LATENT ENERGY ENVIRONMENTS

F. Menczer and R. K. Belew

Cognitive Computer Science Research Group
Computer Science and Engineering Department
University of California, San Diego
La Jolla, CA 92093-0114
{fil,rik}@cs.ucsd.edu

Abstract. A novel ALife model, called LEE, is introduced and described. The motivation lies in the need for a measure of complexity across different ALife experiments. This goal is achieved through a careful characterization of environments in which different forms of energy are well-defined. A steady-state genetic algorithm is used to model the evolutionary process. Plastic organisms in the population are modeled by neural networks with non-Lamarckian learning during life. Behaviors are shown to be crucial in the interactions between organisms and their environment. Some emerging properties of the model are illustrated and discussed.

1. Introduction

From the tradition of artificial intelligence, the relationship between a plastic individual and the environment within which it must operate is often viewed in relatively simplistic, stimulus-response terms. The environment provides input on the basis of which the individual then responds. Perhaps because it has grown out of the same tradition, the Genetic Algorithm (GA) [Holland 1992] often relies upon an equally simplistic view of the relationship between evolving populations and their environment: the GA generates individuals of a population and the environment must only return the fitness of each of these individuals. The GA can then be viewed as a function optimizer.

In spite of the many reasons to recommend this evolution-as-search perspective, any such model fails to address one of evolution's most obvious characteristics: its *creative* power. One very useful insight arising from an artificial life (ALife) perspective is that the above dichotomy is replaced by a more holistic view of organism and environment as part of a single system to be modeled. In order to allow the necessary evolutionary creativity, the central problem facing every designer of ALife environments is the specification of an artificial task admitting strategies of arbitrary and unanticipated complexity. Unfortunately, this coupling of organisms with environments has created a major methodological problem within ALife research: results that report behaviors of different organisms in different environments are incommensurable, making it very hard to decide whether an observed difference in adaptation is really significant, or it is only due to one of the environments being more *complex* [Rössler 1974, Wilson 1991, Belew 1991].

Some models that try to address the environmental complexity issue, such as *Echo* [Holland 1992], are so inclusive that their results are very hard to interpret. Other recent works have began to analyze simpler environments in open-ended evolutionary simulations [e.g., Todd and Wilson 1993]. The next section will introduce our approach, called "latent energy environments" (LEE), which attempts to extend the characterization of complexity to a more general class of environments, while maintaining the analysis manageable by carefully enforcing appropriate physical constraints. A view of the model is offered in Figure 1. The arrow labeled "behavior" represents all types of interactions between organisms and environments, described in section 3. The arrow labeled "evolution" represents, roughly, all genetics-mediated interactions among organisms. Section 4 will describe the GA used to model such evolutionary process. Section 5 illustrates some general emerging properties of the model, while the final section exemplifies some of our objectives for the use of LEE.

2. Modeling environmental complexity

Our central goal is the specification of a series of ALife environments of graduated complexity. To accomplish this we begin by replacing the standard model of "food" as a spatially localized element of the environment required for survival, 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 base elements results in an energy gain or loss, plus some other by-product element(s). Let the base elements belong to an artificial set of types $\{A_1, A_2, ..., A_t\}$. Furthermore, let elements of each of these types be generated according to some spatiotemporal distribution. A reaction will be denoted as:

$$A_i + A_j \rightarrow E_{ij} + P_{ij}$$

where E_{ij} is the resulting energy ($E_{ij}>0$ for exothermic reactions, $E_{ij}<0$ for endothermic ones) and P_{ij} indicates one or more by-products. All the possible reactions can be represented by the entries of a symmetric reaction table like the one in Figure 2. This table can be viewed as that part of the LEE model which is neither directly nor indirectly affected by adaptation, and remains constant throughout the entire evolutionary process. In other words, it represents the laws of physics and chemistry, which by definition do not change over time.

We can now define two critical features of this artificial world. First, let us define a sequence, or chain, of reactions. For any world configuration, there may be many possible combinations of existing elements according to the possible reactions in the table. 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. Eventually, the chain terminates if no possible reactions exist among the remaining elements. Every sequence of reactions has a corresponding potential energy, given by the arithmetic sum of energies released and/or absorbed by all the reactions in the sequence. For each world configuration, we call *latent energy* the set of potential energies corresponding to all reactions chains available to organisms starting from that configuration. By controlling the rate at which elements of each type are introduced, we can regulate the amount of potential energy available in its latent form, while actual energy can ultimately be realized only upon dynamically choosing a reaction sequence.

Second, the amount of *work* required to release energy is defined in terms of the distance one of the elements in a reaction must be moved in order to occupy the cell of the other. This is the correct metric for evaluating organisms who must harvest energy to survive, and whose fundamental behavior, as will be shown in the next section, is movement. By controlling the spatial distributions of element types in the world, we can dynamically regulate the amount of work required to combine them. Similarly to the way we have associated potential energy to a reaction sequence, we can also associate to a chain its corresponding work. Subtracting this from latent energy and maximizing over all possible reaction sequences, we can estimate optimal energy efficiencies. In short, we have control and quantitative knowledge of how much latent energy is available in a world, how much work must be done in order to realize this energy, and how difficult it therefore is for an organism (or ecology of organisms) to survive!

These definitions 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. Different species may exploit different reaction chains in the same environment --

but we should expect the most efficient one to prevail if their uses of atomic elements are mutually exclusive.

