

Environmental Effects on Minimal Behaviors in the Minimat World

Peter M. Todd
Center for Adaptive Behavior and Cognition
Max Planck Institute for Psychological Research
Leopoldstrasse 24, 80802 Munich, GERMANY
ptodd@mpipf-muenchen.mpg.de

Holly A. Yanco
MIT Artificial Intelligence Laboratory
545 Technology Square, Room 705
Cambridge, MA 02139, USA
holly@ai.mit.edu

Abstract

The structure of an environment affects the behaviors of the organisms that have evolved in it. How is that structure to be described, and how can its behavioral consequences be explained and predicted? We aim to establish initial answers to these questions by simulating the evolution of very simple organisms in simple environments with different structures. Our artificial creatures, called “minimats,” have neither sensors nor memory and behave solely by picking amongst the actions of moving, eating, reproducing, and sitting, according to an inherited probability distribution. Our simulated environments contain only food (and multiple minimats) and are structured in terms of their spatial and temporal food density and the patchiness with which the food appears. Changes in these environmental parameters affect the evolved behaviors of minimats in different ways, and all three parameters are of importance in describing the minimat world. One of the most useful behavioral strategies that evolves is “looping” movement, which allows minimats—despite their lack of internal state—to match their behavior to the temporal (and spatial) structure of their environment. Ultimately we find that minimats construct their own environments through their individual behaviors, making the study of the impact of global environment structure on individual behavior much more complex.

Key Words: environment structure, evolved behavior, resource density, resource regrowth, resource patches, simulation

1 Introduction

What matters most in life? That is, what aspects of an organism’s external world make the most difference to the success or failure of its life? This is the basic question underlying the study of the effect of environment structure on behavior. Some ways of organizing or structuring the world make a difference in an organism’s life, and some do not. We seek

to discover which kinds of environment structure *do* make a difference, how much, and why.

Which environmental characteristics effect life’s success or failure clearly depends on how such success or failure is defined. For humans, this is complicated by the fact that we can decide on our own goals, our own criteria of life-fulfillment, and thereby change what we think matters most: For epicureans (if not Epicurus himself), the enjoyment of sensuous gustatory delights is the goal of life, and so good food and drink are what matter most; for monks, the ultimate goal is abandonment of earthly distractions for spiritual insight, so stationary, celibate solitude is best; while for the average “thirtysomething” American, raising a family defines a successful life, and having enough resources in a child-friendly place to live are most important.

For the majority of other animal species, however, a successful life is less a matter of individual decision or contemplation, and more a matter of survival and reproduction. It is then evolution’s task to “figure out” what matters most for these two goals—which aspects of the environment can most affect an organism’s success in attaining resources and spreading its genes. Environmental characteristics that matter most will have the most effect on the evolved (and any learned) behavior of the organism. That is, if an important aspect of the environment had been different during the evolution of the organism’s ancestors, its current evolved behavior would be different as well. But this is not to say that changing such a characteristic of the organism’s *current* environment will have a similar effect on its current behavior—it may not have evolved sensitivity to this characteristic, but rather only appropriate, unchangeable behavioral adaptations. For instance, if our three-dimensional world were suddenly made five-dimensional, we might not have any behavioral response to this change, because we did not evolve sensitivity to such changes—even though such an environmental characteristic is undoubtedly central to our behavior (Shepard, 1987).

We are interested here in which environmental structures affect the evolved behaviors of organisms, and how. To study these questions, we will construct simulations of simple organisms behaving and evolving in simple environments and

test how different environments influence the behaviors that ultimately evolve. We start at the beginning, with arguably the most basic environment structure parameters possible: resource density in space and time. And we start with the most primitive instantiation of organisms that we can imagine: creatures with no sensors, and no internal state or memory of past events, but only the ability to emit actions according to an evolved probability distribution. As a consequence, the distinction mentioned in the previous paragraph between changes in the evolutionary environment and changes in the current environment does not affect our research here: Because these simulated organisms cannot sense their current environment, no changes in it could influence their behavior, and only their evolutionary environment can have an effect. Moreover, the environments we use are stable and stationary, in the sense that we do not change their structural parameters over time—only the actions of the simulated creatures themselves can change the current state of their world.

