle = rE\left[\eta + \eta\beta - \beta\right] - pc = 0. \tag{10}$$

Equation (10) provides the link between efficiency of behavioral strategies, expressed through the probability distribution over catalyzed reactions (η), and population size *(p).* The remaining variables *(r, E, β, c)* are environmental parameters. If η is known, we can solve Eq. (10) for the expected population size:

$$p = \frac{rE}{c}(\eta + \eta\beta - \beta). \tag{11}$$

In particular, the case of optimal behavior, $\eta=1$, corresponds to the maximum sustainable population:

$$p_{\max} = \frac{rE}{c}. \tag{12}$$

The converse case is useful for estimating the optimality of a population's behavior at equilibrium, by measuring the population size and solving Eq. (10) for η:

$$\eta = \frac{\frac{pc}{rE} + \beta}{1 + \beta}. \tag{13}$$

As an illustration of this simple analysis, in Figure 6 population size is plotted versus time for three simulations with different entries (β and E values) in a reaction matrix like (2). No evolution is allowed in these runs. Disregarding the initial damped oscillations (in agreement with population dynamics models) and the noise fluctuations, we readily observe a steady-state regime in which the population size depends on the rate at which energy is introduced into the world. Simulations are labeled by the quantity K, which is the carrying capacity for a population of individuals with random behaviors (obtained from Eq. (11) with $\eta=1/2$):

$$K = p_{\eta=1/2} = \frac{rE}{2c}(1 - \beta). \tag{14}$$

We can use the equations to make predictions about the outcomes of the simulations, and to compare behaviors in the different environments. The smallest K corresponds to a random behavior population smaller than the amplitude of the fluctuations, so extinction occurs rapidly. Larger K values result in different stable population levels. Substituting the measured population sizes into Eq. (13), we find that the corresponding strategies are not significantly different (η~1/2, or random behavior, in both cases). Then the observed difference is to be attributed to the different environments, rather than to different behaviors.

Of course, the linear relation (10) holds only at equilibrium and for this simple example: the more general non-equilibrium case and more complex environments may yield systems of differential equations hard to solve analytically. In the remainder of this paper, environments will be designed so as to keep their analytical characterization simple. We want to stress that carefully designed latent energy environments allow us to maintain an accurate connection between environmental and behavioral complexity.

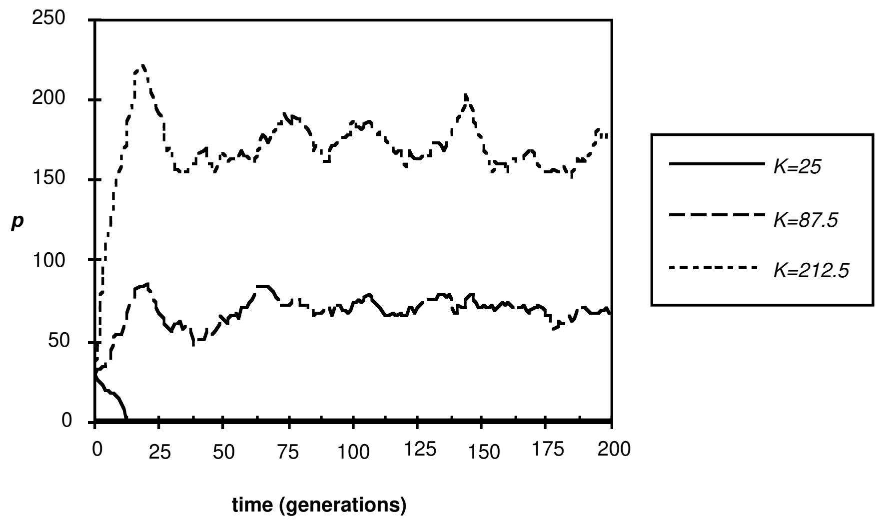

Figure 6: Carrying capacity for three simulations with different reaction matrices. K is a measure of latent energy: when enough latent energy is supplied by the environment to avoid extinction, an early equilibrium is reached. (From [Menczer and Belew 1996].)

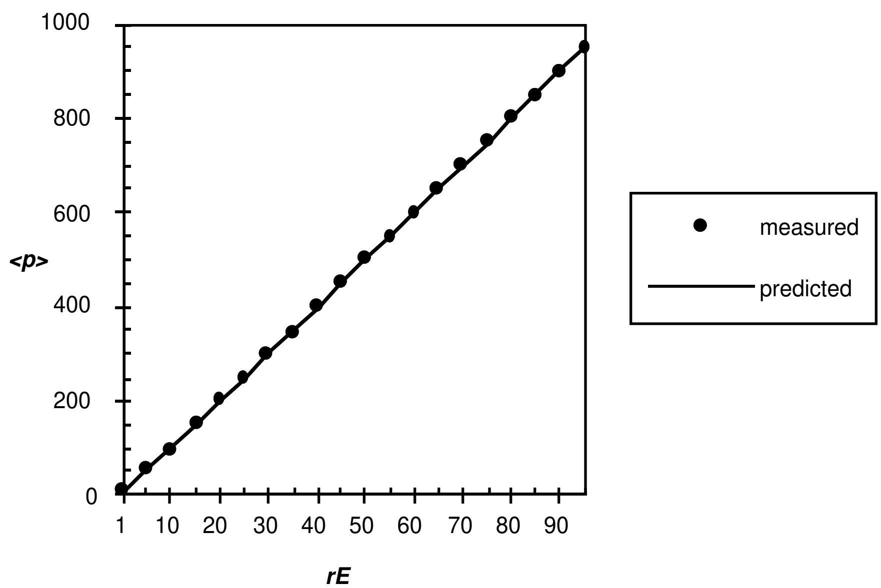

Figure 7: Correlation of population size with environmental carrying capacity. Measures of population size are time averages in the steady-state regime; predicted population size is from Eq. (15). (From [Menczer and Belew 1996].)

Since individuals interact solely by sharing finite environmental resources, fitness is density dependent and thus affected by the size of the population. Under such conditions, carrying capacity is commonly considered in biology the best measure of fitness [Stearns 1992]. We now want to point to the strong role of the environment under density dependent fitness, by showing that the carrying capacity of certain environments can be predicted very reliably from knowledge about the environment alone, without considering behaviors. One such environment, even simpler than the one described in matrix (2), is easily characterized: assume that there is just one element with, say, uniform distribution and rate of replenishment *r;* let each atom be associated with energy E (again, no by-products).[4] Equation (8) still holds at equilibrium, and Eq. (9) for energy conservation yields in this case:

$$K \equiv p = \frac{rE}{c} \tag{15}$$

where the carrying capacity K corresponds to the size of the population, whatever the behavior. Equation (15) is different from Eq. (11) in that there is no independent variable connected with behavioral effects. The prediction, by Eq. (15), of perfect correlation between population size and environmental carrying capacity (parameterized by rE) is confirmed by the measures illustrated in Figure 7.

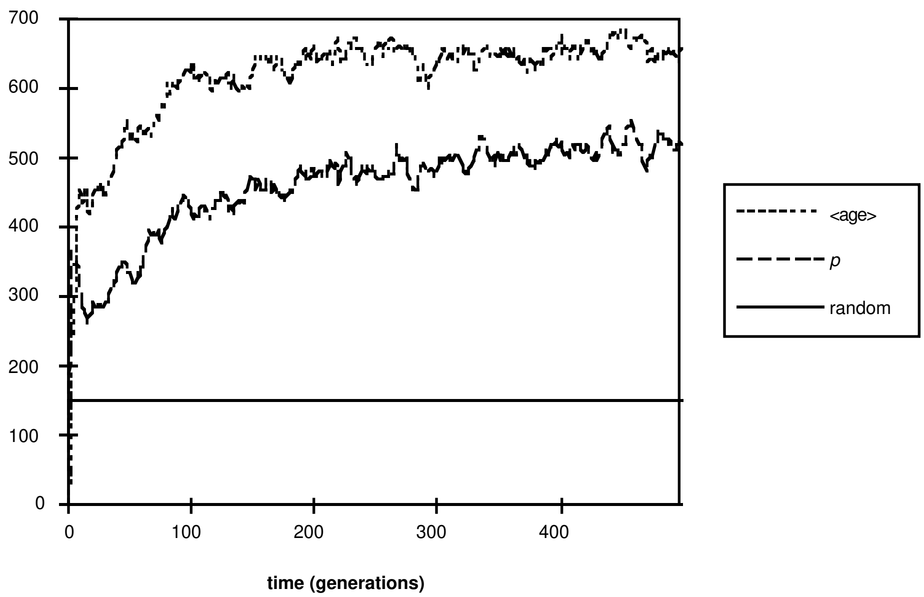

Figure 8: The role of behavior in environmental carrying capacity. Here Eq. (12) yields p_{max}=900, while the size of a population with random behaviors (from Eq. (14)) is marked "random". Average age is also shown to be correlated with population size in this experiment (units are cycles). (From [Menczer and Belew 1996].)

When organisms are allowed to evolve behaviors that may use the latent energy more efficiently, an increase in average age and in population size is observed. The situation is illustrated in Figure 8, where measures are from LEE simulations with a reaction matrix like (2). During these non-equilibrium phases Eq. (8) does not hold, and the environment is no longer sufficient to predict carrying capacity. Population size, on the other hand, can still be used as a fitness measure at the population level. Its increase is due to the evolution of better behaviors, that bring η from 0.5 to approximately 0.73 (from Eq. (13)). It turns out that population size is well correlated with *expected*

[4]In the LEE simulator, the oxymoron “unary reaction” refers to the realization of energy from a single atom, without the need for combinations with other atoms.

reproductive success, a fitness measure used in biology under density dependence. Therefore we consider population size our (dependent!) measure of fitness in LEE, and will use it in the remainder of the paper to analyze the interplay between environments and behaviors of evolving populations.[5]

5 Experiments

In this section we outline the main results of three LEE experiments. Each experiment explores a different dimension of environmental complexity, by a set of environments designed to characterize and vary such complexity. We then analyze the evolving population to find the ways in which it responds and becomes adapted to the different complexities. These experiments are aimed at identifying the degree to which the environment can create and shape the selective pressures driving the adaptive process.