3. Modeling the life process

Each organism in the evolving populations of LEE has a "brain" and a "body". The former interacts with the environment external to the body through a sensory-motor system: it receives input corresponding to the organism's stimuli and produces output controlling its behaviors. 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. The sensory information is elaborated by the brain to produce an output, interpreted as an action (movement) in the world and executed by a set of motors that make up the motor system. Motors may have different characteristics such as energy cost, power, orientation, accuracy, etc. Figure 3 illustrates an organism's body structure with a typical sensory-motor system. The latter provides a strong coupling between environmental complexity and difficulty of the survival task [Menczer and Parisi 1992].

3.1. Behavior

One of the fundamental features of LEE is that energy in an environment can only be released through the *behaviors* of organisms. Behaviors induce reaction sequences, thus catalyzing the transformation of latent energy into usable energy. The notion of behavior immediately binds an organism to its environment in a much more intrinsic way than allowed by simple fitness-function GA models.

Ethologists have long argued that behaviors be treated as first-class phenotypic characteristics of an organism. A second desirable consequence of modeling behaviors, then, is that it demands a more elaborate characterization of the genotype/phenotype distinction [Belew 1993]. 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, cross-over, etc.) and the phenotype's fitness (i.e., number of offspring) ultimately evaluated by the environment. The appropriateness of 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.

We take *movement* to be our canonical behavior, in part because this overt action is most easily observed. However, certain other actions on the part of the modeled organisms are also part of the model and may be under an organism's discretionary control. One such example is the ingestion of elements in the world: an organism can only carry around a limited number of elements in its "gut" in order to cause reactions, so it may have to discriminate among elements to ingest along its path.

3.2. Learning

We wish to allow the behavior of an organism to improve over the course of its lifetime as a consequence of its experience, i.e., to *learn*. Littman [1994] provides a survey of architectures for learning in evolutionary frameworks. One approach is to use a well-studied type of neural network (NNet) as our computational model of an organism's brain. The LEE context helps to distinguish the particular type of NNet learning of interest to us.

First, we make no assumptions about additional "teaching" input being available to the organism, beyond the stimulus of the organism's environment. This would seem to preclude many successful "supervised" learning techniques, such as back-propagation of error [Rumelhart, Hinton and Williams 1986]. Such error correction algorithms require that the NNet be given information about the "correct" output/action. Not only are these correct actions unknown, but even attempting such a normative stance violates the spirit of creative discovery we treat as central.

Fortunately, if the organism is simply forced to *predict* the expected outcome of its actions, differences between its expectations and the actual outcome can generate the same sort of error information, without any additional teacher [Sutton 1988]. Prediction learning has been used successfully in a number of ALife simulations [Nolfi, Elman and Parisi 1990, Parisi, Cecconi and Nolfi 1990], and the NNets used in LEE will exploit this predictive device. Figure 3 shows a possible organism's NNet architecture[1]. Note that only a part of the weights of the last connection layer (marked in dark) is learnable through the prediction task.

Besides learning to predict good actions, organisms controlled by NNets must also survive in order for them to reproduce and for the population to evolve. For example, an organism could generate perfect predictions simply by choosing to make no action whatsoever (assuming for the moment a static environment in which it is the only actor). The goal of accurate prediction is therefore secondary, subordinate to the organism's primary goal of surviving.

The second important characteristic of our NNet learning situation is that the organism is an *active* selector of its own training experience. One aspect of this problem is to perform "critical experiments" that most rapidly allow the learning system to converge on a consistent explanation of the world. However, an experiment which may immediately help to identify an important pattern may also be deadly to perform. Holland [1992] has characterized this as the "exploitation/exploration" dilemma, where an organism's ability to exploit already identified regularities must be balanced against the acquisition of new information.

4. Modeling the evolutionary process

The LEE model uses a *steady state* GA [De Jong and Sarma 1993] to simulate the evolutionary process. This means that successive generations are interleaved through time, shifting the control of differential reproduction from the experimenter to the adaptive process itself. Rather than being determined by the arbitrary ranking mediation of an explicit fitness measure, selective pressure results directly from intrinsic competition for finite resources within the LEE population. The only currency for such resources in the model is the energy latent in the environment. This characteristic places LEE in the class of *endogenous fitness models* [Mitchell and Forrest 1994].

For each time step, every organism executes a basic life cycle[2]. The operations performed in one such cycle are outlined in Figure 4. This simple steady state version of the GA has a few important consequences. The most important one is that the population size does *not* remain constant throughout an experiment. We consider this an important feature of the model, better founded biologically than the fixed-population GA. As shown in the next section, the population size can become stable spontaneously when the environmental conditions allow it (and quite robustly indeed), without this being imposed externally by the model. Another biologically plausible consequence of having a variable population size is that *extinction* can occur.

Finally, the LEE evolutionary process is *noisy*, for two reasons: the high variance caused by the steady state GA, and the fact that every parameter for which there is no value motivated by the model is treated as a random variable. Conforming to such stochastic-oriented principle allows the LEE model to be free of unwanted bias from the experimenter. Noise is the cost paid for this: results may be harder to interpret, and adaptation may be slowed down. Furthermore, stochasticity may cause many types of drift in the dynamics of population genetics. Once again, we are satisfied that all these effects are well present in nature and thus the LEE model is strengthened by its stochastic features as well.

---

[1] Other unsupervised learning architectures have been used in the LEE framework and described elsewhere [Menczer and Belew 1994].

[2] 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 bias of the sequential simulation.

