

The blind breeding the blind: Adaptive behavior without looking

Peter M. Todd, Stewart W. Wilson

The Rowland Institute for Science
100 Edwin H. Land Boulevard
Cambridge, MA 02142 USA
ptodd@spo.rowland.org
wilson@smith.rowland.org

Anil B. Somayaji

Department of Mathematics
MIT
Cambridge, MA 02139 USA
somayaji@mit.edu

Holly A. Yanco

MIT AI Lab
545 Technology Square
Cambridge, MA 02139
holly@ai.mit.edu

Abstract

Sensors and internal states are often considered necessary components of any adaptively behaving organism, providing the information needed to adapt a creature's behavior in response to conditions in its external or internal environment. But adaptive, survival-enhancing behavior is possible even in simple simulated creatures lacking all direct contact with their environment — evolutionarily shaped blind action may suffice to keep a population of creatures alive and reproducing. In this paper, we consider the evolution of the behavioral repertoires of such sensor-less creatures in response to environments of various types. Different spatial and temporal distributions of food result in the evolution of very different behavioral strategies, including the use of looping movements as time-keepers in these otherwise cognitively-challenged creatures. Exploring the level of adaptiveness available in even such simple creatures as these serves to establish a baseline to which the adaptive behavior of animats with sensors and internal states can be compared.

1 Introduction

Adaptive behavior is usually thought of as behavior that can change in response to conditions in an organism's external or internal environment, with the result that the organism's survival is enhanced. As Meyer and Guillot (1991, p. 2) have put this, "In a changing, unpredictable, and more or less threatening environment, the behavior of an animal is adaptive as long as the behavior allows the animal to survive." Almost universally, sensory systems are assumed to be the organism's link with conditions in its external environment, allowing it to respond adaptively to the situation in which it finds itself. Adaptive behavior can also be generated as a consequence of internal conditions, as Meyer and Guillot point out; homeostatic mechanisms that sense an organism's internal state with reference to some desired fixed point can cause

behavior that will lead to a return to that fixed point. These sorts of adaptive responses to environmental conditions, the kinds of things that nervous systems and computer circuits do, form the basis of research on adaptive behavior. Without some way of knowing what's out (or in) there, without knowledge of the prevailing conditions to which behavior *should* adapt, how can there be any adaptive behavior?

In this paper, we construct a series of "changing, unpredictable, and more or less threatening environments" populated by simple simulated creatures that have *no* sensors nor internal states, and yet still manage to survive and prosper in their worlds. Through our simulation system, *SPO* (Simulation of Primitive Organisms), we demonstrate that, at least according to the definition given in the previous paragraph, adaptive behavior *is* possible without any knowledge of the environment. True, the behavior of any individual creature is fixed throughout its lifetime and does not change in response to any internal or external stimuli, so that in the strictest sense the use of the term "adaptive" is arguable. But the creatures still behave in a way particularly suited to the environment they inhabit, so that we may say that their behavior is at least "adapted," if not adaptive.

In actuality, the behaviors of the creatures in the SPO system *do* change over time in response to the environment — but only *between* individual lifetimes, rather than within them. Instead of shaping behavior through the adaptive processes of development, learning, or sensor-guided action, all of which bring about changes in the behavioral responses of an individual, we employ only the adaptive process of evolution, which brings about changes in the aggregate behavioral responses of a population of individuals. The individuals in our system each act blindly in their world (surrounded by a population of other simultaneously-acting individuals), choosing randomly according to some unchanging weighted probability distribution among a small set of actions, including eating, moving, and reproducing. Individuals who find and eat the food-energy in their world can survive and produce slightly modified offspring, while those that run out of energy will die. Over the course of time, the evolutionary process of descent with modification and selection through competition will result in

a gradual tuning of the behavioral repertoires of individuals in the population. After this adaptive process has run long enough, the behaviors of individual creatures will be adapted to their particular environment.

One of the simplest ways to show that the behavior of these creatures is indeed adapted to their particular environment is to evolve several populations under *different* environmental conditions, and look for evidence that the final adapted behaviors differ as well. This is the approach we take here, constructing a variety of very simple environments, each of which is composed solely of some pattern of food placed at various locations in a square grid. The spatiotemporal structure of these environments is determined by a small set of parameters that control food growth and distribution. We then investigate how different environmental structures lead to the evolution of different adapted behavioral strategies in individual creatures. We are also interested in discovering what environmental differences do *not* lead to differences in evolved behaviors — that is, the structural invariances that the adaptive process of evolution ignores. Separating aspects of environmental structure into those that matter for the creatures (e.g., the distribution of food in the world) and those that do not (e.g., the color of the food) is of central importance both in applying the results of this work to understanding natural systems, and in designing adaptively behaving artificial systems in light of the critical features of their operating environments.

This work thus fits into the general research program of characterizing the important features of environment structure in terms of the adaptive behavior they elicit, as outlined in Todd and Wilson (1993). By beginning with the simplest possible form of adaptive (or adapted) behavior, that of blind stateless creatures, we can establish a baseline against which the adaptiveness of sensors and internal state or memory systems will be much clearer in comparison. In the rest of this paper, we cover the results of this first foray in this research direction, beginning in the next section with a discussion of similar work on environment/creature simulations. In section 3, we present the way the world works, and the various parameters that control environment structure in this system. Section 4 covers the types of actions that creatures can perform, and the way they are chosen amongst at each time-step. In section 5, we present the results of many runs of this system, first discussing the process of adaptation of the behavioral repertoires over time (in section 5.1), and then going into more detail on the nature of the final adapted behavior patterns themselves (in section 5.2). Finally, in section 6, we draw conclusions and indicate the next steps to be taken with this work.

2 Past Work

Before going into the details of our system, it is important first to review the design and goals of recent related work on the simulation of creatures behaving in more or less naturalistic environments. Most of the systems surveyed here use “gridworld” environments composed of a square grid of locations (rather than continuous-space models) that can contain

various objects in addition to the creatures. This overview should make apparent how our system differs in motivation and emphasis from past work.