Noise plays an important role in the LEE simulator, and this is reflected in stochastic fluctuations of all the monitored variables. By repeating a simulation for several runs, each with different initial conditions, we can average out such noise and find statistically significant effects. The results reported here are all found to be statistically significant across runs, but will be visualized more easily by plotting results of single runs.

5.1 Patchiness

5.1.1 Simulation setting

In the first experiment, the latent energy of the environment is as described in the reaction matrix (2) of the previous section. Since the spatial dimension is crucial for environmental characterization [Kareiva 1994], we have explored this dimension by observing the response of an evolving population to variations of the spatial distributions of elements in the world.

The environment is replenished with atoms of both elements at a constant rate and according to pseudo-gaussian distributions along the two spatial axes.[6] By decreasing the variance of these distributions, we consider four such environments of increasing patchiness. These are shown in Figure 9. This LEE world is 25 by 25 cells wide; however, due to its toroidal edge conditions, the environment appears as arbitrarily large and periodic. The patch centers of the two elements are offset diagonally. The first, uniform environment has no spatial structure and thus no patch size; in this sense, it is the simplest. The patchy environments have increasing spatial structure, which we

---

[5]It is interesting to note that in this population-based analysis, a population of efficient foraging *individuals* (e.g., the "superman" hand-crafted by Ackley and Littman [1992]) would be considered a less-fit species if this efficiency meant that the environment could not support as large a population.

[6]We obtain a pseudo-gaussian probability density in LEE by adding one or more uniform probability densities on some interval along each axis. The more of these uniform probability densities are added, the more peaked the resulting distribution (i.e., the smaller its variance). This is a standard statistical method, derived from the central limit theorem. The four environments of Figure 9, for example, are obtained by adding up to four uniform probability densities, respectively, for each element and for each axis.

consider a source of complexity. The patches overlap largely in the second environment, less in the third, and are completely separated in the fourth and last environment. Therefore the patch size decreases from the entire world size to about half of the world's diagonal.

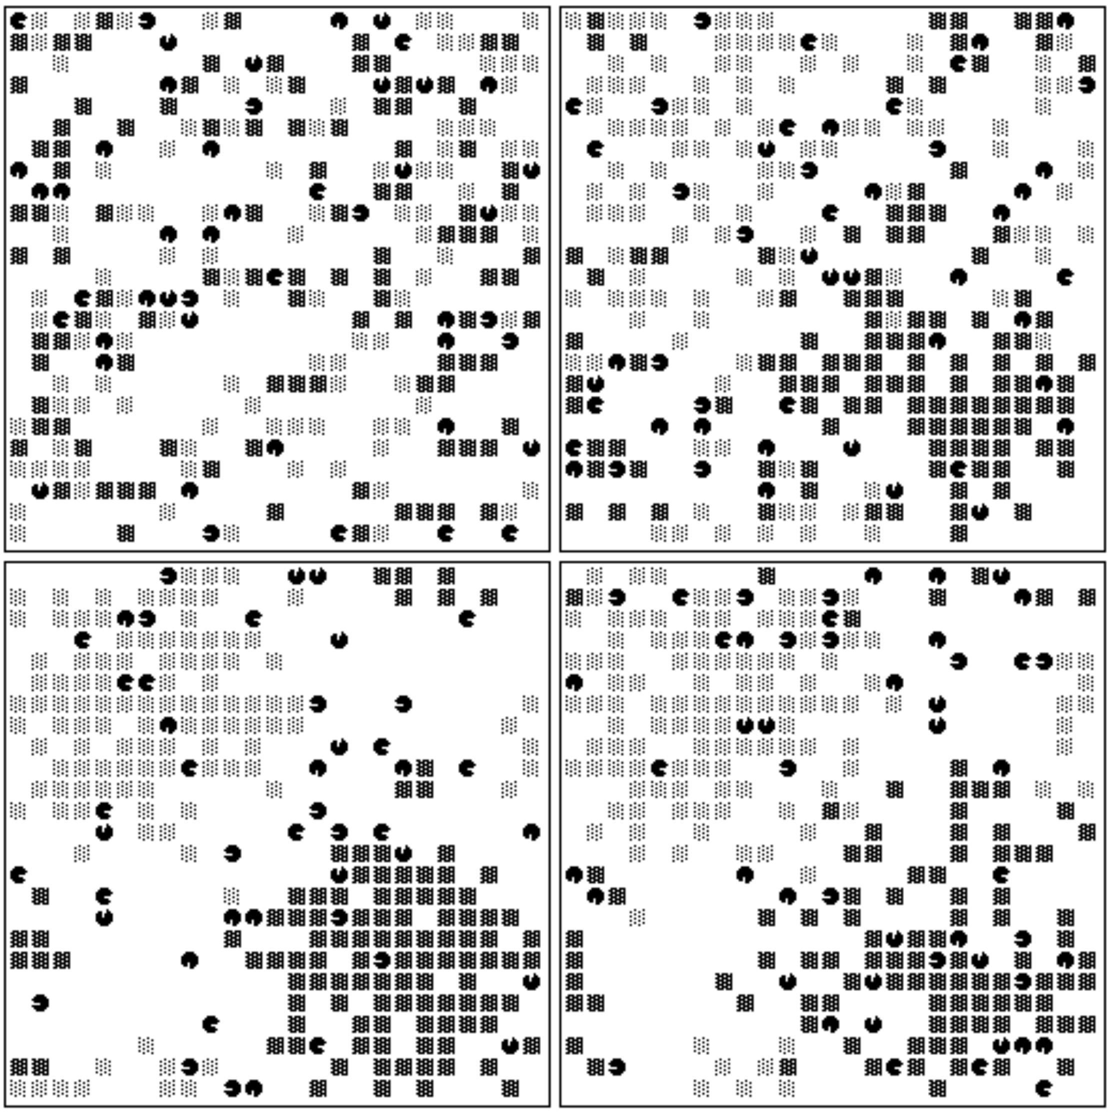

Figure 9: Snapshots of four environments used for the first experiment. They are ordered (clockwise from top left) according to increasing patchiness: from the simplest case of uniform element distributions to the most structured environment with sharply separated element patches. A and B atoms are marked as squares of different shades of gray.

To signal the presence of either element in its neighborhood, each organism has two ambient sensors, oriented forward (cf. Figure 4). Like all ambient sensors used in this paper, they have range 5. By comparing the two signals, an organism can discern whether it is located closer to an A or B patch. There are also two contact sensors, identifying the content of the cell facing the organism. Finally, two gut sensors make it possible to compare external patch signals with previously ingested atoms and decide whether they are of the same or different elements. The four sensors provide input to the neural net, that has six hidden units and one output. The latter drives a single jump motor (cf. Figure 5). Like all jump motors used in this paper, this has power 10.

Given this representation, adaptive individuals can evolve behaviors yielding to the different carrying capacities obtained by Eq. (11). The parameters (r, E, c) of the simulations are set in such a way that a population with random behaviors, catalyzing all reactions with equal probabilities, has a size $p_{\eta=1/2} = 50$ (from Eq. (14)). Given the characterization of the environments in this experiment, it is easy to see that the most efficient behavior consists of jumping far from patches of the same element as an atom in the gut, and otherwise foraging with small jumps: this way exothermic reactions are catalyzed. The carrying capacity of this collective behavior is $p_{\max} = 200$ (from Eq. (12)).

Contact sensors	Gut sensors	Jump motor
A	A	Not nearest
B	A	Nearest
A	B	Nearest
B	B	Not nearest
Any condition	Empty	Any move

Table 2: Behavior of maximum energy efficiency for uniform environments in the first experiment. No cell contains both A and B atoms because they would react together, so that contact signals cannot be ambiguous. 'Nearest' means jumping to the cell immediately facing the organism, i.e., the one sensed by contact sensors.

Ambient sensors	Gut sensors	Jump motor
A>B	A	Far
B>A	A	Near
A>B	B	Near
B>A	B	Far
Any condition	Empty	Any move

Table 3: Behavior of maximum energy efficiency for patchy environments in the first experiment. Hidden units must perform the comparisons between ambient signals. 'Far' and 'near' mean jumping distances greater or smaller than patch size, respectively.

Up to a point, patch size of the environment determines how complex such efficient behavior can be. In uniform environments, contact sensors alone provide all the information needed to implement such behavior, as shown in Table 2. No patch size information is available, because there are no patches. Nothing is to be lost from jumping too far or too near, as long as endothermic reactions are avoided. In environments with spatially clustered elements, on the other hand, space is more structured and thus more complex. The population needs to evolve an "understanding" of environmental structure (e.g., patch size) to gauge appropriate moves. Organisms need to make use of the more ambiguous information provided by ambient sensors — and perform preliminary comparisons between their signals — to decide how far to jump (recall that the direction of a jump cannot be decided). Due to the two additional requirements (comparison of inputs and gauging patch size) this is a more complex behavior, as illustrated in Table 3. Of course, the complexity of the task would be greatly affected by changes in the sensors, as discussed elsewhere [Menczer and Belew 1994]. With the given sensory system, however, such is the direction in which we expect environmental complexity to drive the evolution of adaptive behaviors in this experiment. More complicated spatial organizations (e.g., checker-board patterns and other layouts with even higher spatial order statistics) can reasonably be expected to require still greater sophistication from individuals' cognition.

5.1.2 Results

We have run four simulation with different environmental conditions. The first two simulations have stationary environments, whose complexity remains constant through evolutionary time. The last two have changing environments, whose complexity increases in different ways.

Populations adapting in the two fixed environments reach carrying capacities shown in Figure 10. The patchy environment is more complex and the population is unable to evolve any behavior superior to the random one, as seen by comparing population size with the predicted carrying capacity of the random behavior. The uniform environment is more benign and (a good part of) the population can evolve the more efficient behavior of Table 2, resulting in a carrying capacity about three times as large (η~0.83 from Eq. (13)). The residual selective pressure does not appear sufficient for convergence to the optimal strategy within the observed evolutionary rime, i.e., to push the population to its maximum size.