5. Emerging features

While classical GAs constitute the standard paradigm for adaptation viewed as an optimizing process, endogenous fitness models are being used increasingly often as a paradigm for open-ended evolution [Brooks and Maes 1994]. However, the creative power of such models has exposed the lack of appropriate instruments of analysis [Mitchell and Forrest 1994]. In this section we report on some emerging properties of our simulations, to illustrate that the task of analyzing endogenous fitness models is greatly facilitated by relying on well-founded constraints such as the conservation of energy.

5.1. Carrying capacity

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: we identify this maximum population size with the "carrying capacity" of the environment. For example, extinction indicates that the carrying capacity of the environment is insufficient to support a population large enough to withstand stochastic fluctuations.

LEE allows to quantitatively estimate optimality 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. Consider the following reaction energy table for the simple case of an environment with only two types of atomic elements, a and b:

$$\begin{array}{c|cc} & a & b \\ \hline a & -\beta E & E \\ b & E & -\beta E \end{array} \tag{0}$$

where

$$\begin{aligned} &E > 0 \\ &|\beta| < 1. \end{aligned} \tag{1}$$

and there are no by-products. Table (0), along with the spatiotemporal distributions of the elements, determines the environmental complexity of the survival task. The minimal set of elements for alternative behavioral strategies is $\{a,a,b,b\}$. In fact, with these elements an organism can catalyze the following sets of reactions:

$$(a+a),(b+b) \to -2\beta E \tag{2}$$

$$(a+b),(a+b) \to +2E \tag{3}$$

where (3) is clearly a more advantageous strategy than (2), given conditions (1). If r is the expected rate of replenishment for both a and b elements, i.e.,

$$\begin{aligned} &r_a \equiv \# a \,/\, lifecycle \\ &r_b \equiv \# b \,/\, lifecycle \\ &r_a = r_b \equiv r \end{aligned}$$

then a set $\{a,a,b,b\}$ is produced every $2/r$ lifecycles. Thus strategies (2) and (3) produce energy changes

$$\Delta E_2 = \frac{-2\beta E}{2/r} = -r\beta E \tag{4}$$

$$\Delta E_3 = \frac{2E}{2/r} = rE \tag{5}$$

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

$$\langle population \rangle = const \equiv p \tag{6}$$

$$\langle \Delta E \rangle = 0. \tag{7}$$

Since energy is always conserved, the only consumed energy is that lost in the form of work, that is, used up for moving in the world. If we call c the cost per move for any organism and use (6), 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. Let us then introduce a probability distribution over strategies. This is quite simple in the present example, with only two strategies: let us call η the probability of strategy (3), so that $(1-\eta)$ is that of strategy (2). Then the average energy produced per unit time by the population is, using (4) and (5):

$$\eta \Delta E_3 + (1-\eta)\Delta E_2 = r\left[\eta E + (\eta - 1)\beta E\right]$$

so that we can finally rewrite (7):

$$\langle \Delta E \rangle = rE\left[\eta + \eta\beta - \beta\right] - pc = 0. \tag{8}$$

Equation (8) provides the link between optimality of behavioral strategies, expressed through the probability distribution over catalyzed reactions, and population size. If the former is known, we can solve (8) for the expected population size p:

$$p = \frac{rE}{c}(\eta + \eta\beta - \beta). \tag{9}$$

In particular, the case of optimal behavior $\eta=1$, corresponds to the maximum sustainable population, i.e., the carrying capacity:

$$p_{\max} = \frac{rE}{c}. \tag{10}$$

The converse case is useful for estimating the optimality of a population's behavior at equilibrium, by measuring the population size and solving (8) for η:

$$\eta = \frac{\frac{pc}{rE} + \beta}{1+\beta}. \tag{11}$$

As an illustration of this simple analysis, Figure 5 plots population size versus time in three simulations with different entries in a reaction table like (0). Note the damped oscillations (plus noise), in agreement with population dynamics models. In the steady state regime, 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 predicted size for a population of random walkers (obtained from (9) with $\eta=1/2$):

$$K = p_{\eta=1/2} = \frac{rE}{2c}(1-\beta). \tag{12}$$

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 walk population smaller than the amplitude of the fluctuations, so extinction occurs rapidly. Larger K values result in different stable population levels; the carrying capacities of the two environments, however, are also different (from (10)). So we can substitute the measured populations into (11) to find that the two performances are not significantly different ($\eta \sim 1/2$, the random walk value, 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 (8) holds only at equilibrium and for this simple example: the more general non-equilibrium case and more complex environments will yield systems of differential equations, which may be too hard to solve analytically. The important point is that carefully designed latent energy environments allow to maintain an accurate connection between ALife models and simulations.

5.2. Fitness measures

We have claimed by equation (9) that the population size in a LEE simulation is an effective measure of the average optimality of the organisms' behaviors. We want to show how reliable population size really is in estimating population fitness. Since the environmental resources are finite, fitness is *density dependent*; furthermore, since organisms interact only by indirectly competing for the resources, carrying capacity is the best measure of fitness [Stearns 1992]. Let us consider an environment (simpler than the one described in table (0)) with just one element, where each atom is associated with energy E (again, no by-products). Equation (6) still holds at equilibrium, and equation (7) for energy conservation yields in this case:

$$K \equiv p = \frac{rE}{c} \tag{13}$$

where K is defined as the carrying capacity. The difference between (12) and (13) is that, in the former, fitness corresponds to one of the behaviors made possible by the binary reactions, while in the latter, the fitness is purely density-dependent, i.e., behavioral effects are negligible with respect to the environmental carrying capacity. Equation (13) predicts perfect correlation between population size and density-dependent fitness. Figure 6 illustrates how well the population size measured in LEE simulations with different carrying capacities (determined by rE) matches the predicted fitness.