Among systems that use gridworlds to model animats (artificial creatures) interacting with an environment, a significant number place only one creature in an environment at a time. The environment defines a fitness function that is used for some sort of genetic search over the space of possible behaviors of individual creatures. Systems of this type include Wilson’s original animat (1985), the Genesys/Tracker system (Jefferson et al., 1992), Floreano’s (1993a) work on nest-based foraging strategies, the work of Parisi, Nolfi, and Cecconi (1992) and Todd and Miller (1991a,b) on neural network-controlled creatures that evolve learning abilities, and Cecconi and Parisi’s (1993) neural networks with motivational units. In AntFarm (Collins & Jefferson, 1992a), and Koza’s (1992) ant colony simulations, ants within a given colony interact; however, the ants in each colony have a single common genome, and selection occurs between these single genomes based on the fitness of each colony determined in isolation. All of these systems differ from ours in that multiple creatures with different behaviors never interact in a common world, and as a result much of the richness of both social behavior and the effects of whole populations on a shared environment is left out.

The RAM system (Taylor et al., 1989) is notable in having the potential to model arbitrary interactions between multiple animals in a gridworld by representing each part of the system as a distinct computer program. The modeler chooses what features are relevant and hand-codes programs that embody them, as the model itself has almost no predetermined “biological smarts.” For instance, reproduction must be explicitly written into each program. In contrast, all of the behavior of creatures in our SPO system is determined solely by evolution; if a creature never specifically chooses to reproduce in our system (in this case by having too low an evolved probability of randomly selecting the reproduction action), it will die (or live) childless.

Many gridworlds that allow multiple creatures to interact have been constructed with the goal of illustrating certain kinds of interactions between creatures or certain kinds of system dynamics. Ackley and Littman (1992) focus on the relation between learning and evolution. Bedau and Packard (1992) create a simple system with which they study measures of evolutionary activity. Werner and Dyer (1992, 1993) explore systems that illustrate specific kinds of communication and herding activity. Floreano’s (1993b) shared environment research examines the differences in behavior that occur when creatures can or cannot sense each other. This work in particular bears on our current interest in the effects of multiple creatures in a commonly grazed environment, but his methods are quite distinct and his system starts at a higher level of behavioral (and sensory) complexity. Each of these systems comes much closer to the work we describe here, but all differ in being set up to explicate some particular form of behavior.

Our approach instead is to build in the possibility of evolving a wide set of different types of behaviors, and then manipulate general features of environment structure to see what types of behaviors actually do evolve as adaptive.

Packard’s (1989) work on intrinsic evolution shares similar motivations with ours, striving for simplicity in order to bring a variety of basic evolutionary questions to the forefront. Some of the main differences are that his creatures are born a lot smarter (they know how to find food) and have significant sensory input, and only a few types of food distribution and growth are examined. Littman (1992) has approached the problem of characterizing environments in terms of behavioral capabilities by expanding on Wilson’s (1991) scheme for classifying environmental indeterminacy, and he has recently explored half of the problem we pose here, analyzing the goal-seeking abilities of agents with no memory or internal state, only current sensation (Littman, in press). Mason (1993; Erdmann and Mason, 1988) has discussed the other half of the equation, designing state-guided robots that have no sensors. Other more recent projects by Yeager (1994) and Grant (1993) and the LEE system of Menczer and Belew (in press) introduce rich creature simulation environments in which a variety of behaviors can evolve. But each come with predetermined decisions about the cognitive, sensory, and memory abilities of the creatures involved and the structures of the environments in which they live. The SPO system, of course, also makes such choices about the evolvable capabilities of the simulated creatures, but at what we hope is a lower level of complication (if not also sophistication), more appropriate for the kind of study we are undertaking here into the simplest forms of blind adaptive behavior evolved across a range of environments.

3 The World

We begin the exploration of the evolution of sensor-less, state-less behavior in response to environmental structure with the creation of the environments themselves. The world our creatures inhabit is at present a very simple (and commonly used) one, a two-dimensional 64x64 grid of distinct square locations, toroidally connected at opposite edges so that movement wraps around and nobody will fall off. Any number of creatures can occupy any of the 4096 locations in the world. The only other type of object found in the world is food, which can also be present in different locations in varying amounts. (“Food” is equivalent to energy, the common currency in the world, which creatures absorb through eating, use up in all their activities, and die without.) The only direct control we have over their environment is the spatial and temporal distribution of the food they encounter.

There are an endless number of ways to parameterize the possible patterns of food distribution over space and time in our simple world. We have started with just a few parameters that are particularly salient and more or less clearly connected to patterns of plant growth in natural environments (see e.g. Barbour, Burk, & Pitts, 1987; Bell, 1991). First of all, we specify the maximum overall density of food in the world –

that is, the percentage of locations in the world in which food can be present simultaneously. If food is distributed purely randomly in the world with a given density (as it is in many of the experiments described in this paper), then the density also determines how far on average a creature will have to travel from one food-bearing location before encountering another.

The pattern over time with which food regrows determines its temporal distribution. We control this aspect of the environment with two parameters. First, there is a food regrowth delay, that specifies how long a just-grazed location will remain empty (food-less) before food can regrow there. Where the maximum food density specifies how far a creature must go to find food, the regrowth delay specifies how long a creature must wait at a given location before food reappears there. Once food can sprout again in a given empty location, the second parameter — the food regrowth rate — says how much food energy will be added to this location at each time-step, that is, how fast the food-plants (re)grow.

Food does not go on growing indefinitely in a given location in this world – we also specify a maximum amount of energy per location (that is, maximum food-plant “size”). This maximum food amount can be specified directly, but we usually use a more indirect computation of this parameter, instead specifying the desired energy “flux” in the world — that is, the average amount of energy appearing over all locations in the world and across all time-steps. Say for instance that we have two environments in which food regrows instantly (zero regrowth delay) to a level of 2.0 units of energy. In one case, we specify a maximum food density of 100%, so that there are 2.0 units of energy everywhere in the world. In the second case, we specify a maximum food density of 50%, so that half of the locations have 2.0 units of energy, and half have 0.0 units. Thus, in the second environment, we have the same maximum food amount per location, but only half as much energy is ever available for the population of creatures, and this can seriously affect their livelihood. We would like to have two environments with different densities (and other parameters) that can support the *same* numbers of creatures, by having the same *average* amount of energy per location per time-step. If we want 2.0 units of energy per location per time-step in the 50% density environment, we can achieve this by doubling the maximum food amount that can grow at each location.

The same thing applies if we have a food regrowth delay. For instance, in an environment with 100% food density and 1 time-step regrowth delay, food appears only half as often (every second time-step), so to give this environment an equal energy flux, we again have to double the maximum food amount used in the zero delay case. The energy flux can thus be computed as

$$energy_flux = food_max_amount * \frac{food_max_density}{1 + regrowth_delay}$$

Since the energy flux is what is critical for maintaining creature populations in the world, rather than the maximum food amount per location, we can set the energy flux and then

switch around the above equation to compute the maximum food amount from the flux and the density and delay parameters. Again we set the food regrowth rate to be equal to the maximum food amount per location, so that food regrows to its full level as soon as it can.