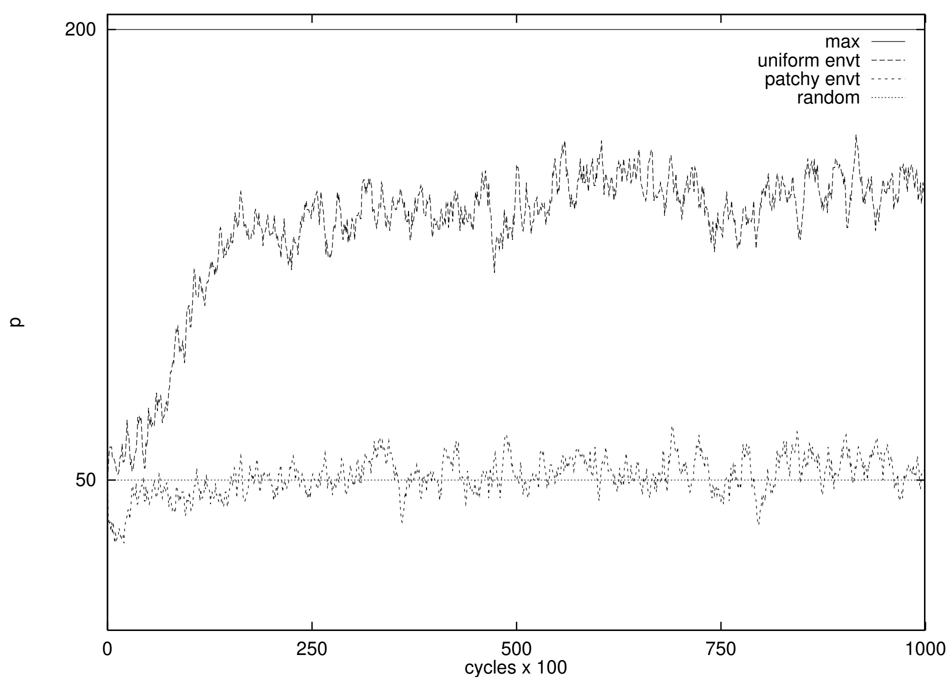

Figure 10: Population size for two fixed environments. These are the first (uniform) and the last (most patchy) worlds of Figure 9. The predicted carrying capacities corresponding to random and most efficient behaviors are also shown.

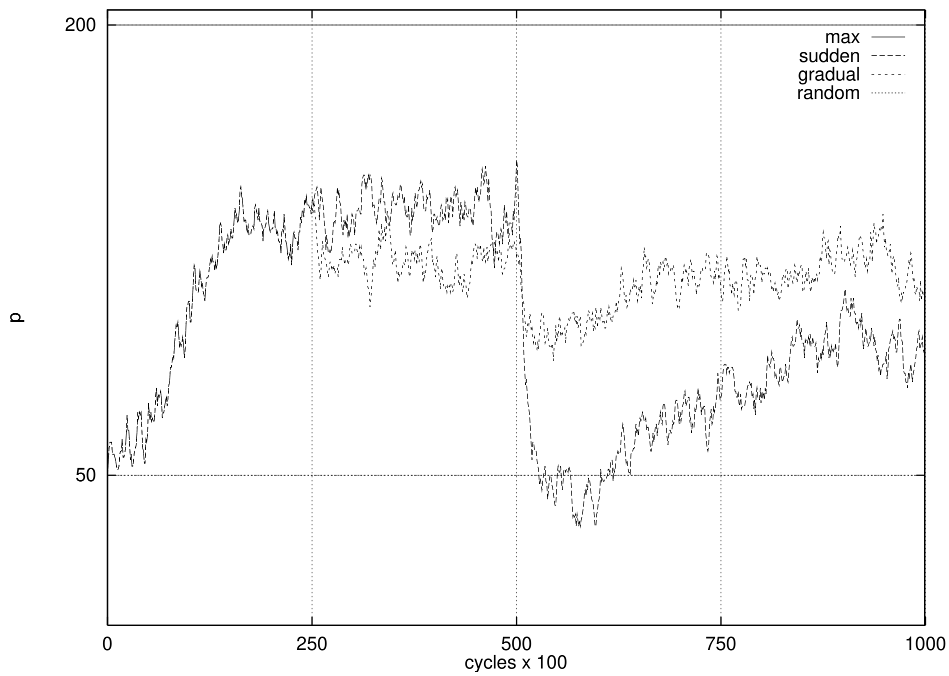

Figure 11: Population dynamics for two changing environments. The 'sudden' change is a single transition, at 50,000 cycles, from the uniform to the most patchy worlds of Figure 9. The 'gradual' change is made of three transitions, every 25,000 cycles (see vertical grid), through all of the four worlds of Figure 9 in order of increasing patchiness. The predicted carrying capacities corresponding to random and most efficient behaviors are also shown.

Figure 11 illustrates what happens to the populations evolving in the two environments of increasing complexity. When the environment suddenly becomes very patchy, the population continues to apply the evolved behavior of Table 2 with catastrophic consequences; the population size decreases to the random level. Since the individuals are essentially behaving the same way right before and right after the environmental transition, we could say that the carrying capacity of the environment has dropped. However, the population then grows again, becoming significantly higher than random ($\eta \sim 0.63$). This is evidence that (part of) the population is able to evolve a behavior similar to the one of Table 3. In the simulation where the environment becomes gradually more patchy, adaptation is not so catastrophic. There are small performance degradations that are quickly offset by the evolution of improved behaviors, so that finally the population can cope quite well with the most patchy environment — over doubling with respect to the population starting in the patchy environment ($\eta \sim 0.7$).

5.1.3 Discussion

The above results demonstrate, in the first place, that patchiness can provide a useful measure of environmental complexity. Further, spatial heterogeneity affects evolved behaviors. A uniform world is simple because there is nothing to be learned about different parts of it: it is the same everywhere. Here, as shown by Figure 10, the simple behavior of Table 2 is discovered rapidly by the evolving population. Conversely, a patchy world is complex because of its spatial structure. If such complexity is too large for the population to make sense of it, then no behavior evolves to exploit its resources better than by moving at random. The evolutionary leap from random behavior to the strategy of Table 3, in terms of individual neural networks, is much larger than to the strategy of Table 2. We conclude that in the most patchy environment, the selective pressure away from the random behavior is too small to discover more complex behaviors within the observed evolutionary time.

Simpler environments have additional advantages for the adaptive individuals. As Figure 11 indicates, behaviors adaptive in the more complex environments can indeed be evolved but only if the population has previously adapted to less complex environments. Thus a gradual increase in environmental complexity generates an adaptive response in the collective behavior evolved.

If the environment undergoes a large increase in complexity, any advantage of previously evolved behaviors seems lost: after the sudden transition occurring at 50,000 cycles, the population goes down to the random level, evidence that the simple behavior appropriate for the uniform environment (cf. Table 2) is not adaptive in the patchy environment. This catastrophe is only illusory, however: the fact that population size grows again in the following phase represents evidence that the more complex behavior (cf. Table 3) has evolved. This did not happen when the environment was patchy to start with (Figure 10); the adaptive behavior evolved in the milder environmental phase therefore provides an evolutionary stepping point.

5.2 Biodiversity

5.2.1 Simulation setting

The second experiment explores the role of environmental complexity in enforcing ecological balance. The environment permits the population a more robust survival when this is structured in such a way to maximize its biodiversity. The population is divided into two "species" distinguishable by a genetic marker; parents pass their species gene to offspring. Some degree of ecological stability might be provided if it were possible for this genes to mutate occasionally. However, in this experiment we are interested in the conditions under which biodiversity is maintained by environmental pressures alone. Therefore the species gene is not mutated. Since reproduction occurs asexually by cloning, species are completely determined by philogeny. Let us call a and b the alleles of the species gene. Let us further assume that the two species have different metabolisms, given by the following reaction matrices:

a	A	B
A	$E+2B$	non reactive
B		$-\beta E$

(16)

b	A	B
A	$-\beta E$	non reactive
B		$E+2A$

(17)

for species a and b respectively. Both A and B elements are uniformly distributed and replenished at the same, low rate.[7]

Since different elements are never reactive, and conditions (3) apply, all individuals have only two (non-exclusive) possible actions. This simplifies the range of behaviors that we can explore in this experiment. The first action is to ingest atoms of the element whose symbol is the (capitalized) letter of their species allele, and gain energy; the second possibility is to ingest atoms of the other element, and lose energy.[8] The "optimal" (most efficient) behavior requires that species a and b catalyze exothermic reaction by foraging for A or B atoms, respectively, and avoid other atoms by turning away. The consequences of this optimal collective behavior are conditional upon the biodiversity of the population; in fact, if one species goes extinct, the optimal carrying capacity halves because the remaining species can no longer use the other's by-products. The random behavior is to always move ahead, irrespective of input, and eat every atom along the way. Biodiversity has no effect on a population with random behavior.

The representation of individuals is quite simple. Each organism has two contact sensors signalling the presence of A or B atoms in the facing cell. Its neural net has four hidden units and two outputs driving a single binary (power=1) motor. The gut has size=1, i.e., it can hold at most one previously ingested atom between actions. Since only same-element atoms react, there is no need of gut sensors.

5.2.2 Results

We have run two simulations for this experiment, using β to model two environments of different complexity. In the first 'easy' environment, $\beta = 0.1$. The parameters are such that the carrying capacity achievable in this environment by the optimal collective behavior described above is $p_{\max} = 95$. The population cannot go beyond this limit. This environment is quite benign and thus there is small pressure to push the population above carrying capacity corresponding to the random behavior ($p_{\eta=1/2} = 42.75$). In fact, as shown in Figure 12, the population does not go beyond such level within the observed 50,000 cycles.

In the second 'hard' environment, $\beta = 0.5$ and thus the situation is more harsh. The maximum carrying capacity of this environment, for an optimal population, is $p_{\max} = 75$; the random carrying capacity is ($p_{\eta=1/2} = 18.75$). The ratio between the former and the latter is almost twice as large as in the easy environment, therefore there is larger selective pressure toward more efficient behaviors. Figure 13 shows the population

[7] In this and the following experiment, the world can fill up with unused elements. If this occurs, replenishment of those elements is suspended.