When evolved behaviors allow organisms to make a more efficient use of the latent energy, an increase in average age and in population size is observed. The situation is illustrated in Figure 7, where measures are from LEE simulations with a reaction table like (0). During these phases (6) does not hold, environmental complexity is no longer the sole fitness factor, and thus carrying capacity is neither a sufficient predictor of population size nor a satisfactory fitness measure. Similarly, other biological fitness measures which ignore density dependence, such as *net reproductive rate*, are poorly correlated with population growth in our experiments. Population size and average age, on the other hand, can still be used as fitness measures at the population level. Their increase is well correlated with *expected reproductive success*, a fitness measure useful in biology with density dependence. Therefore, as long as energy is conserved, density dependence provides a key to analyzing behaviors of evolving populations in endogenous fitness models.

5.3. Age distribution

Carrying capacity arguments permit to predict other emerging features of our model. For example, LEE has no control over the age of organisms. It would therefore be possible for organisms to stay alive and keep reproducing forever, introducing a sort of "evolutionary inertia" that would weaken the effect of selective pressure and thus slow down adaptation. The negative evolutionary consequences of immortal organism have been addressed by Todd [1993].

However, we observe that the distribution of age in the LEE population quickly becomes *stable*, according to the Euler-Lotka classic demographic equation. When fitness is density dependent --a consequence, as we have seen, of sharing finite environmental resources-- the population reaches zero-growth and the Euler-Lotka equation predicts that the age distribution becomes *stationary*: this is in agreement with the age distributions measured during a LEE simulation and shown in Figure 8. Furthermore, in this situation, the fraction of individuals with age x becomes a direct measure of the probability of survival to age x [Stearns 1992]. The latter by definition is a non-increasing function of age, again in agreement with the measured age distributions (Figure 8). This explains the fact that, while the average age of the LEE population grows initially, in the course of evolution the organisms begin to die "spontaneously" at a finite age. Therefore we argue that immortals will not evolve in endogenous fitness models, as long as energy is conserved.

6. Current and future directions

The LEE model has been implemented to run simulations on UNIX and Macintosh platforms [Menczer and Belew 1993]. Figure 9 shows a portion of screen in a typical interactive simulation[3]. We believe that LEE represents both a rich theoretical framework and a useful simulation software for the ALife community. This chapter has shown some of its emerging features relevant to population dynamics, aging, and fitness. In the near future we intend to use it to study several other issues.

The first experiment we have proposed for LEE has the objective of studying the evolution of sensory systems. The question is whether efficient encodings of information from sensors should facilitate learning and thus emerge by means of an influence of learning on evolution [Hinton and Nowlan 1987, Nolfi, Elman and Parisi 1990, Belew 1990]. The idea for the experiment comes from a simulation where NNets must look for food while moving in an environment with different zones [Miglino and Parisi 1991]. Analysis of the interaction between evolution and learning in that setting leads to a non-linear relation between fitness expectancy and sensory information about zones. This inefficient sensory encoding seems to make the task unnecessarily difficult to learn, and suggests that evolution should favor an encoding that would facilitate learning. To test such hypothesis, the simulation with zones has been recast into the LEE framework -- good evidence for its flexibility, allowing to capture different ALife experiments while evaluating their complexity. Organisms in a population adapting in a latent energy environment use a very simple sensory system that is allowed to evolve. Simulation results are in agreement with our hypothesis [Menczer and Belew 1994].

In another set of experiments [Menczer 1994], dealing more directly with the adaptive advantages of individual plasticity, the LEE model and simulator have been used to study the evolutionary interaction between temporal environmental changes and phenotypic plasticity. Both were modeled quite crudely, the former by random changes in the reaction table and the latter by a form of "developmental noise". The adaptive rate of such noise was found correlated with the rate of changes in the environment, building partial quantitative support for the claim that plasticity is advantageous for populations evolving in non-stationary environments. Yet another direction of current research is the evolution of reproductive maturity in the presence of cultural transmission by imitation [Cecconi *et al.* in preparation].

In the near future, by extending the model to include sexual reproduction and development, we intend to address further topics such as morphology, speciation, and multiple-species population dynamics. It is our hope that the LEE framework can serve as a computational tool to investigate complex problems in theoretical biology.

---

[3] C source code and documentation for release 1.* is available from URL http://www-cse.ucsd.edu/users/fil/ or ftp://cs.ucsd.edu/pub/LEE/.

Acknowledgments

This project originated from conversations with Stefano Nolfi and Domenico Parisi at the Institute of Psychology of the National Research Center in Rome, Italy. Stefano Nolfi, Jeff Elman, Greg Linden and Federico Cecconi have assisted in writing the code. Thanks to William Hart, Federico Cecconi, Alex Guazelli, and Jonathan Roughgarden for reviewing earlier drafts of this chapter, Mark Land for discussion on Figure 6, and everyone in the Cognitive Computer Science Research Group for helpful discussions on many issues of the model. Finally, we thank SFI for support during the Workshop.