Simulated organisms of the type we use here are often dubbed “animats,” after Wilson’s (1985, 1991) first use of the term to describe his artificial animals that use local sensors to search for food in a grid-based environment. Since Wilson’s work, the term’s meaning has been extended through its use by others following in this research tradition. As a consequence, “animat” now conjures up notions (to those for whom it conjures up anything) of simulated creatures that use simple sensors and cognitive architectures, often with restricted learning and memory abilities, to behave adaptively and take advantage of their surroundings. But the artificial organisms we study here, having neither sensors nor memory of any sort, are simpler than even the most mentally challenged animat. To emphasize this distinction, and not to bring too much disrepute on other more respectable animats by association, we name our simulated creatures here *minimats*.

Our goal with such humble beginnings is to establish a baseline of what effect environmental structure has on even the simplest possible behaving organisms that have no direct contact with that environment. If we find, for instance, that food density has a particular effect on even their behavior, we can suppose that food density will have an effect on the behavior of any more realistically complex, sensor-bearing organisms as well. By starting with the simplest possible forms of blind, stateless behavior in the simplest possible environments, we also hope to make clearer the impact of richer environmental structure on more sophisticated sensor- and memory-guided behavior in comparison.

But by focusing initially on simulated organisms without sensors or memory, we also intend to emphasize the primacy of action in behavior. Rather than viewing organisms as input-output machines that operate on information from the environment to produce optimal adaptive output behavior, we should take the lesson from minimats that acting appropriately in the environment is the main (evolutionary) objective, and sensory and memory systems only evolve to help guide that action. This puts our approach in sympathy with the dynamical systems view of cognition and behavior (see, e.g., Yamauchi & Beer, 1994; Cliff, Harvey, & Husbands, 1993; Hendriks-Jansen, 1996), even though minimat “brains” are hardly dynamical systems.

1.1 Our approach in this paper—what we will do

To establish which environmental structure parameters make a difference to the evolved behaviors of minimats, we use our SPO system (Simulation of Primitive Organisms) to simulate the evolution of populations of minimats in a variety of different environments. The behaviors of minimats in the population change over time in response to their environment—but only *between* individual lifetimes, rather than within them. Minimats lack the adaptive processes of development, learning, and sensor-guided action, all of which bring about changes in the ongoing behavioral responses of an individual. Rather, they are subject only to the adapting process of evolution, which brings about changes in the aggregate behavioral responses of the minimat population. Thus, minimat behaviors are *adapted* to their ancestral environment, rather than *adaptive* to changes in their current environment on a moment-by-moment basis.

The individual minimats 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. Minimats who find and eat the food-energy in their world can survive and produce offspring with slightly modified behaviors, 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 minimat behavioral repertoires. After this adapting process has run long enough, and minimats become well adapted to their environment (or at least as adapted as they can be, given their abilities), we analyze their behaviors and look for patterns in the fit to the prevalent environmental structure.

We construct a variety of very simple environments for the minimats to evolve in, 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. 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 (at least on the minimat’s behalf). Separating aspects of environmental structure into those that matter for minimats (e.g., the distribution of food in the world) from those that do not (e.g., whether food occurs in the “north” half of the environment or the “south” half) 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.

We begin by looking at how the spatial density of food in the environment—how common food is in a given area—influences evolved minimat behavior, expecting this most profound aspect of environment structure to have a large impact. Once this relationship is established, we proceed to look at the effect of temporal density—how often food appears—on evolved behavior. If temporal density also has an effect, there are two

possibilities for how it might influence behavioral evolution: First, temporal density could cause a change in spatial density, which then affects evolution via the first relationship we found; or second, temporal density can have an effect that is independent of any simultaneous change in spatial density. It is important to establish which of these is the case, because if the former condition holds, then we can ignore temporal density as an environmental parameter of importance to minimat behavior and just analyze environments in terms of spatial density (and possibly other parameters). If we find that spatial density has an independent effect, then we must describe any minimat environment with at least these two structure parameters together, and we have reason to suspect that they will both influence the behavior of more cognitively rich animals and organisms.

We continue with this logic for testing the independence of environment structure parameters as we explore the effects of higher-order structure. In particular, we vary another set of minimat environments according to the patchiness of food present—whether food-bearing locations occur independently or together in clusters. We then test whether this affects minimat behavior evolution at all, and if so, whether the effect can be explained through a corresponding change in spatial density. (The temporal density does not change.) As soon as we find a level of environmental organization or structure that our evolved minimats no longer respond to uniquely (differently from their response to lower-order structure), we can stop our search for environment structure parameters that matter to minimats, and we only need to describe minimat environments with the previously established set of parameters.