[8] Notice, however, that in the former case there are by-products which may still react in the gut and consume part of the energy that was acquired with the original action. This will happen half of the times on average, due to the digestion algorithm (cf. Section 3.1); gut contents are shuffled, and then reactions occur in the order determined by a single pass through the atoms in the gut. The effect is easily accounted for in the computation of environmental carrying capacities.

dynamics in this simulation. After 50,000 cycles, the population has evolved a collective behavior consistently superior to the random one (η~0.69).

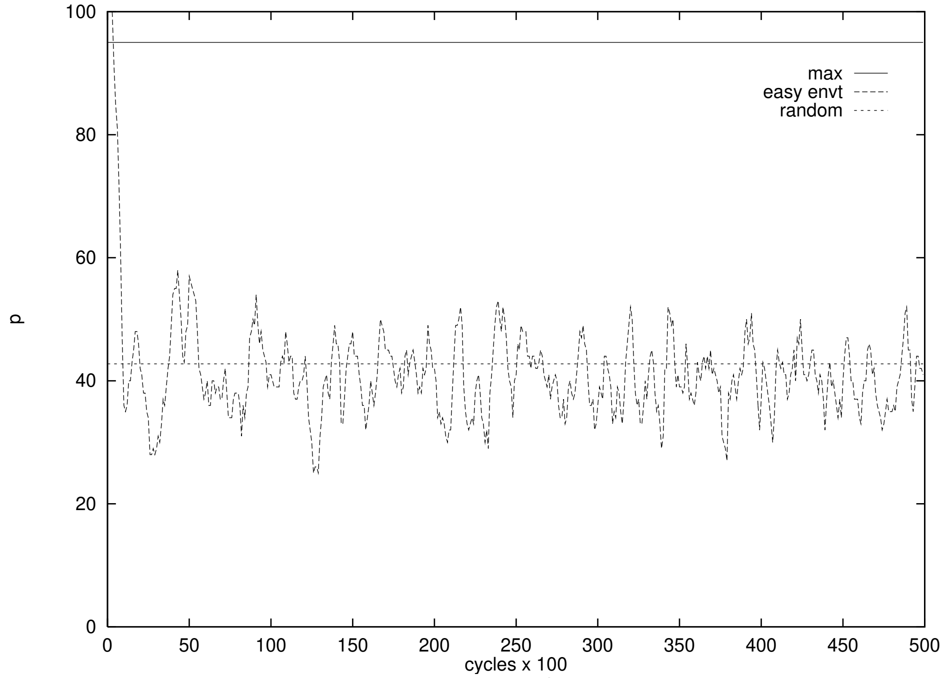

Figure 12: Population size in the easy (β=0.1) environment. The predicted carrying capacities corresponding to random and most efficient behaviors are also shown.

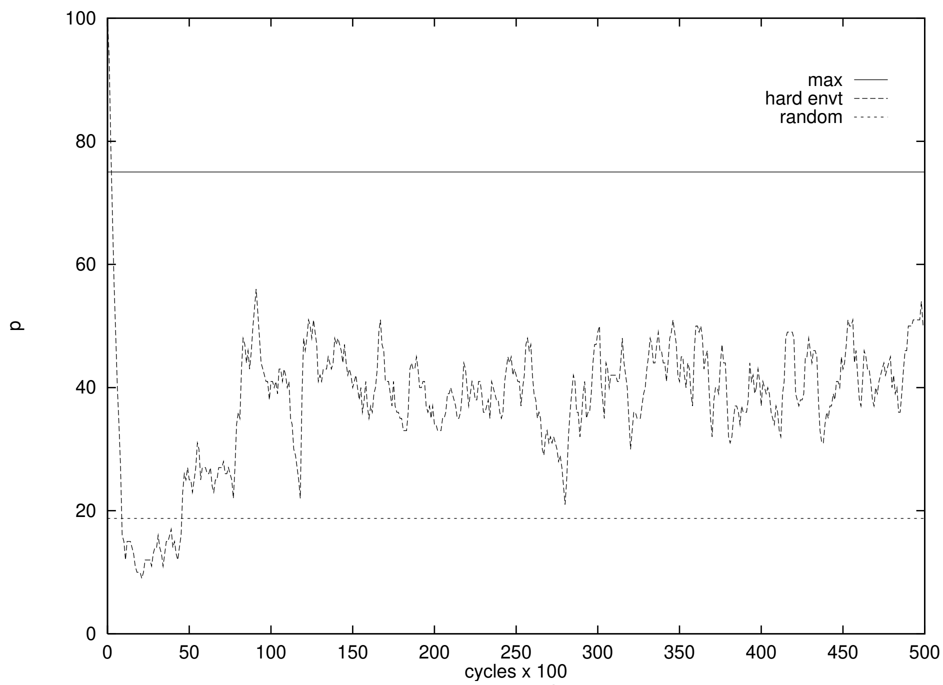

Figure 13: Population size in the hard (β=0.5) environment. The predicted carrying capacities corresponding to random and most efficient behaviors are also shown.

A more interesting statistics in this experiment is the biodiversity of the population, based on the species gene. We have plotted in Figure 14 the evolving biodiversity of the two simulations with easy and harsh environment. In the former the biodiversity goes to zero, which is an absorbing state due to the lack of mutations on the species gene. Therefore one of the two species has gone extinct for good. In the latter, however, the biodiversity is maintained at its maximum (not an absorbing state) showing perfect balance between the two species thanks to their capability to consume and replenish each other's resources.

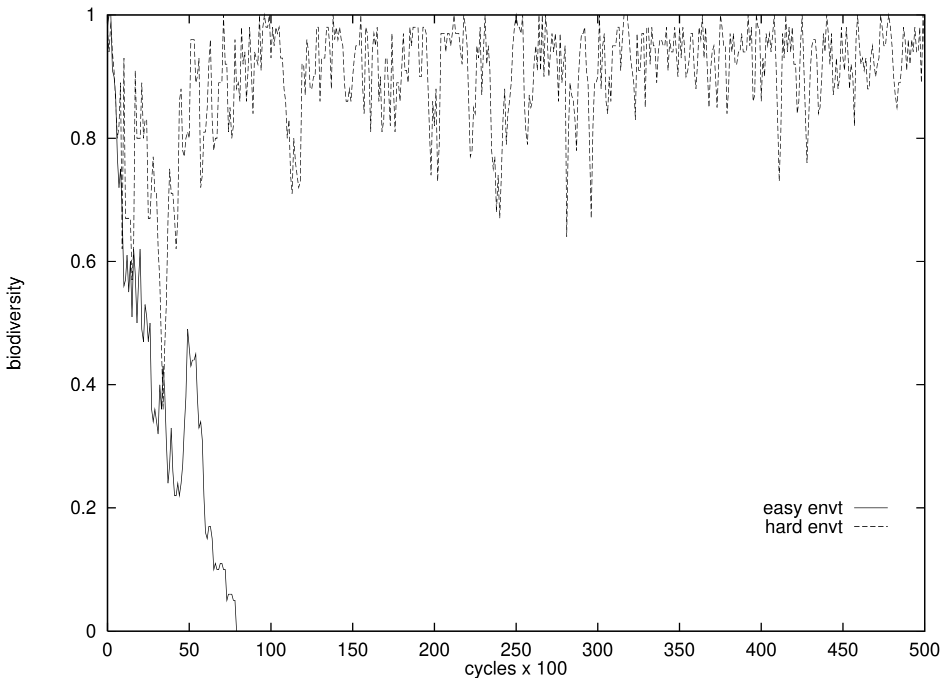

Figure 14: Biodiversity of the population in the two environments of the second experiment. This is measured based on the species gene and normalized so that 1, the maximum value, means that half of the total population belongs to each species; 0 means that all individuals belong to one and the same species, while the other has gone extinct.

We have arbitrarily killed the biodiversity in a variant of the hard environment simulation, by setting all species genes to the same allele after 50,000 cycles. This a poor man's model of some "environmentally unsound" external action, like an oil spill or some other disaster causing the fracture of the food chain. Figure 15 shows that the consequence is catastrophic; complete extinction follows swiftly.

We have extended the simulation in the harsh environment to compare the consequences of self-preserved vs. killed biodiversity on population dynamics, but also to see whether the population size can further increase toward the optimal level if given sufficient evolutionary time. The result of a 500,000 cycle-run (shown only in small part in Figure 15, for clarity) is a very slow improvement that eventually reaches a population size over 2.5 times larger than the random behavior level ($\eta \sim 0.78$).

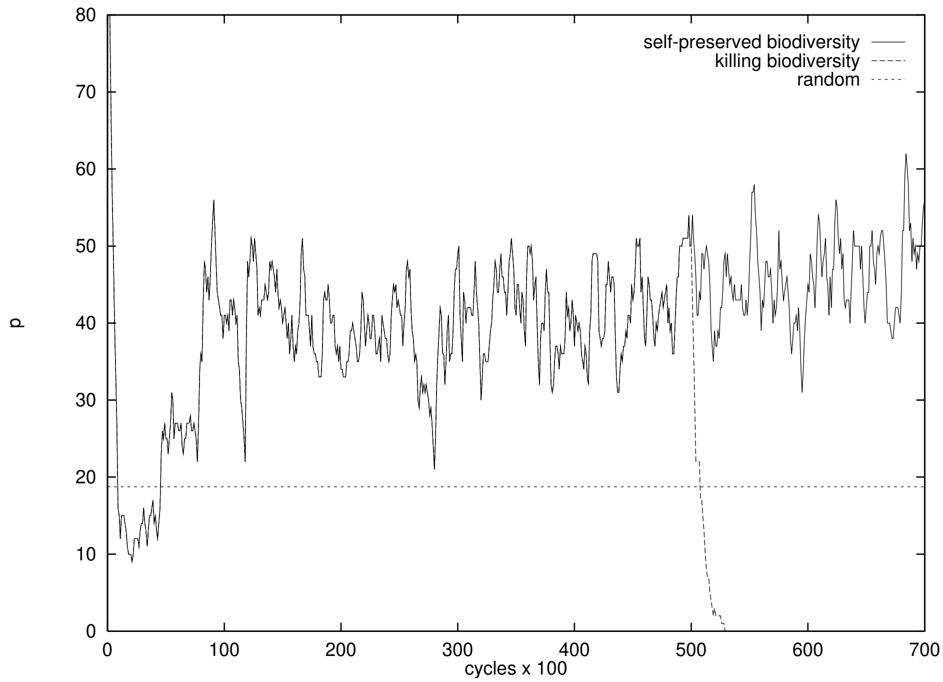

Figure 15: Population size in part of the extended simulations with the hard environment (cf. Figure 13). The normal case with evolving biodiversity is compared with the one in which the biodiversity is killed after 50,000 cycles (see text).