References

Belew RK (1990) Evolution, Learning, and Culture: Computational Metaphors for Adaptive Algorithms, *Complex Systems* 4, 11-49

Belew RK (1991) Artificial Life: A Constructive Lower Bound for AI, *IEEE Expert* 6-1, 8-15

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

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

Cecconi F, Menczer F, and Belew RK (in preparation) Learning and the evolution of age at maturity

De Jong KA and Sarma J (1993) Generation Gaps Revisited, in: Foundations of Genetic Algorithms 2, Whitley LD, ed., San Mateo, CA: Morgan Kaufmann, 19-28

Hinton GE and Nowlan SJ (1987) How Learning Can Guide Evolution, *Complex Systems* 1, 495-502

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

Littman ML (1994) How evolution can guide learning, in this volume

Menczer F (1994) Changing Latent Energy Environments: A Case for the Evolution of Plasticity, Technical Report CS94-336, University of California, San Diego

Menczer F and Belew RK (1994) Evolving sensors in environments of controlled complexity, in: Proc. *Artificial Life IV*, Brooks R and Maes P, eds., Cambridge, MA: MIT Press

Menczer F and Belew RK (1993) LEE: A Tool for Artificial Life Simulations, Technical Report CS93-301, University of California, San Diego

Menczer F and Parisi D (1992) Recombination and unsupervised learning: effects of crossover in the genetic optimization of neural networks, *Network* 3, 423-442

Miglino O and Parisi D (1991) Evolutionary stable and unstable strategies in neural networks, Proc. IJCNN, Piscataway, NJ: IEEE, Vol. 2, 1448-1453

Mitchell M and Forrest S (1994) Genetic algorithms and artificial life. *Artificial Life* 1:3

Nolfi S, Elman JL, and Parisi D (1990) Learning and Evolution in Neural Networks, Technical Report CRL-TR-9019, University of California, San Diego

Parisi D, Cecconi F, and Nolfi S (1990) Econets: Neural networks that learn in an environment, *Network* 1, 149-168

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

Rumelhart DE, Hinton GE, and Williams RJ (1986) Learning internal representations by error propagation, in: *Parallel Distributed Processing: Explorations in the Microstructure of Cognition*, Rumelhart DE and McClelland JL, eds., Cambridge, MA: Bradford Books (MIT Press), Vol. 1, Ch. 8

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

Sutton RS (1988) Learning to predict by the methods of temporal differences, *Machine Learning* 3, 9-44

Todd PM (1993) Artificial death, in: Pre-proc. Second European Conf. on Artificial Life, Brussels, Belgium, May 1993

Todd PM and Wilson SW (1993) Environment structure and adaptive behavior from the ground up, in: *From Animals to Animats 2*: Proc. Second Intl. Conf. on Simulation of Adaptive Behavior, Meyer J-A, Roitblat HL, and Wilson SW, eds., Cambridge, MA: Bradford Books (MIT Press), 11-20

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

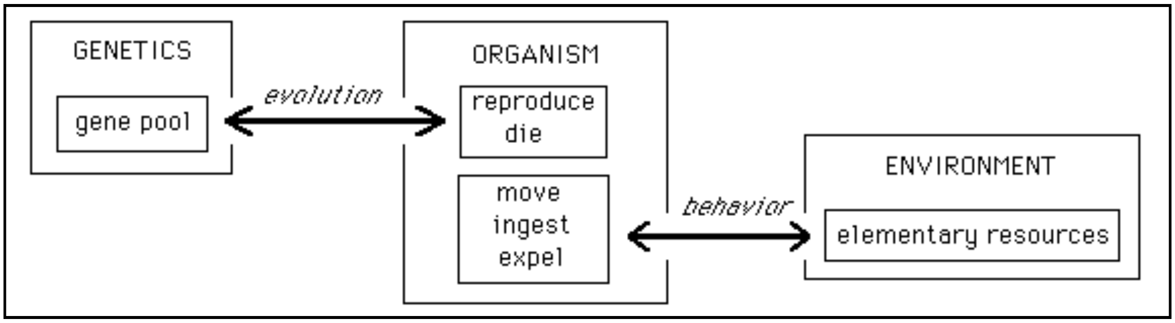

Figure 1: LEE model.

	A_1	A_2	A_i	A_t
A_1	0	$+1, A_3$		$+3$
A_2		-2		$-1, A_t$
A_j			E_{ij}, P_{ij}	
A_t				0

Figure 2: example of a reaction table.