4 The Creatures

As we described in our original statement of this research program (Todd & Wilson, 1993), we are interested in the evolution of three components of behavior in response to different environment structures, namely action, sensation, and memory. Only the first, action, is actually necessary for behavior, and indeed adaptive behavior. That is, creatures need not be able to sense nor remember anything about their world (or themselves) in order to behave adaptively in it, provided the world is generous and benign enough. Simply performing different actions with certain probabilities can suffice for survival and reproduction. There are probably no creatures alive today that adhere to this strategy – the competitive real world is probably no longer kind enough to allow such blind-faith behavior, if it ever was — but it is possible that such mindless action could play a part in earlier periods of the evolution of life, terrestrial or otherwise. (In any case, the simple creatures we simulate here could be likened to nothing more complicated than a primitive unicellular organism; even plants have more sensory abilities and behavioral options than the individuals in the SPO system!) And since such unguided action is the logically simplest case we can construct, it is sensible to begin here in our exploration of the evolution of adaptive behavior. Once we have established this base-level state, it will also be much easier to show clearly the adaptiveness of evolved sensors and memory systems in comparison.

4.1 Action types

There are only four kinds of actions that we allow these first simple creatures to perform. To absorb the food-energy in their environment upon which they subsist, the creatures can eat; to travel from location to location in search of new food, the creatures can move; to populate the world with more of their own kind, the creatures can split; and to pass time until conditions have changed (or maybe not changed), the creatures can sit and do nothing. We will discuss each of these in greater detail.

The only way that creatures can get the energy they need to live is to absorb it from the environment through eating, and the only option available to them for eating at present is the food-plants growing in the world. Plants can only be consumed whole in these worlds; thus, if there are multiple creatures all trying to eat at one particular food-bearing location, only one of the creatures actually succeeds in eating, and it alone will get all the food-energy in that location, while all the others get no new energy. The winning creature — the one that gets to eat — is chosen randomly from among all the contenders attempting to eat in a location. The winning creature's own

internal energy is incremented by the amount of food-energy in the location, and the location's food-energy is set to zero, ready to sprout again after the specified regrowth delay.

Creatures might not need to move about in their environment in order to survive — if they are lucky enough to be born on a food-growing location, they might be able to get by with just sitting and eating their whole lives — but since movement is the basis of all outwardly observable behavior, without it we would end up with very dull creatures. The creatures in SPO can move one location at a time in one of the four orthogonal directions in the square grid. Movement is relative to the current heading of the creature, so that an individual does not choose to move north, south, east, or west, but rather forward, backward, left, or right. Every move-action resets the current heading to make the creature face the direction it is going to move (e.g. if the creature chooses to move backward, it first turns around and then moves one step *forward* in that new heading). This heading-relative movement has implications for the creatures' behavior, as we will see in section 5.2.

The ultimate arbiter of what behavior patterns have been adaptive in a given environment is those creatures that are still around populating the environment after a suitable length of time. This could be accomplished simply by filling the world with randomly generated creatures and seeing who remains after the selective filtering process of death weeds out the maladaptive ones. But to enable powerful evolutionary search through behavior-space by means of small adjustments to existing behavior patterns, we need replication with modification. In SPO, we allow the creatures to reproduce by asexual splitting, which creates a modified copy of the “parent” individual and divides that individual's current energy equally between it and its new “offspring.” The newborn is placed in the same location as the original splitting creature, and they go on to lead separate lives from that instant. The modifications in the newborn consist of rare changes in the probabilities of choosing the various actions in the creature's behavioral repertoire, as will be described in the next section on action selection. (The mutation rate — the chance that any given action probability will be replaced with a completely new probability value — is 0.05 throughout the runs described here.) This results in an individual who will behave slightly differently from the parent it split off from, allowing evolution to search behavior-space appropriately.

Finally, we allow creatures to do nothing in a particular time-step, just sit and wait for the moment to pass. In this way, creatures can “travel” through time in the same way they can traverse space, and the two are separable (at least insofar as sitting merely covers time, not space — movement at present crosses both time and space simultaneously, since it takes time to move a given distance). We hope that including this ability will allow interesting aspects of the duality of space and time in these environments to manifest themselves in the behaviors of the creatures.

Every action a creature performs has energetic side-effects determined by the bioenergetic characteristics of the world

(see Todd & Wilson, 1993; Menczer and Belew, in press, present another approach). Each type of action uses up a certain amount of energy; varying these action-costs across different worlds can yield very different behavioral repertoires, as we will discuss in section 5.2. Beyond these action-costs, some actions can further increase or decrease an individual's energy. Successful eating attempts increase the individual's energy by the amount of food-energy absorbed from the current location, while splitting cuts the individual's energy in half. Thus energy flows into the simulated world through food growth, continues into creatures when they eat, and flows back out of the creatures by being burned up through the actions they perform. Since creatures currently die only when they run out of energy, no energy-filled carcasses ever end up in the world — dead creatures just disappear.

4.2 Action selection

Since our very simple creatures have neither sensors nor memories or internal states to distinguish any location from any other or any time-step from the rest of eternity, there is not much they can use to decide what actions to perform at a given instant. In fact, all they *can* use is the roll of a die — they select among all the possible actions randomly. But the die can be weighted, with different probabilities (all summing to 1.0) assigned to each of the seven possible actions (move forward, backward, left, or right, eat, split, or sit). It is these probabilities, distinct in each individual creature, that determine the makeup of their behavioral repertoire, and allow them to behave more or less adaptively in response to the different environments they live in.

Each creature therefore has an internal list of seven probabilities for selecting among the seven possible actions. These probabilities remain constant during the creature's lifetime. When a creature splits, it is these probabilities that are modified slightly in the newborn offspring. At every time-step, each creature chooses one action at random according to its own internal probability distribution, and then performs that action. These probability distributions are thus all we can look at to assess the behavioral repertoire of the creatures in our evolving populations, but this situation has the advantage of allowing us a very easy and straightforward way of analyzing what is happening in the world. In contrast, trying to make sense of the behaviors of a population of sensing, remembering, and acting neural network-based simulated creatures is a far more daunting proposition. And as we will see in the next section, even so simple a behavioral repertoire as this set of seven actions weighted relative to each other can still lead to some interesting evolved results.