5.2.3 Discussion

This experiment also shows how the environment creates the selective pressures for the evolution of adaptive behaviors. Environmental complexity is measurable here by the ratio between the energetic outcomes of the different reactions. This heterogeneity of selective pressures has a direct impact on the carrying capacity of random behaviors: the more complex the environment, the harder to survive without adaptation.

When the environment is so benign that the random behavior has a carrying capacity close enough to that of the optimal behavior, there is not enough pressure for organism to improve their efficiency. Biodiversity in this experiment is not a genetic trait that can be selected for; it is a collective property of the population. Therefore it is subject to genetic drift in the course of evolution. Since zero is the only absorbing state — one from which the population cannot escape — the biodiversity drifts to zero after about 8,000 cycles, as shown in Figure 14. From this moment on, all individuals are of the same species, so only one element can provide positive energy. The carrying capacity of the optimal behavior, as it turns out, decreases to only about 11% higher than the random one, and predictably nothing else happens (Figure 12).

Harsher environments make a random population smaller. Just like too much environmental complexity prevented evolution in the previous experiment, too large a value for β would lead to extinction in this experiment. A moderate increase in selective pressure, however, causes better than random behaviors to evolve, as Figure 13 illustrates. At least part of the population catalyzes exothermic reactions, creating by-products that help the other species. This collective behavior, requiring the presence of both species, can be viewed as a form of mutualism and increases the carrying capacity thanks to the extra energy from by-products. This indirectly pushes biodiversity to a maximum, as it happens (Figure 14): each species needs the other.

In other words, we see two selective pressures reinforcing each other. The presence of individuals with efficient foraging behaviors in a certain area creates selective pressure in favor of individuals of whatever species is less frequent in that area (they have more food available), thereby increasing biodiversity. The ecological balance between species maintains carrying capacity at its maximum and thus creates selective pressure in favor of individuals with efficient foraging behaviors. This selective “loop” occurs locally in the environment, mediated by the resources, yet we observe a larger role of the individuals locally shaping each other’s environment.

Once this coupled adaptive process is “bootstrapped” by a sufficiently complex initial environment, it gives rise to behaviors that depend upon the diverse ecosystem. In fact, as Figure 15 shows, the biodiversity annulment has a much more catastrophic consequence in the complex environment — extinction. This phenomenon is due to the population no longer behaving randomly; each species relies on the existence of the other for providing the resources that can sustain its size. When those resources disappear, the carrying capacity of the environment drops dramatically and the violent population fluctuation that follows leads to extinction. Therefore the environment determines the necessity of mutualism and biodiversity. Ecological balance can provide for robust adaptation in complex environments.

5.3 Seasonality

5.3.1 Simulation setting

In the third experiment, we consider continuously changing environments. Under such conditions it makes no sense to speak of stable carrying capacity in the sense of Section 4; instead, we expect environmental fluctuations to drive the adaptive process in an equally continuous fashion. In particular, we want to model seasons and therefore choose environments characterized by periodic, sinusoidal patterns of change. To this end, consider a world with two uniformly distributed elements, with constant replenishment rate, that give off energy as follows:

$$A \rightarrow E_A = E_0 + \lambda \sin(2\pi t/\tau), \quad (18)$$

$$B \rightarrow E_B = E_0 - \lambda \sin(2\pi t/\tau), \quad (19)$$

without need of binary reactions and without by-products. Here t represents time (measured in cycles) and τ is a seasonal time constant that we will assume to be fixed throughout a simulation. We can think of the elements as two food sources, e.g., fruits, whose caloric contents vary with the passing of seasons, one reaching its high peak in the summer and the other in the winter. Note that if $\lambda > E_0$, each food becomes “poisonous” in its bad season. We are interested in observing how the population adapts to the seasons when either environmental resource alone is insufficient to guarantee survival.

Each individual has two contact sensors, signalling A and B atoms respectively in the facing cell. Because of "unary" reactions, there is no need for gut sensors. There are two hidden units and two outputs driving a binary motor. As in the previous experiment, ignoring input leads to the random behavior of always moving forward, eating all atoms. We can easily derive a "time-averaged" carrying capacity for this behavior and the seasonal environment, in the same form of Eq. (15), setting

$$E = E_A + E_B = 2E_0. \tag{20}$$

Behavior (20), however, neglects the complexity of the environment provided by the seasonal fluctuations.

At any given moment (except the instants when $E_A = E_B = E_0$) one element is more energetic, so one possible action would be for individuals to avoid the other element. As long as the other element remains exothermic, this is a less efficient behavior than the random one; but as soon as the other element becomes endothermic, such poison-avoidance strategy is a winner. It requires, however, that organisms can "track" the environment to tell the energetic content of an atom before ingesting it, a capability that sensors do not provide in our experiment. It is instead possible for the population (or a part of it) to converge on a behavior avoiding one of the two elements (say, B, the one that first goes through winter). This (sub)population has a carrying capacity again given by Eq. (15), but with

$$E = E_A. \tag{21}$$

The rest of the population (or a part of it) could in theory converge on the opposite behavior (say, avoiding A) with carrying capacity following

$$E = E_B. \tag{22}$$

Following Roughgarden [1971], we expect the adaptive process by endogenous fitness to favor individuals who are robust in the face of environmental change. Given the sensory limitations, and the fact that a single individual's behavior is fixed throughout its lifetime, selective pressure will shift seasonally on those individuals making the best use of the current season's resources. We therefore predict the formation of subpopulations of individuals adapted to the different seasons, shifting in relative size with seasonal changes.

5.3.2 Results

In this experiment we analyze two simulations. The amplitude of environmental oscillations remains constant in the first, a simpler condition that makes it possible to interpret the behavior evolved by the population. In the second simulation the oscillations become more and more violent, putting the robustness of the collective behavior to the test.

In the first simulation $\lambda = E_0$, i.e., the energetic value of each element has a minimum of zero. The situation is illustrated in Figure 16. The fact that the population oscillates, well correlated with A energy, indicates that it is differentiated. Inspection confirms that the following is going on: one subpopulation follows strategy (21) and a second one the random behavior (20). These are adaptive because both realize the higher energy provided by A atoms; individuals do not evolve behavior (22) because initially it is maladaptive to avoid energetic A atoms. Since all A is consumed, the population level is very close to the case in which there is no B; and since some B is consumed as well (by the subpopulation with random behavior), it is slightly higher. However, many B atoms are left over (by the subpopulation with behavior (21)) and later consumed (by behavior (20)) when their energy content is higher. This joint strategy by the two subpopulations turns out to be a very robust collective behavior, because the available resources are used efficiently. Another source of robustness is the energy contained in the individual reservoirs. The observed oscillations of its population average reflect seasonal fluctuations: energy accumulates in good times and then is spent slowly in bad times.

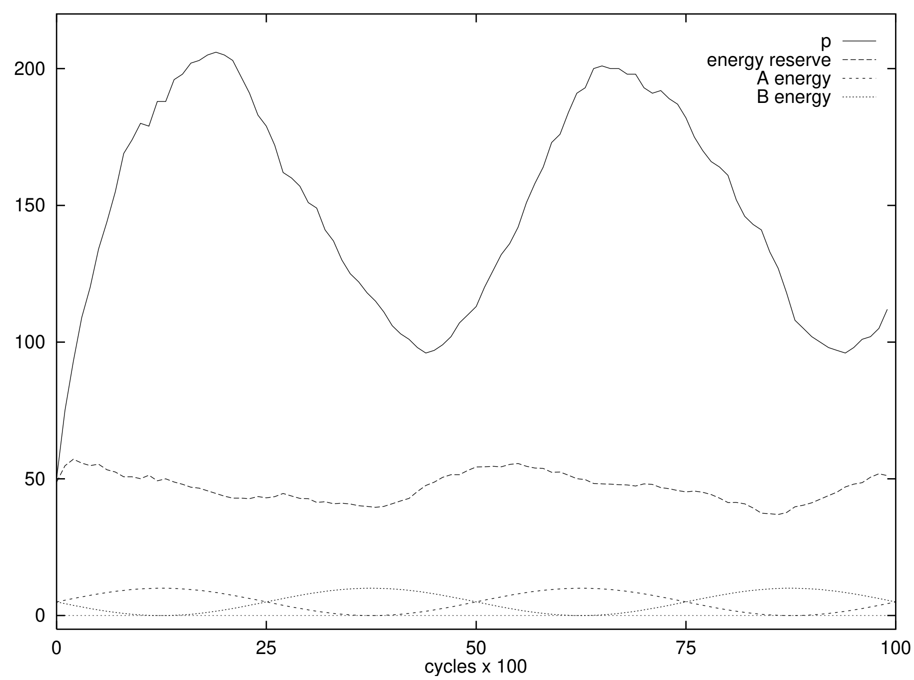

Figure 16: Results of the simulation with fixed-amplitude energy oscillations. The energy of the two elements is shown on an energy scale (bottom). E_0=5, as in the following simulation. The period is τ=5,000. A and B atoms have equal rates of replenishment. The resulting population size is plotted, along with the average energy reservoir levels.

In the second simulation $\lambda = \varepsilon t$, i.e., the amplitude of the energy oscillations in this environment increase linearly with time. The situation is illustrated in Figure 17. At first, when the environment is almost stationary ($\lambda << E_0$), the random behavior is good enough. As the environment becomes increasingly harsh, however, individuals evolve into subpopulations with differentiated behaviors just like in the previous simulation. Eventually the two elements become in turn poisonous and their ingestion extremely dangerous. Nevertheless, the population survives for a long time after $\lambda > E_0$, providing further evidence for the robustness of the evolved collective behavior. Finally, after 150,000 cycles, the environment is just too violent and the population goes extinct.