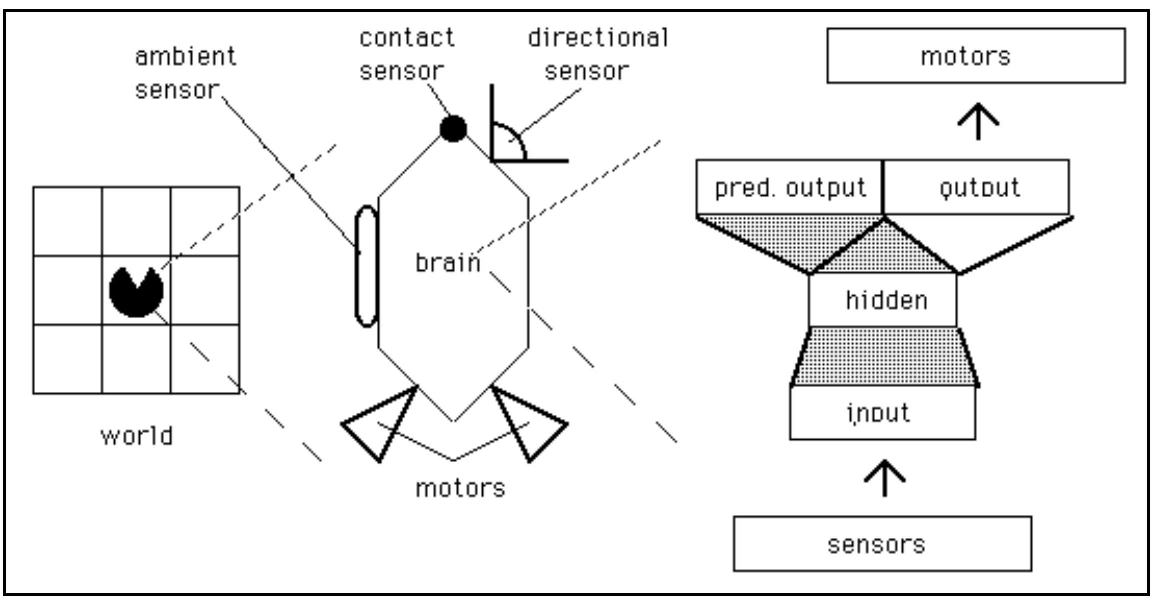

Figure 3: architecture of a LEE organism.

```
for each time step {
      for each alive organism {
            sense world;
            feed forward activations;
            move;
            learn;                  /* changes phenotype        */
            digest;                 /* catalyze reactions       */
            if (energy > α) {
                  reproduce;        /* copy genotype            */
                  mutate;           /* change new genotype      */
                  develop;          /* get new phenotype        */
            }
            else if (energy < ω) die;
      }
      replenish world;
}
```
Figure 4: pseudo-code for the main loop of the LEE GA. The constants α and ω are related to the energy distribution over the population: at initialization, organisms are given random energy uniformly distributed in the interval [α,ω]; at reproduction, energy is conserved by parent and offspring splitting the parent's energy evenly.

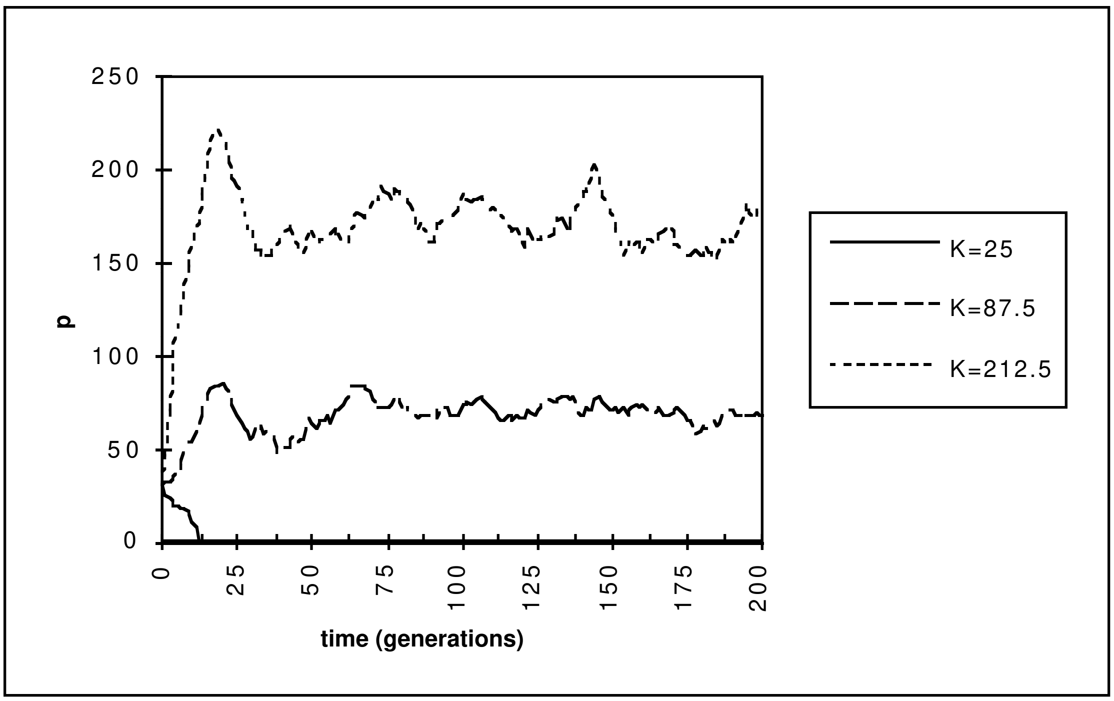

Figure 5: carrying capacity for three simulations with different reaction tables. K is a measure of latent energy: when enough latent energy is supplied by the environment to avoid extinction, an early equilibrium is reached.