5 Results

We have currently implemented the SPO system on a 4096-processor 7 MHz Connection Machine 2 (CM2), using the *Lisp data-parallel language. Each processor maintains one location in the world (hence the current 64x64 size), as well

as one or more creatures (which may be located anywhere in the world). The simulation churns along at a rate of about five time-steps per second for a population of up to 16,000 creatures. We have watched the evolution of these simple creatures in a wide variety of environments (in the hundreds of different settings to date) created with the parameters described in section 3. In this section, we first describe a typical run embodying general features, and then discuss some of the trends in the evolved behavior patterns that appear with changing parameter settings and a few of the more peculiar behaviors we have encountered.

Each run is begun the same: first, the world is populated with food as described in section 3. Then, to create the initial population, we generate 500 creatures with random action probability distributions and place them randomly across the world. Each creature is started off with an arbitrary 20.0 units of energy, to give everybody an equal foot in the door for survival. This initial population gives the evolutionary process a variety of starting points to begin working from in parallel, and allows multiple possible behavioral strategies to appear simultaneously and compete with each other for domination. In this way, the final population may come represent the most successful — the most adaptive — of several contending patterns of behavior. Once we first populate the world, the only way new creatures can be introduced after this is through the splitting actions of the creatures themselves.

5.1 Adapting over a typical run

The evolutionary course through which a population's behaviors adapt to the environment contains several common features across different environmental structures.

Figure 1 shows aspects of a typical run of a population of creatures evolving in a simple environment over 20,000 time-steps. In this environment, we have specified a maximum food density of 40%, and a food regrowth delay of 4 time-steps, giving abundant food that takes a bit of time to reappear. The energy flux for the world is 0.6 units per location per time-step. These three parameters combine to give a maximum amount of food per location of 7.5 units of energy (which regrows all at once after the regrowth delay).

The first thing that happens as the initial random creatures begin to act in their environment can be seen in Figure 1a: the population size shoots up from its starting point of 500 to a high of 1645 by time-step 7. This is because the initial population contains a large number of creatures with very high probabilities of splitting, so that's what they do, resulting in a miniature population explosion. But since reproducing is basically *all* that these prolific creatures do, they soon run out of energy and die, and the population crashes over the next 70 time-steps to a low of about 300 individuals. From that point on, though, the successful creatures begin to make their presence felt, as they responsibly move about their world, eat, and only very occasionally have offspring. The population grows steadily over the next 800 time-steps, reaching a high of 2188 at time-step 1949 (Figure 1b), and then gradually