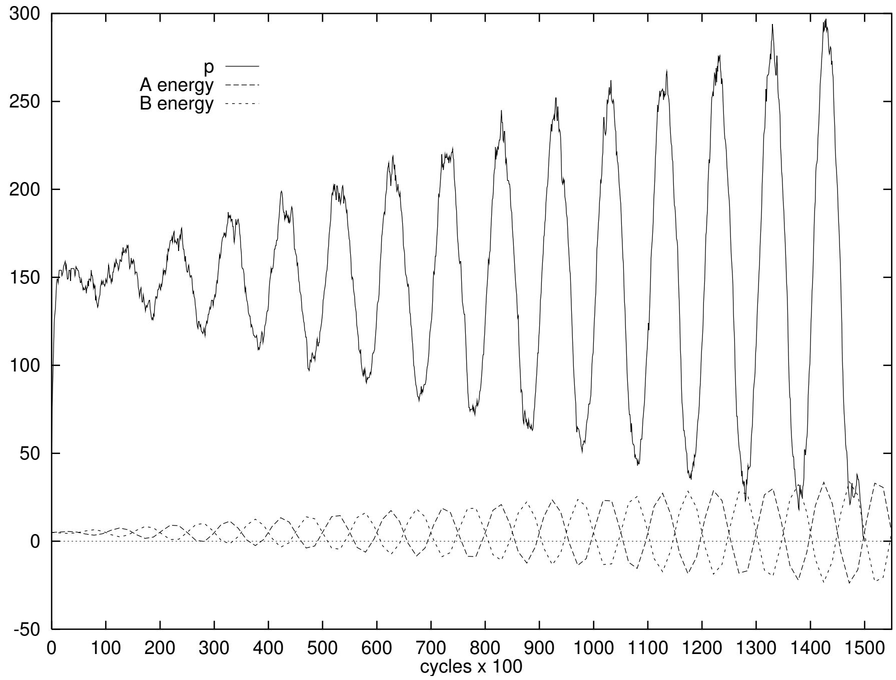

Figure 17: Population dynamics of the simulation with increasing-amplitude energy oscillations. The slope of the amplitude is ε=2x10^{-4} cycles^{-1}. The energy of the two elements, replenished at equal rates, is shown on an energy scale (bottom). The period is τ=10,000.

Figure 18 shows the average age of the population in the same simulation. Observing this changing life-history trait assists us in understanding the dynamics of the adaptive process. From 0 to about 20,000 cycles, age increases steadily. In this phase there is a transition from the initial uniform population with random behavior to a structured population. The following general trend of decreasing age is an indication of shorter lives caused by increasing environmental harshness. A more interesting point comes from inspecting the finer details of age dynamics. Between approximately 20,000 and 90,000 cycles average age fluctuates following a single wave form, anticorrelated with A energy and population size (cf. Figure 17). As we know by inspection, this indicates that the population has split into two subpopulations with behaviors (20) and (21), respectively; average age decreases when population increases (more A births than deaths) and vice versa. After about 90,000 cycles the age fluctuations become more convoluted and reveal a complexity that was not detectable from the population plot. In fact, two low peaks emerge during each period: new individuals are born during both A and B seasons. The explanation is that there are now three subpopulation. The environment has created selective pressure for behavior (22), as adaptive as (21), to emerge. Of course, all subpopulations must eat at least one element and thus eventually become extinct during the season in which it is poisonous.

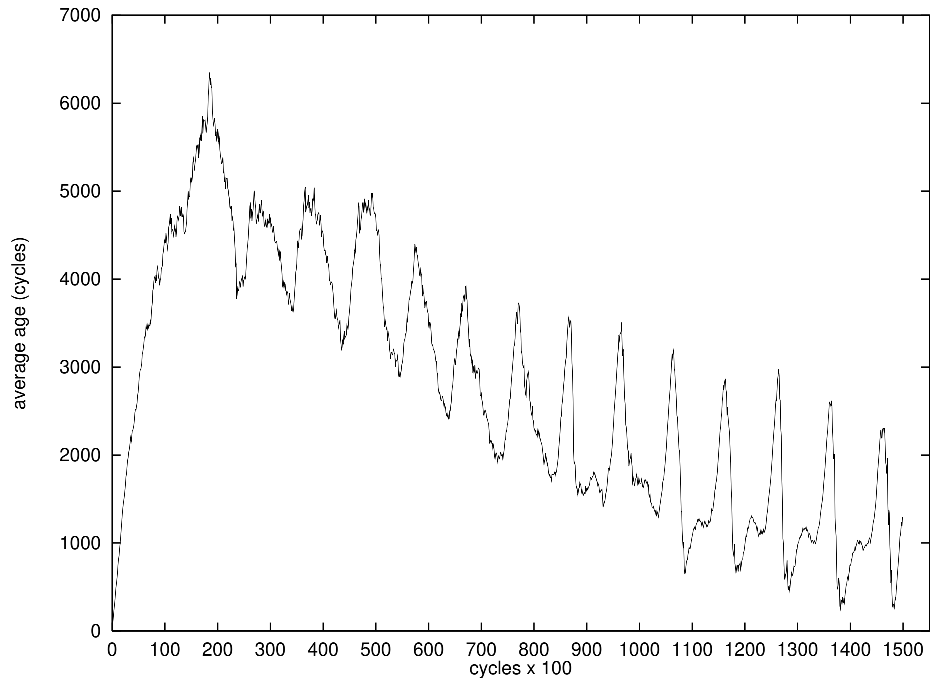

Figure 18: Average age of the population evolving in the environment with increasing-amplitude energy oscillations.

5.3.3 Discussion

In this experiment, the increase of environmental complexity (from stability to fluctuations of growing amplitude) drives the differentiation of individuals into a composite collective behavior. The complexity of the evolving population could be defined as the number of subpopulations, each with a unique individual behavior, it is composed of.

We first observe this (Figure 16) when two subpopulations evolve in a relatively benign environment with energies oscillating always above zero and with fixed amplitude.[9] The synergy of the two subpopulations provides a more robust collective behavior than either strategy alone could accomplish. Survival is guaranteed because the best possible use is made of available resources — given the limitations of the sensory system. It is worth noting that for a uniform population to accomplish the same, it would be necessary for evolution to adapt behaviors at each season, with an ever-young population always behaving optimally in its season and dying — replaced by a new generation — when the next season sets in. While this might have been possible had we used a much lower reproduction threshold, α, thus increasing the reproductive rate, it is virtually impossible with the parameters and mutation rates used in the experiment, due to the short duration ($\tau/2$) of the seasons. On the contrary, some individuals

[9] The simulation tool allows the experimenter to determine the composition of the populations easily, not only by inspection, but also using the analysis based on carrying capacity. Each behavior can be associated with a carrying capacity. By making a linear combination of these and fitting the resulting population curve to the simulation data, we can determine the coefficients that give the relative sizes of the subpopulations.

belonging to each subpopulation survive during their bad season, thanks to the energy accumulated in their internal reservoir during their good season. This is reminiscent of the way some mammals hibernate during low-resource seasons.[10]

Figure 17 illustrates a similar differentiation occurring when the amplitude increases with time and brings the oscillating energies below zero. Further analysis of the age statistics (Figure 18) has shown that as the fluctuations grow and the environment becomes less stable, the population splits again and a new subpopulation fills the one niche that had not yet been taken advantage of — energy from B atoms. This does not improve the overall carrying capacity, since B atoms were already being consumed by the other two subpopulations. In fact, no change is observed in the trend of the population curve (Figure 17). Yet we speculate that the emergence of the last behavior has the beneficial effect of prolonging the survival of the population, based on the observation that the long-term decreasing trend in average age becomes less severe in the last 40,000 cycles (Figure 18).

6 Conclusions

6.1 Related work

The work reported in this paper relates to several previous attempts to characterize environmental complexity in simulations of adaptive behavior. The LEE model itself was developed to overcome some complications observed in similar, more inclusive models [Holland 1992]. Spatial issues have been identified long ago as crucial in characterizing environmental complexity for both natural and artificial systems [Roughgarden 1971, Rössler 1974]. Characterizing the ecological consequences of physical constraints in artificial and natural environments remains an active area of research [Todd and Wilson 1993, Kareiva 1994].

Some of the features incorporated into the LEE model have been studied at much greater depth by others. We have used metabolism, coupled with environmental resource constraints, as a source of complexity in our simple experiments. The evolution of metabolism as a self-sustained complex chemical system is a fascinating open problem [Rössler 1972, Bagley, Farmer and Fontana 1992, Kitano 1994].

Local selection and more generally endogenous fitness [Mitchell and Forrest 1994] are different aspects of LEE that we consider key to modeling populations adapting to an environment — as opposed to explicitly optimizing some predetermined trait. Gould and Lewontin [1979] would assume that evolved behaviors are not optimal as a postulate. We instead don't deny the usefulness of optimality analysis within the LEE framework. In this paper we have indeed been able to assess an evolved behavior's adaptedness by comparing it against the carrying capacity benchmarks computed for random and optimal behaviors.

[10] Other, better solutions exist in the natural world. Insects, for example, only live for one season and then lay eggs that are dormant until the following year.

Previous works with LEE have focused on different issues of adaptive behavior, but are nonetheless related to environmental complexity. In one series of experiments, we simulated the coevolution of behaviors with morphological traits [Menczer and Belew 1994]. The hypothesis considered was whether non-Lamarckian learning could favor the evolution of informative sensors. The environment was crucial in that it provided the signals used for reinforcement by the learning algorithm — an idea previously explored by Ackley and Littman [1992]. It turned out that the environment was in fact an effective teacher, when the sensory apparatus was able to capture the external features whose correlation with fitness had to be discovered by evolution.