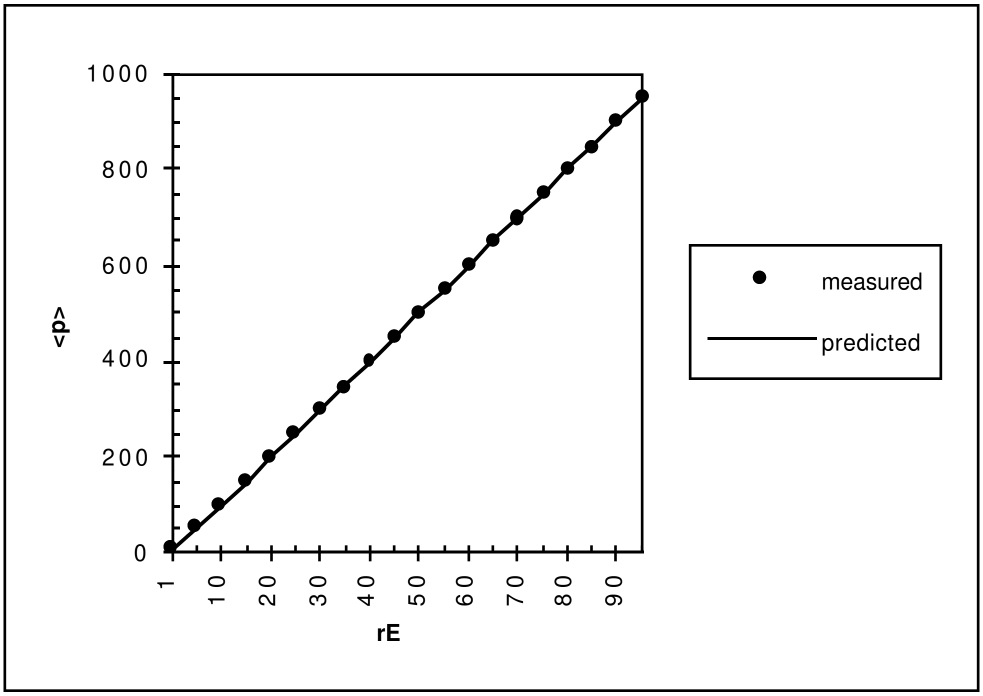

Figure 6: correlation of population size with density-dependent fitness. Measures of population size are time averages in the steady-state regime; predicted population size is from equation (13).

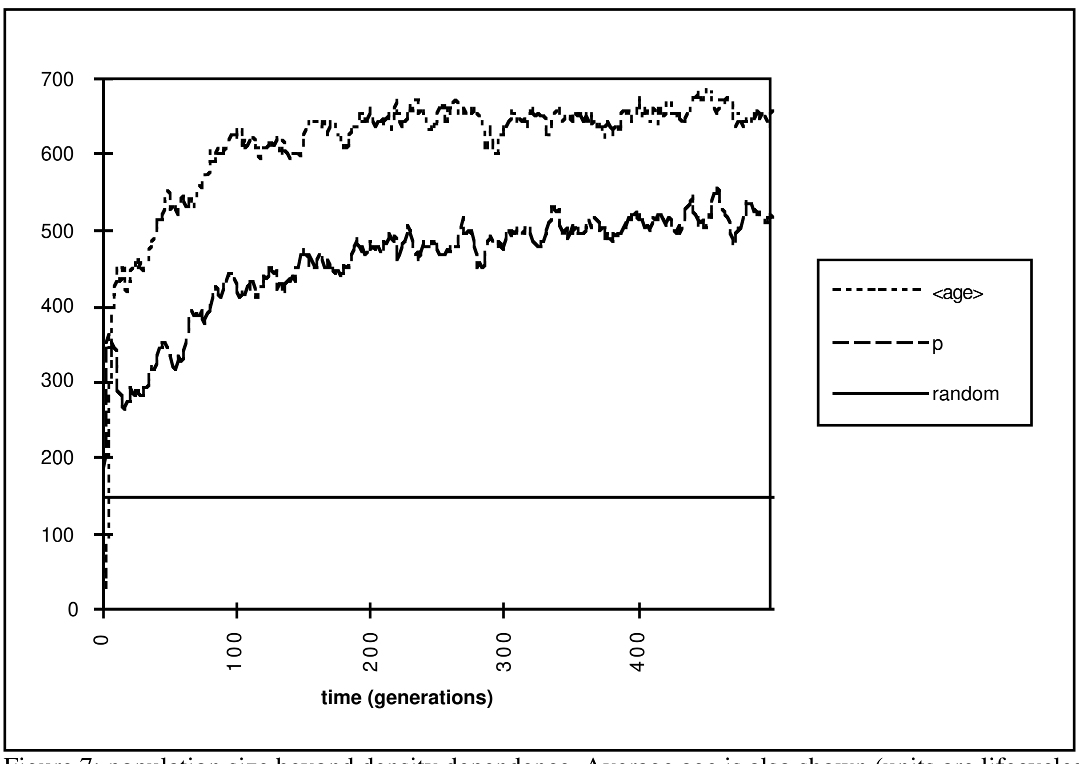

Figure 7: population size beyond density dependence. Average age is also shown (units are lifecycles); for sake of comparison, random refers to the size of a population of random walkers (from equation (12)).