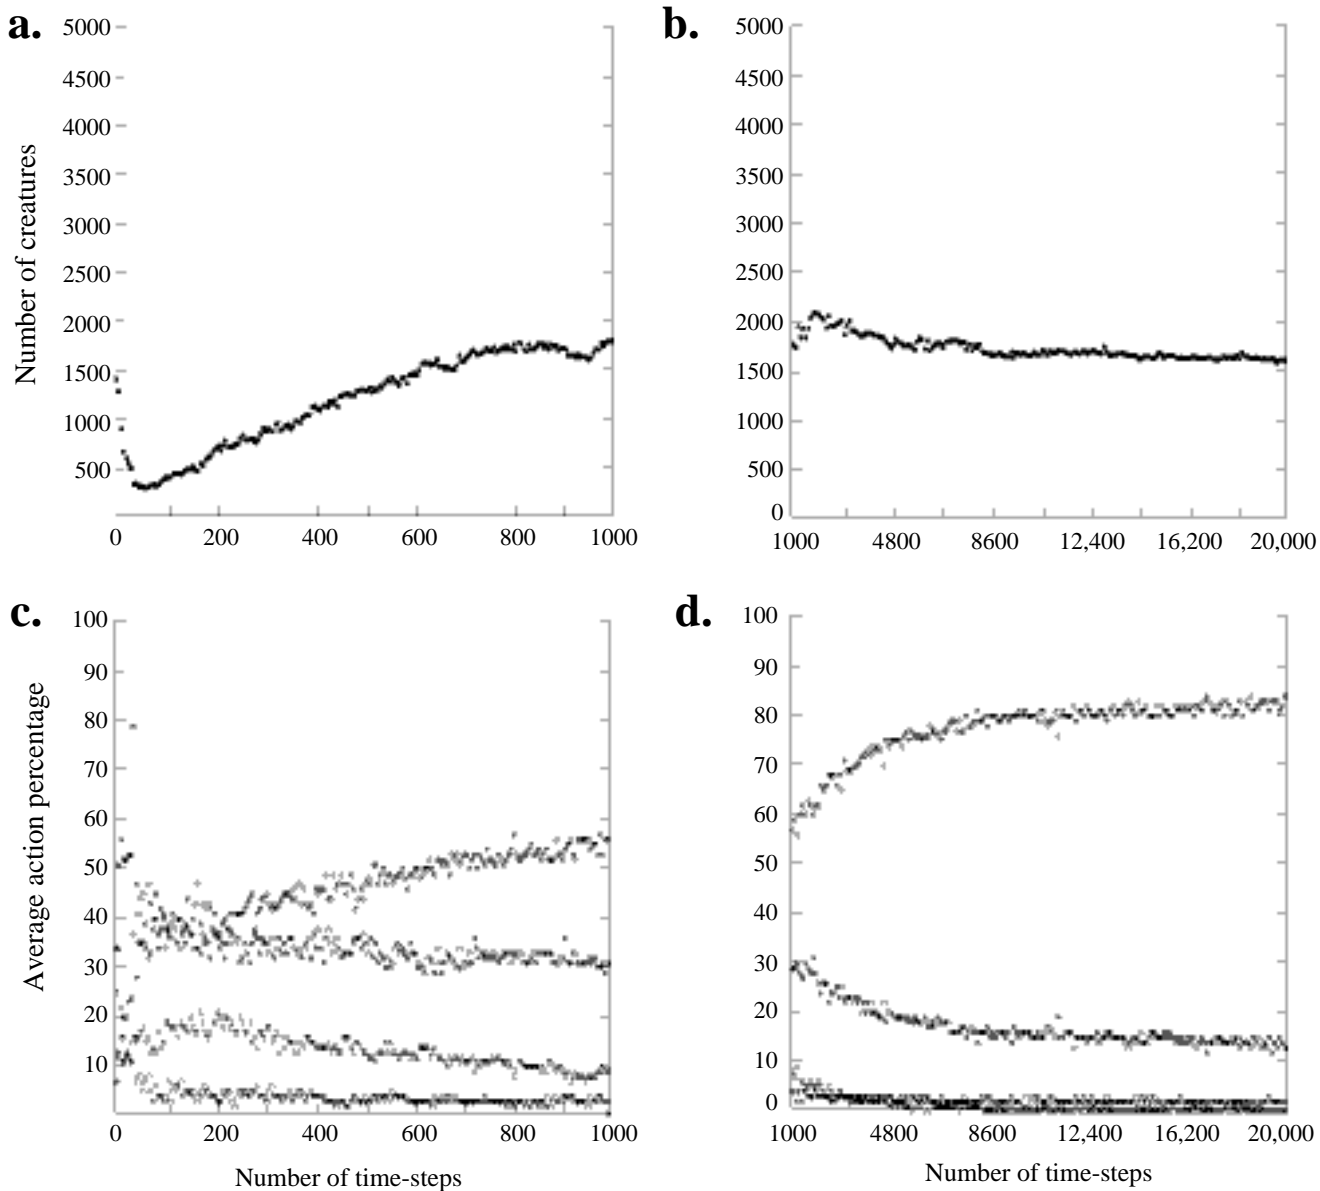


Figure 1: Population features evolving over time in a single typical run. a. Population size from 0-1000 time-steps showing initial population explosion, decline, and slow regrowth. b. Population size from 1000-20,000 time-steps showing eventual stable plateau. c. Proportions of different actions at each time-step averaged across the entire population, from 0-1000 time-steps. From top to bottom (at right edge), “e” represents average percentage of individuals eating, “m” represents combined movement percentage (in all four directions), “s” represents splitting (reproduction), and “n” represents sitting/doing nothing. d. Proportions of different actions, from 1000-20,000 time-steps.

falling to a relatively stable plateau of about 1600 creatures. (If we continue to run this world past 20,000 time-steps, little changes.) This initial population explosion followed by a rapid collapse and slower rebuilding is seen in the majority of runs. The final population size reached, though, depends greatly on the exact structure of the environment.

The reason for this population size growth pattern can be seen in Figures 1c and 1d, showing the average action probability distributions over time. (In each figure, “e” represents percentage of action devoted to eating, “m” is the combined movement percentage, “s” is splitting percentage, and “n” is sitting/do-nothing percentage.) In the initial population, all of the actions are performed about equally, meaning that splitting will make up about 14% (1/7th) of all behavior. Now every time a high-splitting-probability creature splits, it creates an additional new creature that also is likely to have a high splitting probability, so that very rapidly, the probability of splitting increases in the population. By the tenth time-step, 30% of all behavior is splitting. This rampant reproduction cannot be maintained for long — even these simulated creatures must stop to eat — and the splitting rate rapidly falls off as the hyper-splitters fall dead from exhaustion (or at least lack of energy).

The percentage of time spent splitting in fact continues to drop steadily over the course of evolution, eventually reaching a low of 0.17% by time-step 20,000. This is accounted for largely by the fact that relatively few of the creatures in the population ever bother to split at all — the majority of them have a 0% chance of splitting. Indeed, only 175 out of the total 1609 creatures at time-step 20,000 can split. This means that most of the creatures spend their time just eating and occasionally moving, building up greater and greater amounts of energy, reaching an average energy per creature of over 7800 units by the end of this run. This effectively makes these non-splitting creatures immortal, a situation that can lead to such dire consequences as the eventual end of all evolution (as described in Todd, 1993).

In place of the disappearing likelihood of splitting, though, another action must increase in likelihood (since the creatures always select *something* to do at each time-step, even if it is just sitting still). Eating takes up the slack from the decreasing splitting, growing constantly after its initial low probability. In fact, the percentage of time creatures spend eating on average rises at the expense of *all* of the other actions. Eating beats out the combined movement percentage (the sum of movements in each of the four directions) to become the most popular action shortly after 100 time-steps. This all-important activity continues to rise until it plateaus around 83% by time-step 16,000.

The combined movement probability begins very high (since it subsumes four different actions, or 4/7th of the total initial probability distribution), but it quickly falls off. It remains a useful action in this environment, though, leveling at around 14%. (In other environments, as we will see, movement is a much less adaptive behavior, and ends up at much

lower levels.) The amount of time spent moving in each of the four directions is not evenly distributed, however, as we will see in the next section — interesting patterns emerge. Finally, the sit (do nothing) action proves relatively pointless in this environment, rapidly falling to about 2.5% usage. This is still much higher than the final splitting percentage, though, indicating that just sitting still is not as risky a proposition — and thus not as strongly selected down — as is losing half of one’s energy to a new (and competing!) offspring.

These patterns of evolution of the average behavioral repertoire are fairly representative of the range of runs we have performed. (Since the action probability distribution seems to have become more or less stable by 15,000 time-steps, we use this as the length of the runs in the remainder of this paper.) Using averages, though, can hide multiple strategies existing within a single population, a concern to which we will return in the next section; but often the populations that evolve are fairly homogeneous, with one particular behavioral pattern winning out over time.

Other things are afoot in the world besides behaviors changing, however. As the population of creatures adapts to the environment and comes to contain more individuals with behavioral repertoires better tuned to the world they encounter, the world itself changes as well, owing solely to the actions of the creatures. This can be seen in Figure 2, which shows the density of food in the world over time. Two lines are plotted in both panes (Figures 2a and 2b); the bottom one (“*”s) shows the percent of food-bearing locations after the population has grazed, while the top one (“o”s) shows the percent of food-bearing locations after food has regrown in the world. Remember that the maximum food density in this particular environment is 40% — but as can be seen, this value is never actually reached, even following regrowth, after the initial planting of food. The reason for this is the long regrowth delay (4 time-steps), which keeps much of the world infertile for much of the time, so that the maximum food density never can actually regrow. As soon as some of the locations regrow, the creatures graze them down again, keeping the entire percentage low. This in turn has an impact on the future evolution of the creatures themselves, limiting their numbers and probably making movement in search of the sparser food locations more adaptive.

This is a very different situation from simulation systems that run single creatures in their own worlds to test out their adaptiveness. In those cases, each creature experiences the same type of environment over the entire course of evolution, and they become better and better adapted (one hopes) to that fixed environment over evolutionary time. In our situation, though, the environment we initially specify is quite different from the picked-over world the creatures end up making for themselves — they must evolve to a moving target, because the world itself changes over evolutionary time as well. This complicates our task in analyzing evolution in response to environments, but it also makes the model more realistic.