In another set of simulations, also dealing with interactions between phenotypic plasticity and evolution, we considered a more advanced model of (protocultural) learning [Cecconi, Menczer and Belew 1996]. Maturation age was an evolved life history trait; a significant delay of this age was observed when offspring were allowed to imitate their parents' behavior during immaturity. The result points to the role of behavioral improvements in determining the trade-off for optimal age at maturity. In that setting, a different source of environmental complexity is represented by the direct contact among individuals. Such forms of social interaction are not treated in this paper, where individuals interact exclusively through the mediation of environmental resources.

6.2 Relevance to artificial systems

Characterizing environmental complexity can be useful for applications in which a situated agent is to adapt to a real environment. When an agent has to learn appropriate actions by discovering important correlations between performance and environmental signals, the topology of the space biases input sequences and thus has tremendous influence on the time and space complexity of the learning algorithm [Kaelbling 1993]. The different complexity of "ideal" simulated worlds versus noisy, inconsistent, real environments is of particular relevance to evolutionary robotics [Nolfi, Floreano et al. 1994, Jacobi, Husbands and Harvey 1995]. Endogenous fitness also enables the self-evaluation of evolving robot behaviors. Careful modeling of environments where multiple agents can interact by sharing finite resources, or even more directly by signalling or learning from each other, is important for the design of effective agents who must carry out autonomous tasks on behalf of the user [Maes 1990]. The collective behavior of populations of agents sharing a physical environment is also the object of research in distributed robotics [Mataric 1992].

A different type of "virtual" worlds are another field of application of the ALife methods described in this paper. The fast pace of new developments in network technology has made distributed information sources an ideal environment for autonomous agents. In the words of Booker [1994], "Intelligent agents are needed to (...) locate information on wide-area computer networks. In applications such as these, an agent must interact with 'environments' in which spatial structure is not necessarily the most significant basis for organizing behavior." To respond to Booker's challenge, we have applied the endogenous fitness paradigm — indeed, the very algorithm of Figure

2 — to the design of populations of information agents [Menczer, Belew and Willuhn 1995]. The problem is to search and retrieve information relevant to some user query from a large distributed collection of documents such as the World Wide Web. This is a hard task for standard techniques based on manual navigation or indexing, due to the dynamic and distributed nature of the environment and its size. But simulations show that if the relevant documents in the Web are treated as finite and shared information resources, agents competing for them can evolve robust and efficient search strategies. The topology of the Web environment — a graph with nodes (documents) connected by hyperlinks — is, to be sure, quite different from the spaces modeled in this paper. It turns out, however, that by using metrics of locality and velocity appropriate to this environment, it is possible to characterize the Web according to descriptive features that are useful in both types of space, such as density, patchiness, rates of change, etc. This way we are attempting to analyze the algorithmic behavior of the evolving population of information agents and correlate their performance with the local topology of the search environment. One aspect that makes the Web a particularly amenable environment for this approach is the communication bottleneck of the underlying Internet. It is important for distributed algorithms to minimize communication, and this is achieved by populations whose individuals only interact by sharing resources.

The endogenous fitness algorithm of Figure 2 has also been applied to another problem — image compression — in which statistical properties of the environment have to be captured by an adaptive algorithm. Cecconi, Mancuso and Parisi [1996] have considered images as 2-D worlds in which a population of agents evolves by selecting areas rich with information that the user considers worthy of being kept. An extension of the JPEG compression algorithm uses the population to adapt compression factors to local environmental characteristics. Agents evolve collective behaviors that depend on the local texture of the image (e.g., follow edges); their spatial distribution ends up matching the features of the image. This approach saves up to 50% of storage robustly, i.e., preserving the most informative image areas.

6.3 Relevance to natural systems

This paper has discussed a computational model that permits analysis of the complexity of adaptive behaviors with respects to the environments in which those behaviors are evolved. Our results could be seen as quantitative explorations of the thesis put forward by Godfrey-Smith [1996]: “The function of cognition is to enable the agent to deal with environmental complexity.” More specifically, in our experiments, different measures of environmental complexity are dealt with by adaptive populations evolving different varieties of structure in their behaviors.

We believe that models like LEE can become powerful tools for studying a broad range of issues in ecological theory as well as theoretical, behavioral, and evolutionary biology. The methods that we have used to characterize different dimensions of complexity can be applied to understanding the dynamics of adaptive populations in other environments for which classic analytical approaches, such as systems of differential equations, are inadequate.

Ecologists recognize the importance of spatial models, in which relationships between members of a population and their environmental resources are mediated by the distribution of such resources in space [Milinski and Parker 1991]. Ideal free distributions, critical patch size, habitat variability, food gradients, metapopulations and niche adaptation are all examples of important spatial aspects of the environment [Roughgarden, May and Levin 1989]. There is a lack of models that account for the potential influences of space on ecological dynamics. Environmental consequences of the spatial dimension are often treated by adding *ad-hoc* variables into partial differential equation models. Recently, theoretical explorations of spatial effects have involved models of dispersal and patch turnover, reaction-diffusion, and cellular automata [Kareiva 1994]. However, experimental investigations to test the major hypotheses emerging from such models remain missing because field studies are difficult, expensive, or time-consuming. Simulations can assist in this regard.

Roughgarden et al. [1996] have discussed the role and usefulness of computer simulations in biological theory. They identify two classes of minimal models — ideas and systems — in which the contribution of computer simulation can be most effective and relevant for biologists. Idea models need not possess the plausibility of a biological theory; they only need to point to the possibility of a certain idea, which can in turn inspire biologists to devise experiments in the field to test for more formal hypotheses. Minimal system models, though still very simple, are more plausible and aim to provide an explanation for some observed phenomenon; their assumptions and consequences must be testable in the field, so that cooperation with biologists is more essential. LEE is a framework within which both idea and minimal system models can be designed easily, as we have tried to show by the three idea models proposed in this paper. What makes LEE and similar approaches suitable to deal with these issues is the notion of “situatedness:” resource generation and consumption, selection, and other environmental interactions all occur situated in space.

Another important feature of LEE populations and behaviors that makes them useful to model certain classes of natural systems is their robustness. For this reason we are using a variation of the algorithm of Figure 2 for modeling populations of parasitic individuals, such as cancer cells or AIDS viruses [Naviaux et al., in preparation]. The robustness of these populations is believed to lie in their capability to evolve adaptive mutants. The goal of such a model, from a therapeutic viewpoint as opposed to the experiments of this paper and simulations of adaptive behavior in general, is to find ways to most efficiently kill off the population of adapting parasites. In this case it is crucial to characterize the factors that make the environment — the host — harsh enough to reduce the diversity of the population and thus weaken its ability to withstand extinction.

6.4 Future directions

The LEE model has many limitations and possible extensions. To mention just a couple, recombination and development are missing in the current model. Asexual reproduction by cloning has several limits. Recombination may facilitate the adaptive process for many rich environments in which the number of possible behaviors is very large [Menczer and Parisi 1992a,b]. Sexual reproduction would also enable us to use LEE for studying the evolution of speciation; sexual selection models have been proposed in evolutionary biology for this important yet poorly understood problem [Laland 1994, Todd and Miller 1991]. In this paper we have touched on some themes — biodiversity, niche exploitation — that are closely related to speciation (cf. Sections 5.2, 5.3).

Another desirable consequence of modeling behaviors is that it demands a more elaborate characterization of the genotype/phenotype distinction [Belew 1993]. Such a developmental process would strengthen the role of environment, from selective pressure on phenotypes to direct interactions with ontogeny.

Minimal models of behaviors dealing with space, such as foraging, habitat selection, signalling, etc., are buildable with LEE without resorting to many *ad-hoc* assumptions restricting the range of observables. It is also easy to integrate the selective pressures determining within-niche adaptation and across-niche evolution. The environments and behaviors studied in this paper have given limited examples of this approach, that could be extended to include models of selection deriving from more direct interactions among individuals and populations, such as predation, mate choice, communication, and other social behaviors.

Sociality is the last dimension of environmental complexity. Direct interactions among organisms go beyond the characterization of environments grounded on physical space. Examples of social phenomena that we might hope to address [Roughgarden et al. 1996] include interactive signalling, cooperative equilibria, and optimal size of social colonies or other complex societies.

6.5 Summary

We hope that LEE can provide both a rich theoretical framework and a useful simulation tool not only for the ALife community, but also for biologists. In this paper we have demonstrated how it can be used to model selection pressures stemming from different environments and thereby evaluate the complexity of the corresponding behaviors. While not necessarily optimal, these behaviors can be regarded as adaptive in that they let individuals survive in the face of increasing environmental harshness.

We have characterized three forms of environmental complexity and shown how they create selective pressures for collective behaviors that match those environments. Gradual patchiness of resources causes efficient foraging behaviors to be evolved. Metabolisms supporting complementary food chains results in the preservation of biodiversity, but only in the presence of a sufficient carrying capacity differential in favor of mutual behaviors. Finally, dynamic environments with seasonal fluctuations divide the individuals into subpopulations with different behaviors exploiting seasonal niches. Simple demonstrations of these concepts have been achieved through simulation experiments taking advantage of endogenous fitness and local selection, the key features of latent energy environments.

Appendix: LEE software

The LEE software tool was developed to be used for efficient simulations within the framework of the LEE model. Its latest release and documentation is available via World Wide Web at URL http://www-cse.ucsd.edu/users/fil/ or anonymous ftp://cs.ucsd.edu/pub/LEE. The source code is copyright of the University of California and its use if free, except for commercial purposes. Copies of the software are also available at mirror sites on ALife Online, ENCORE/SAFIER, CMU/AI, and other repositories as well as on CD-ROM freeware.