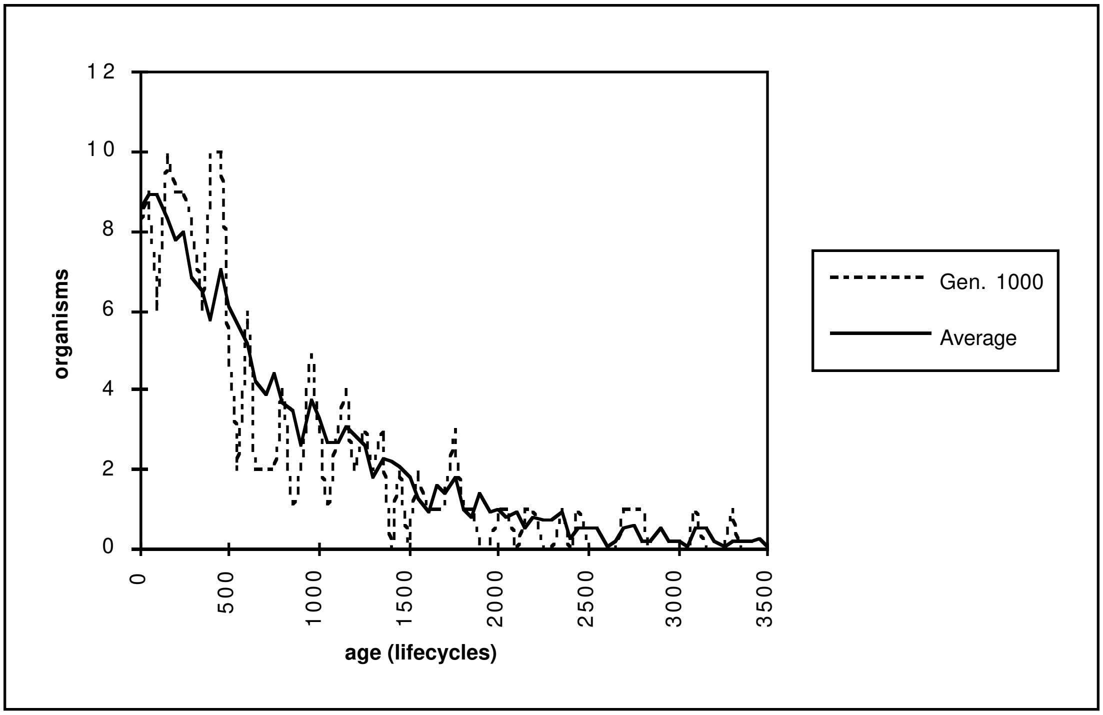

Figure 8: histograms showing age distributions. Average refers to a time average of bin values between generations 1000 and 10000.

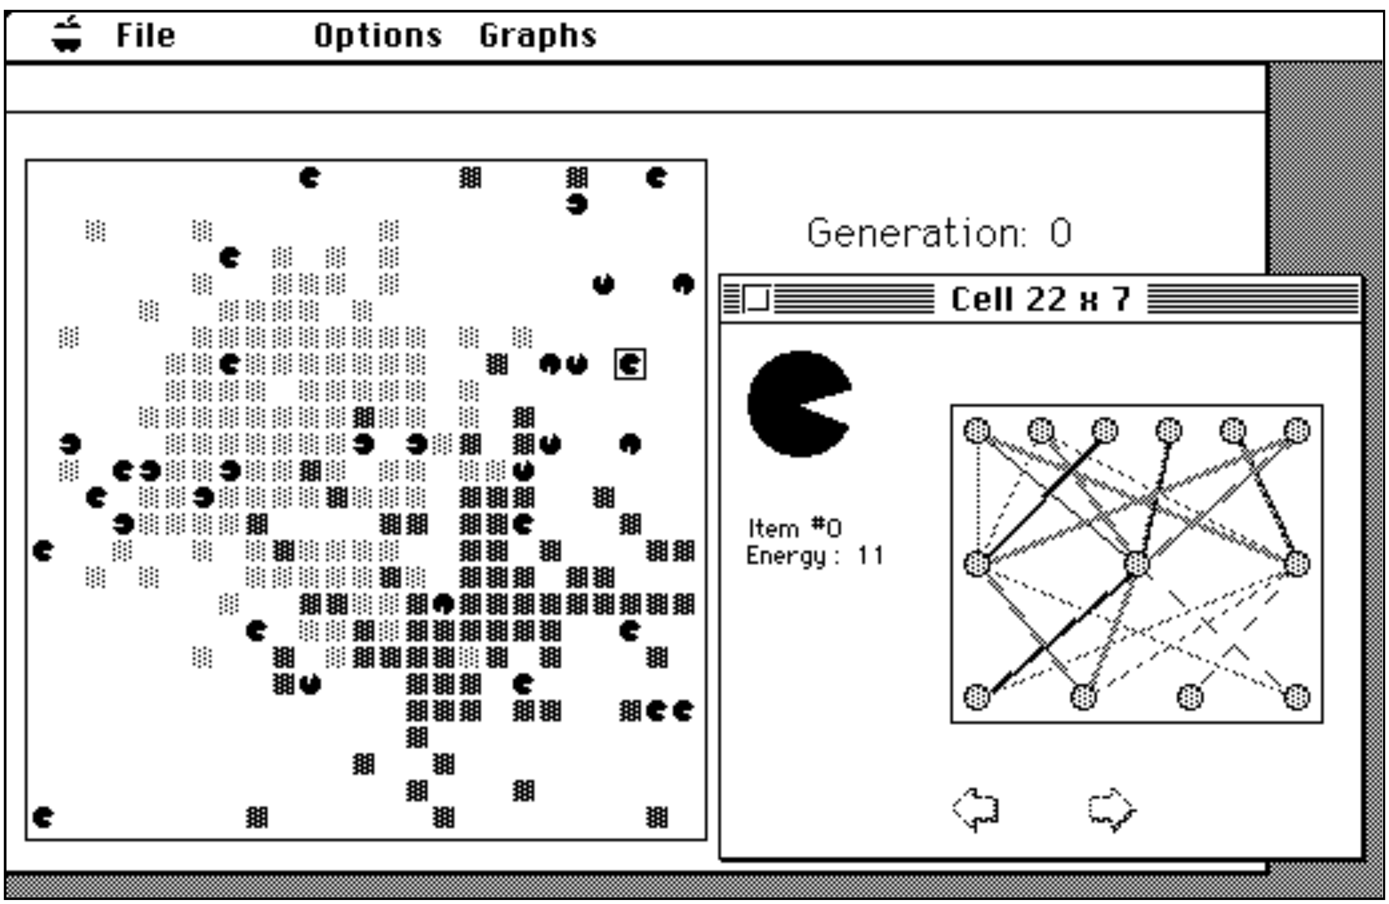

Figure 9: interactive screen during a LEE simulation on the Macintosh.