1.2 Possible differences in evolved behavior – what we will look for

In an early and elegant precursor to our work here, Herbert Simon (1956) analyzed the nature of “satisficing” behavior of a simple food-seeking organism in a minimally structured environment. He stressed the importance of knowing the needs and goals of an organism in order to analyze the relationship between its behavior and its environment:

We are not interested in describing some physically objective world in its totality, but only those aspects of the totality that have relevance as the “life space” of the organism considered. Hence, what we call the “environment” will depend upon the “needs,” “drives,” or “goals” of the organism, and upon its perceptual apparatus. (p. 130)

We take the same approach here¹. Earlier we spoke of the “goals” toward which animals strive as survival and reproduction. For minimats, the situation is largely the same. Because we assess their evolved behavior after a certain fixed amount of time has passed in their world, the “goal” of a minimat—or a minimat behavioral repertoire—is to (still) be around at that final time, when the whistle (or trumpet) blows and the scores are tallied. To accomplish this, a minimat can either be long-lived, conserving its energy and avoiding death so that it will itself be included in the final reckoning; or it can

be prolific, producing many prolific offspring, so that at least its descendants will be present at the final time step. Thus, for minimats survival and reproduction are also the keys to a successful life, though for different “species” of minimats the two may be weighted differently (as they are indeed for animals as well—compare the long-lived seldom-reproducing primates with the short-lived many-egged mayfly).

But in any given environment, organisms can approach the problems of survival and reproduction in a variety of ways. Herbivores, carnivores, parasites, and plants each employ very different strategies for gathering energy and creating young. Evolution not only shapes the behaviors of the species adopting each of these (and other) lifestyles, but it also balances the relative presence of each lifestyle in a given environment. Thus, when we analyze the evolved behaviors of minimats, we will also look for these two evolutionary effects: how behaviors in a particular strategy are adapted to a particular environment, and how the mix of different behavioral strategies also varies from one environment to the next.

Of course, telling whether two behavioral strategies are qualitatively different or just modified versions of the same strategy might not be a straightforward proposition. We have to consider the full range of behavioral possibilities from the organism’s perspective and see if they fall into easily distinguishable categories. (See Blythe, Miller, & Todd, 1996, for a description of the basic-level behavioral categories we can generally expect to evolve, and some of the cues that can be used to distinguish them.) Luckily, this is easily done in the case of minimats: Given the small set of possible actions (and their limited perspective), we can distinguish a few obvious strategies. Epicurean minimats can adopt a “move in search of new food and eat” strategy to secure their next meal, possibly reproducing along the way. Monkish minimats can sit immobile in one location and never reproduce. And thirtysomething, family-oriented minimats can also settle immobile in one location, but reproduce in a sensible fashion, creating offspring who may stay at home or who may themselves move away in search of other rewards.

These are the different behavioral strategies (whimsically described) that we will watch for in our evolved minimat populations. How they mix together—or perhaps form homogeneous populations of a single behavioral strategy—will be one important aspect to consider: Do we get ecosystems with different “species” of minimats? We will also look at the specific adapted rates of the seven different actions within any given behavioral strategy, to see how these too are shaped in response to environment structure. One of our most surprising findings in this regard is the curious behavior of “looping”—moving repeatedly along a small path that loops back to where it began—which allows minimats to “time” their environment in a way that is well-matched to its temporal structure.

In the next section, we indicate our intellectual backdrop in terms of related simulations of simple organisms in grid-based environments. We present the environment structure parameters used in this research in section 3, and the descriptions of the minimats and their behavioral capabilities in section 4. Section 5 gives the methods used to collect the data we report in section 6. In section 6, we follow the general course of investigating the behavioral impact of first food location density, next regrowth delay, and finally environmental

¹Note that we make the analysis situation that Simon describes simpler by eliminating the need to consider perceptual apparatus in minimats.

patchiness, seeing which of these structural parameters make a difference in the life of minimats. But along the way, we take detours to explore mysteries that arise, such as population extinctions, the usefulness of moving in general, and of different types of looping behavior in particular, and what the environment “looks like” to different sorts of minimats. We have highlighted the main results in the subsection titles to make the way easier to follow.

2 Past Work

The work reported in this article fits into the general research program of characterizing the important features of environment structure in terms of the adapted (and often –but not here–adaptive) behaviors they elicit, as outlined in Todd and Wilson (1993). In our previous article (Todd, Wilson, Somayaji, & Yanco, 1994), we showed that the behavior of blind, stateless simulated organisms (also minimats, though we did not yet call them that) with no sensory contact with their current environment could still be adapted to different environments through the course of evolution. This time around, we give the details of what those adaptations actually are. But first, we want to indicate how our efforts compare with those of other researchers in this area.