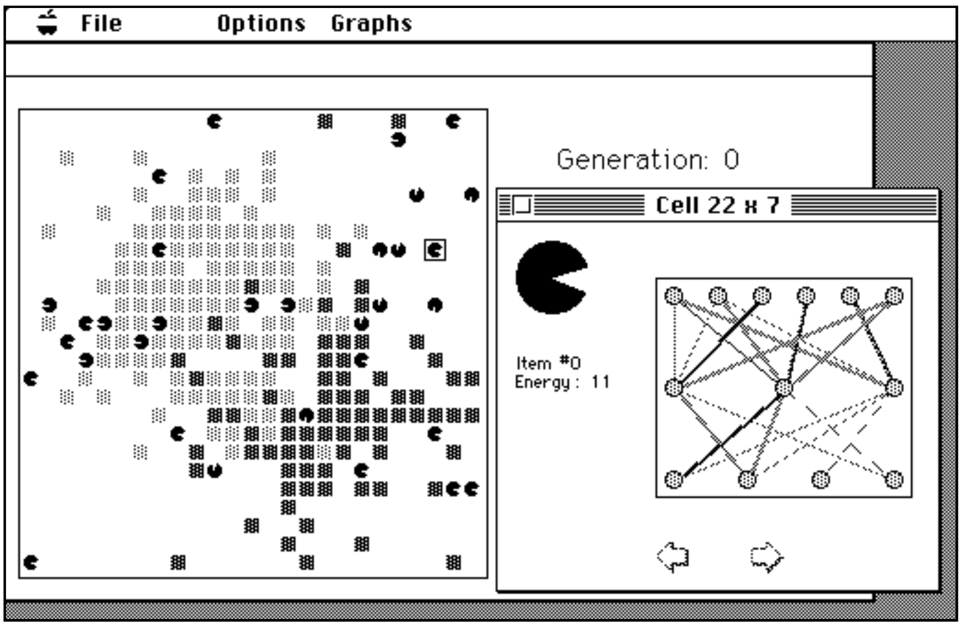

Figure A1: A typical interactive screen during a LEE simulation on the Macintosh.

The LEE package consists of approximately 7,000 lines of C code and runs on both UNIX and Macintosh platforms. The documentation contains directions on how to compile LEE for different systems. The only machine dependencies are in the input and output interfaces. While most of the default and user-driven I/O is implemented through files, an interactive graphical interface exists for the Macintosh and one might follow in the future for X Windows. The typical graphical output is shown in Fig. A1. Complete state files allow for checkpointing and continuations of experiments across platforms.

Acknowledgements

The ideas behind the LEE model have emerged from conversations with Domenico Parisi and Stefano Nolfi. We are particularly indebted for useful discussions and advice to Federico Cecconi, Peter Todd, an anonymous reviewer, and everyone in the GRAL group at the Institute of Psychology of CNR in Rome. This work was partly supported by a NATO fellowship to the first author.

References

Ackley, D. and M. Littman (1992) Interactions between learning and evolution. In C. Langton, C. Taylor, J. Farmer and S. Rasmussen (eds.)*Artificial Life II*. Reading, MA: Addison Wesley. 487-507

Bagley, R.J., J.D. Farmer, and W. Fontana (1992) Evolution of a metabolism. In C. Langton, C. Taylor, J. Farmer and S. Rasmussen (eds.)*Artificial Life II*. Reading, MA: Addison Wesley. 141-158

Belew, R.K. (1993) Interposing an ontogenic model between Genetic Algorithms and Neural Networks. In J. Cowan (ed.) *Advances in Neural Information Processing* (NIPS5). San Mateo, CA: M. Kaufmann

Booker, L.B. (1994) Editorial. *Adaptive Behavior* 3(1): 1-2

Brooks, R. and P. Maes, eds. (1994) *Artificial Life IV.* Cambridge, MA: MIT Press

Cecconi, F., S. Mancuso, and D. Parisi (1996) Edge eaters: An artificial life approach to image compression. To appear in Proc. Adaptive Computing in Engineering Design and Control (ACEDC'96). PEDC, University of Plymouth, UK

Cecconi, F., F. Menczer, and R.K. Belew (1996) Maturation and the evolution of imitative learning in artificial organisms. *Adaptive Behavior* 4(1): 27-48

De Jong, K.A. and J. Sarma (1993) Generation Gaps Revisited. In L.D. Whitley (ed.) *Foundations of Genetic Algorithms 2*. San Mateo, CA: Morgan Kaufmann

Godfrey-Smith, P. (1996) *Complexity and the Function of Mind in Nature.* Cambridge University Press

Gould, S. and R. Lewontin (1979) The spandrels of San Marco and the Panglossian paradigm: A critique of the adaptionist programme. *Proceedings of the Royal Society of London* B205: 581-598

Holland, J.H. (1992)*Adaptation in Natural and Artificial Systems*. Cambridge, MA: MIT Press

Huberman, B.A. and N.S. Glance (1993) Evolutionary Games And Computer Simulations. *Proc. Natl. Acad. Sciences (USA)* 90 (16): 7716-7718

Jacobi, N., P. Husbands, and I. Harvey (1995) Noise and the reality gap: The use of simulation in evolutionary robotics. In F. Moran, A. Moreno, J.J. Merelo and P. Chacon (eds.) *Advances in Artificial Life.* Proc. Third European Conference on Artificial Life. Berlin, Germany: Springer-Verlag. 704-720

Jones, T. and M. Mitchell (1993) Introduction to the ECHO model. Santa Fe Institute. Working paper 93-12-074

Kaelbling, L.P. (1993) *Learning in Embedded Systems.* Cambridge, MA: MIT Press

Kareiva, P. (1994) Space: The Final Frontier for Ecological Theory. Special Feature, *Ecology* 75:1

Kitano, H. (1994) Evolution of metabolism for morphogenesis. In R. Brooks and P. Maes (eds.)*Artificial Life IV.* Cambridge, MA: MIT Press. 49-58

Laland, K.N. (1994) On the evolutionary consequences of sexual imprinting.*Evolution* 48(2): 477-489

Maes, P. (1990) *Designing autonomous agents: theory and practice from biology to engineering and back.* Cambridge, MA: MIT Press

Mataric, M.J. (1992) Designing emergent behaviors: From local interactions to collective intelligence. In J.-A. Meyer, H.L. Roitblat, and S.W. Wilson (eds.) *From Animals to Animats,* Proc. Intl. Conf. on Simulation of Adaptive Behavior. Cambridge, MA: Bradford Books (MIT Press)

Menczer, F. and R.K. Belew (1996) Latent Energy Environments. In R.K. Belew and M. Mitchell (eds.) *Adaptive Individuals in Evolving Populations: Models and Algorithms.* Reading, MA: Addison Wesley

Menczer, F. and R.K. Belew (1994) Evolving Sensors in Environments of Controlled Complexity. In R. Brooks and P. Maes (eds.)*Artificial Life IV.* Cambridge, MA: MIT Press. 210-221

Menczer, F., R.K. Belew, and W. Willuhn (1995) Artificial Life Applied to Adaptive Information Agents. Working Notes of the AAAI Spring Symposium on Information Gathering from Heterogeneous, Distributed Environments. Palo Alto, CA: Stanford University

Menczer, F. and D. Parisi (1992a) Evidence of hyperplanes in the genetic learning of neural networks. *Biological Cybernetics* 66(3): 283-289

Menczer, F. and D. Parisi (1992b) Recombination and Unsupervised Learning: Effects of Crossover in the Genetic Optimization of Neural Networks. *Network* 3: 423-442

Milinski, M. and G. A. Parker (1991) Competition for resources. In J. R. Krebs and N. B. Davies (eds.) *Behavioral Ecology* (3rd Ed.). Oxford, UK: Blackwell

Mitchell, M. and S. Forrest (1994) Genetic Algorithms and Artificial Life. *Artificial Life* 1(3): 267-289

Naviaux, R., F. Menczer, J. Hilton, and R.K. Belew (in preparation) Computational models of adaptive mutations in populations of cancer cells and AIDS viruses.

Nolfi, S., D. Floreano, O. Miglino, and F. Mondada (1994) How to evolve autonomous robots: Different approaches in evolutionary robotics. In R. Brooks and P. Maes (eds.)*Artificial Life IV.* Cambridge, MA: MIT Press. 190-197

Nolfi, S., D. Parisi, and J.L. Elman (1994) Learning and evolution in neural networks. *Adaptive Behavior* 3(1): 5-28

Rössler, O.E. (1974) Adequate Locomotion Strategies for an Abstract Environment -- A Relational Approach. In M. Conrad, W. Guttinger, and M.D. Cin (eds.) *Physics and Mathematics of the Nervous System.* New York, NY: Springer-Verlag. 399-418

Rössler, O.E. (1972) Design for autonomous chemical growth under different environmental constraints. In R. Rosen and F.M. Snell (eds.) *Progress in Theoretical Biology.*

Roughgarden, J., A. Bergman, S. Shafir, and C. Taylor (1996) Adaptive Computation in Ecology and Evolution: A Guide for Future Research. In R.K. Belew and M. Mitchell (eds.) *Adaptive Individuals in Evolving Populations: Models and Algorithms.* Santa Fe Institute Studies in the Sciences of Complexity. Reading, MA: Addison Wesley

Roughgarden, J., R.M. May, and S.A. Levin, eds. (1989) *Perspectives in Ecological Theory.* Princeton University Press

Roughgarden, J. (1971) Density-dependent natural selection. *Ecology* 52(3): 453-468

Rumelhart, D.E. and J.L. McClelland, eds. (1986) *Parallel Distributed Processing: Explorations in the Microstructure of Cognition*, Vol. 1. Cambridge, MA: Bradford Books (MIT Press)

Stearns, S.C. (1992) *The Evolution of Life Histories.* New York, NY: Oxford University Press. 20-38

Todd, P. M. and G. F. Miller (1991) On the sympatric origin of species: Mercurial mating in the Quicksilver Model. In R. K. Belew and L. B. Booker (eds.) Proc. Fourth Intl. Conf. on Genetic Algorithms. San Mateo, CA: Morgan Kaufmann. 547-554

Todd, P.M. and S.W. Wilson (1993) Environment Structure and Adaptive Behavior from the Ground Up. In J.-A. Meyer, H.L. Roitblat, and S.W. Wilson (eds.) *From Animals to Animats 2,* Proc. Second Intl. Conf. on Simulation of Adaptive Behavior. Cambridge, MA: Bradford Books (MIT Press). 11-20

Wilson, S.W. (1991) The Animat Path to AI. In J.-A. Meyer and S.W. Wilson (eds.) *From Animals to Animats,* Proc. First Intl. Conf. on Simulation of Adaptive Behavior. Cambridge, MA: Bradford Books (MIT Press). 15-21