But the creatures never eat *all* of the food at a given instant;

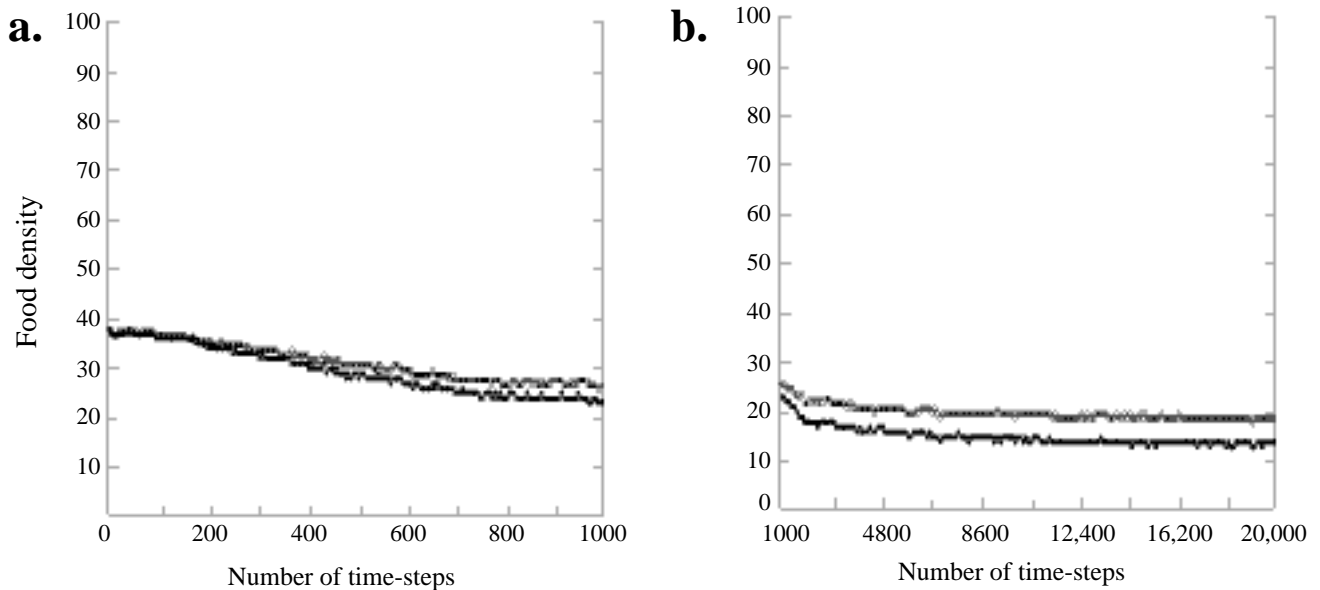


Figure 2: Environmental changes occurring over time in a single typical run. a. Food density in percent of locations containing food-plants at each time-step from 0-1000; the upper line (“o”s) shows density after food regrowth, while the lower line (“*”s) shows density before regrowth. b. Food density from 1000-20,000 time-steps showing eventual leveling off and large difference before and after regrowth.

as can be seen in Figure 2b, even after 20,000 time-steps the grazing creatures leave behind food in about 14% of the locations (out of about 20% of locations that regrow each time-step). This is because there are too few creatures in the final population to graze everywhere. The 1600 creatures in the population at 20,000 time-steps in fact end up only occupying 29% of the locations in the world. The creatures are relatively well dispersed, though, as there are only an average of 1.36 creatures per occupied location, compared to the stack of 2.3 per location that occurs at time-step 10 (right after the initial population explosion, when many creatures and their offspring are sitting on the same spot).

5.2 Adapted behavior patterns and trends

The typical run just discussed shows that the behavioral repertoires of creatures are changed over time through the simulated evolutionary process – but are the creatures actually adapted to their particular environment? Or are we just witnessing the unwinding of some general dynamic process independent of the pattern of food distribution in the world? The most direct way to answer this question is to vary the structure of the environment, and see whether or not the evolved behaviors also change as a result. We have varied the maximum food density, the food regrowth delay, the average energy flux per location per time-step, and the cost of movement relative to the costs of the other actions, and have found that the be-

havioral strategies evolved do in fact change significantly in response. This indicates that the challenges of different types of environments are best (most adaptively) met by different probability distributions of blind actions. In this section we briefly discuss some of the major trends we see as we vary the environmental parameters, and a few of the adaptational oddities that pop up.

The most significant trends we find as the environments change are in the big-ticket actions, the ones that creatures spend most of their time doing: moving and eating. Holding other parameters constant, the rate of movement in general increases both as the density of food increases, and (independently) as the food regrowth delay increases. The former result is a bit counter-intuitive: why should creatures move *more* when food is denser, that is, closer together? In such environments, it should take *less* movement to find the next fertile food spot. But this is probably just why more movement *is* possible in such environments — there is far less risk to it, because food is likely to be encountered soon. In low-density environments, moving randomly in the way these creatures are restricted to may take them a very long time to stumble from a location where there is certain to be more food growing — if they just wait long enough — to another similarly good oasis, and the possibility of starving on the way is great. In high-density environments, this risk is much reduced, and creatures can afford to allot more of their time to movement from one food-bearing location to another nearby.

As regrowth delay increases from 0 (instant regrowth on the next time-step), again so does the percentage of time spent moving. If food is certain to reappear immediately at your current location, there is little incentive to move from there in hopes of finding another place to eat; but if food won't grow back for some time, wandering away in search of more might be reasonable. This seems to be what the creatures do. But still, why should they give up a good thing if they begin at a food-growing location? If we look carefully at the distribution of individual movement types the creatures are using, it appears that they do *not* simply abandon the location where they begin – instead, the loop back to it after some delay! Analyzing the possibilities, it is clear that only movement solely in direction 0 — straight ahead — takes creatures ever further away from their starting point. Repeated movements in direction 1 — turn right and move forward — bring a creature back to where it started in four moves, as does movement in direction 3 (looping leftwards instead), while movements in direction 2 — turn around and move forward — bring it back in two moves. So if food regrows after a short delay, moving in short two-step loops with direction 2, interspersed with eating, may match the temporal distribution of food most appropriately. If food takes quite a bit longer to regrow, moving in longer four-step loops could be more adaptive.

Remarkably, this pattern is indeed borne out in the data: for instance, in the run described in the previous section, with regrowth delay of 4 time-steps, at the end of 15,000 time-steps of evolution we get 1% of the population only moving in direction 0 (straight ahead), 29% of the population only moving in direction 2 (two-step loops), and 21% of the population either moving only in direction 1 or only in direction 3 (four-step loops) — the rest of the population (49%) has some mixed movement strategy (or no movement). When we look at regrowth delay of 2, the emphasis on two-step loops is even greater, with 32% of the population moving only in direction 2, and just 14% moving in only direction 1 or 3. In contrast, when we change the regrowth delay to 12, these values become roughly inverted: we get 7% movement in direction 0 only, 17% in direction 2 only, and 40% in direction 1 or 3 only, showing that four-step loops have now become much more prevalent, and are therefore probably more adaptive and more often selected for. (Movement only in direction 0 increases as we continue to raise the regrowth delay, showing that the creatures eventually start to give up on the looping strategy when the delays between food reappearances are too large.) Thus, even these extraordinarily simple creatures are able to invent a behaviorally-based clock to adaptively time their behaviors in response to the structure of the environment, something we were quite surprised to find. (Of course, these timers are *not that* useful, in some sense, because the creatures don't really know when to start them, nor when to stop them, nor even how to keep them running — everything is done probabilistically, and a rogue “sit” action coming in the middle of a four-step loop could throw off the timing of the

entire sequence, for instance. Only through the use of internal states can such timers be made more accurate.)

The percentage of time spent trying to eat basically follows trends opposite to those for movement, falling with both increasing food density (because it is less necessary) and increasing regrowth delays (because it is useful less often). Splitting and doing nothing both remain quite low-percentage actions in most worlds, splitting because it is a very costly act, and sitting because it does not offer much positive return (why sit when one can try to eat instead, since both have the same energy cost in this case?). All of these trends are modified somewhat by the cost of movement relative to the costs of the other actions, but basically they are just scaled up or down in magnitude, and the direction of change remains the same. For instance, when we run another world like the one described in the previous section, but this time movement costs only 0.1 unit of energy instead of 0.5 (i.e., it costs one-fifth of what all the other actions cost), after 15,000 time-steps creatures spend about 27% of the time moving and 72% of the time eating (compared to 14% and 83% respectively above). When movement costs 1.0 unit of energy (twice the other action costs), these values change to 8% moving and 88% eating, reflecting the fact that it's cheaper to attempt to eat than to move. Finally, in an extreme case, when movement costs *nothing* (0.0 unit of energy), creatures switch to spending 74% of their time skidding joyfully around their frictionless air-hockey world, and only pausing to try to grab a mouthful of food 25% of the time.

Despite these general trends, there are still occasions at particular parameter settings when very curious behavioral strategies appear and take over the population. For instance, with movement cost of 1.0, food density of 10%, and regrowth delay of 2 time-steps, we find great variation in the kinds of strategies that different populations (on different runs of the system) converge upon, indicating different “species” of creatures. One particularly odd species spends 32% of its time just sitting in place (with 18% movement and 50% eating), far higher than any other population ever encountered. In another case, with food density 5% and regrowth delay of 1 time-step, one species spends 38% of its time splitting, again an unheard-of figure. The reasons for these strange strategies are being investigated, as well as their prevalence at other parameter settings. It is possible that they are very rare but widespread, and only occasionally take over an entire population (indeed, we find considerable variation in many evolved populations — convergence is often not attained so long as splitting continues). The extent of speciation in general, and how much within-population variation is hidden by our use of behavioral repertoire averaging, remains to be explored. But the fact that multiple behavioral strategies can be found for dealing with certain kinds of environments, and the peculiarity of some of these solutions, further indicates the extent to which the behaviors of individual creatures have been adapted to the particular situations they evolve in.

6 Conclusions and Further Directions

Despite the simplicity of our simulation model for the evolution of behavior in blind, amnesic creatures, we have found a surprising amount of richness in the results. The behaviors of creatures are adapted by evolution over time to the particular spatiotemporal structure of food growth in their environment. Even in creatures with nothing else to do but sit, eat, move about, and reproduce in the dark, unexpectedly sophisticated strategies emerged, like the use of movement-based time-keeping to deal with temporal delays in the world. Sensors, internal states, and memory systems are thus clearly unnecessary for at least the grossest level of adapted behavior.

But such enhanced information-gathering abilities will obviously help, and that is the direction we are heading in with the SPO research program. The next step is to add food-detecting sensors which can focus at any relative position in the world, to investigate the particular deployment of sensors that prove adaptive in different types of environments. Furthermore, in keeping with the philosophy that every action must be chosen by the creature (and not by some external system bias), we want to let sensation itself be a selectable action, so that creatures will “have their eyes shut” unless they specify otherwise, rather than always having a constant stream of input data impinging on them from ever-active sensors. Monitoring the pattern of “eye-opening” will also allow us to get a better handle on just when creatures need to use sensory input, in a way that is impossible with constantly-sensing creature models.

After this, we will add internal state (memory) systems that allow creatures to “look back in time” at the state of a particular world location in the past (e.g. whether or not it held food four time-steps ago), in a manner analogous to the way their sensors look “across” time to locations that will take several time-steps into the future to reach. In this way, we hope to explore both the temporal and spatial aspects of environmental variation and behavioral responses to them more accurately. But before these enhancements, we will continue to establish the baseline of adaptive behavior capable in sensor-less, memory-less creatures of the kind described here, showing that adaptive behavior *is* possible without looking.

References

- Ackley, D., and Littman, M. (1992). Interactions between learning and evolution. In C.G. Langton, C. Taylor, J.D. Farmer, and S. Rasmussen (Eds.), *Artificial life II* (pp. 487-509). Reading, MA: Addison-Wesley.
- Barbour, M.G., Burk, J.H., and Pitts, W.D. (1987). *Terrestrial plant ecology*, 2nd ed. Menlo Park, CA: Benjamin/Cummings.
- Bedau, M.A., and Packard, N.H. (1992). Measurement of evolutionary activity, teleology, and life. In C.G. Langton, C. Taylor, J.D. Farmer, and S. Rasmussen (Eds.), *Artificial life II* (pp. 431-461). Reading, MA: Addison-Wesley.
- Bell, W.J. (1991). *Searching behaviour: The behavioural ecology of finding resources*. London: Chapman and Hall.
- Cecconi, F., and Parisi, D. (1993). Neural networks with motivational units. In J.-A. Meyer, H.L. Roitblat, and S.W. Wilson (Eds.), *From animals to animats 2: Proceedings of the Second International Conference on Simulation of Adaptive Behavior* (pp. 346-355). Cambridge, MA: MIT Press/Bradford Books.
- Collins, R.J., and Jefferson, D.R. (1992). AntFarm: Towards simulated evolution. In C.G. Langton, C. Taylor, J.D. Farmer, and S. Rasmussen (Eds.), *Artificial life II* (pp. 579-601). Reading, MA: Addison-Wesley.
- Floreano, D. (1993a). Emergence of nest-based foraging strategies in ecosystems of neural networks. In J.-A. Meyer, H.L. Roitblat, and S.W. Wilson (Eds.), *From animals to animats 2: Proceedings of the Second International Conference on Simulation of Adaptive Behavior* (pp. 410-416). Cambridge, MA: MIT Press/Bradford Books.
- Floreano, D. (1993b). Patterns of interactions in shared environments. In *Pre-proceedings of the Second European Conference on Artificial Life*, Brussels, Belgium, May 1993.
- Grant, W. (1993). *Artificial life worlds as discovery environments for learning*. Talk presented at the Artificial Life III Workshop, Santa Fe, NM.
- Jefferson, D., Collins, R., Cooper, C., Dyer, M., Flowers, M., Korf, R., Taylor, C., and Wang, A. (1992). Evolution as a theme in artificial life: The Genesys/Tracker system. In C.G. Langton, C. Taylor, J.D. Farmer, and S. Rasmussen (Eds.), *Artificial life II* (pp. 549-578). Reading, MA: Addison-Wesley.
- Koza, J.R. (1992). Genetic evolution and co-evolution of computer programs. In C.G. Langton, C. Taylor, J.D. Farmer, and S. Rasmussen (Eds.), *Artificial life II* (pp. 603-629). Reading, MA: Addison-Wesley.
- Littman, M.L. (1992). An optimization-based categorization of reinforcement learning environments. In J.-A. Meyer, H.L. Roitblat, and S.W. Wilson (Eds.), *From animals to animats 2: Proceedings of the Second International Conference on Simulation of Adaptive Behavior* (pp.). Cambridge, MA: MIT Press/Bradford Books.
- Littman, M.L. (in press). Memoryless policies: Theoretical limitations and practical results. In J.-A. Meyer, P. Husbands, and S.W. Wilson (Eds.), *From animals to animats 3: Proceedings of the Third International Conference on Simulation of Adaptive Behavior*. Cambridge, MA: MIT Press/Bradford Books.
- Mason, M.T. (1993). Kicking the sensing habit. *AI Magazine*, 1, 58-59.
- Erdmann, M.E., and Mason, M.T. (1988). An exploration of sensorless manipulation. *IEEE Journal of Robotics and Automation*, 4, 369-379.

- Menczer, F., and Belew, R.K. (in press). Latent energy environments. In R.K. Belew and M. Mitchell (Eds.), *Plastic individuals in evolving populations*. Reading, MA: Addison-Wesley.
- Meyer, J.-A., and Guillot, A. (1991). Simulation of adaptive behavior in animats: Review and prospect. In J.-A. Meyer and S.W. Wilson (Eds.), *From animals to animats: Proceedings of the First International Conference on Simulation of Adaptive Behavior* (pp. 2-14). Cambridge, MA: MIT Press/Bradford Books.
- Packard, N.H. (1989). Intrinsic adaptation in a simple model for evolution. In C.G. Langton (Ed.), *Artificial life* (pp. 141-155). Redwood City, CA: Addison-Wesley.
- Parisi, D., Nolfi, S., and Cecconi, F. (1992). Learning, behavior, and evolution. In F.J. Varela and P. Bourguine (Eds.), *Towards a practice of autonomous systems: Proceedings of the First European Conference on Artificial Life*. (pp. 207-216). Cambridge, MA: MIT Press/Bradford Books.
- Taylor, C.E., Jefferson, D.R., Turner, S.R., and Goldman, S.R. (1989). RAM: Artificial life for the exploration of complex biological systems. In C.G. Langton (Ed.), *Artificial life* (pp. 275-295). Redwood City, CA: Addison-Wesley.
- Todd, P.M. (1993). Artificial death. In *Pre-proceedings of the Second European Conference on Artificial Life*, Brussels, Belgium, May 1993.
- Todd, P.M., and Miller, G.F. (1991a). Exploring adaptive agency II: Simulating the evolution of associative learning. In J.-A. Meyer and S.W. Wilson (Eds.), *From animals to animats: Proceedings of the First International Conference on Simulation of Adaptive Behavior* (pp. 306-315). Cambridge, MA: MIT Press/Bradford Books.
- Todd, P.M., and Miller, G.F. (1991b). Exploring adaptive agency III: Simulating the evolution of habituation and sensitization. In H.-P. Schwefel and R. Maenner (Eds.), *Proceedings of the First International Conference on Parallel Problem Solving from Nature* (pp. 307-313). Berlin: Springer-Verlag.
- Todd, P.M., and Wilson, S.W. (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: Proceedings of the Second International Conference on Simulation of Adaptive Behavior* (pp. 11-20). Cambridge, MA: MIT Press/Bradford Books.
- Werner, G.M., and Dyer, M.G. (1992). Evolution of communication in artificial organisms. In C.G. Langton, C. Taylor, J.D. Farmer, and S. Rasmussen (Eds.), *Artificial life II* (pp. 659-687). Reading, MA: Addison-Wesley.
- Werner, G.M., and Dyer, M.G. (1993). Evolution of herding behavior in artificial animals. In J.-A. Meyer, H.L. Roitblat, and S.W. Wilson (Eds.), *From animals to animats 2: Proceedings of the Second International Conference on Simulation of Adaptive Behavior* (pp. 393-399). Cambridge, MA: MIT Press/Bradford Books.
- Wilson, S.W. (1985). Knowledge growth in an artificial animal. In J.J. Grefenstette (Ed.), *Proceedings of an International Conference on Genetic Algorithms and their Applications* (pp. 16-23). Pittsburgh, PA: Carnegie-Mellon University.
- Wilson, S.W. (1991). The animat path to AI. In J.-A. Meyer and S.W. Wilson (Eds.), *From animals to animats: Proceedings of the First International Conference on Simulation of Adaptive Behavior* (pp. 15-21). Cambridge, MA: MIT Press/Bradford Books.
- Yeager, L. (1994). Computational genetics, physiology, metabolism, neural systems, learning, vision, and behavior or PolyWorld: Life in a new context. In C.G. Langton (Ed.), *Artificial life III* (pp. 263-298). Reading, MA: Addison-Wesley.