Several other research projects have used methods similar to our SPO system, simulating simple creatures behaving in grid-based environments. These gridworlds are typically composed of a grid of discrete locations (rather than continuous-space models) that can contain various objects (such as food, water, and other resources, and obstacles, predators, and other hazards) in addition to the simulated creatures. Some of these other projects are similar enough to SPO to allow comparison of results across studies. But most are designed with specific behaviors in mind and have not been used to explore the effects of different environment structures on those behaviors. 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) and its close descendants (including Cliff & Bullock, 1993, Wilson, 1994, and Cliff & Ross, 1994). All of these animats start off with greater cognitive complexity than our minimats, however, in terms of their sensors, memory, or learning ability, and the environments they live in are not varied in a systematic way to study their behavioral effects.

Other related approaches can be found in 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. These systems, though, like the other animat systems, differ from ours in that multiple creatures with different behaviors never interact in a common world, and as a result the possibility of social behavior and

the effects of whole populations on a shared environment are left out. In AntFarm (Collins & Jefferson, 1992), and Koza’s (1992) ant colony simulations, ants within a given colony do 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.

The RAM system (Taylor et al., 1989) is notable in having the potential for modeling arbitrary interactions between multiple animals and a gridworld environment by representing each part of the system as an individual computer program. However, reproduction must be explicitly coded into each program, as the model itself has almost no predetermined “biological smarts.” The modeler must choose what features are relevant and hand-code programs that capture them. (Resnick’s simple and elegant StarLogo programming language allows an even more flexible framework within which to develop gridworld ecosystems for a variety of applications—see Resnick, 1994.) In contrast, all of the behavior of minimats in our current system is determined by the process of evolution, and nothing is predetermined; if a minimat never specifically chooses to reproduce, for instance (in this case by having an insufficiently high 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 the effects of certain cognitive or sensory capacities. Of these, a large number ignore the effects of the underlying physical or social environmental structure on behavior, but a few do deal with these effects in ways related to our concerns here. Ackley and Littman (1992) focus on the relation between learning and evolution, and later (1994) on the evolution of communication for altruistic purposes in a spatially structured environment that allows speciation. Werner and Dyer (1992, 1993) explore systems that illustrate specific kinds of communication and herding activity. Toquenaga, Kajitani, and Hoshino (1994) model the effects of resource patchiness on group foraging and flock formation in egrets. Lund and Parisi (1995, 1996) describe the emergence of generalist and specialist strategies in foraging, when the organism’s sensors, behavior, and energy-extraction ability are allowed to evolve together. Hartvigsen and Starmer (1995) discuss the effects of different environment shapes and boundary conditions on the coevolution of plant-herbivore defense and resistance interactions. Maley (1996) studies the effects of spatial environment structure on the success of new species invading existing ecosystems. 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 effect of multiple creatures on 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 (at present, beginning with a small number of primitive action types from which more complex behaviors can be formed) and then to 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 to our work in that he strives for simplicity in his system 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. Bedau and Packard (1992) further explore this type of system to develop measures of evolutionary activity, and Fletcher, Zwick, and Bedau (1996, in this issue) use an extension to study the effects of environment structure on the rate of evolution, yielding results complementary to those we report here. 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, 1994). 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), Gracias et al. (1996), Grant (1992), and the LEE system of Menczer and Belew (1994, 1996a, and 1996b, in this issue) introduce rich creature simulation environments in which a variety of types of behavior can evolve. But each come with predetermined decisions about the cognitive, sensory, or memory abilities of the creatures involved and the structures of the environments in which they live. In the SPO system, of course, we have also made such choices about the evolvable capabilities of minimats. But we hope our particular decisions allow a lower level of complication (if not also sophistication) in the simulations, more appropriate for the kind of study we are undertaking here into the simplest forms of blind but adapted behavior evolved across a variety of environment types.

3 The World

To understand the evolution of minimat behavior in response to different environmental structures, we must first describe the environments themselves. The minimat world is a simple place: It consists of a two-dimensional 64 x 64 grid of distinct square locations, toroidally connected at opposite edges so that movement wraps around and nobody will fall off. Any number of minimats 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 shorthand for energy, the common currency in the world—minimats absorb it through successfully eating, use it up in all their activities, and die when they run out of it.) Each individual minimat thus lives its life in an environment made up solely of other minimats and bits of food—and even the other minimats can largely be ignored, because there is no direct interminimat interaction. One minimat can only affect another by affecting the distribution of food in the world—that is, by eating food that the other individual could otherwise have had. By controlling their own spatial distribution over time, minimats have differing indirect environmental impact on each other. By controlling the


initial and ongoing growth of food in space and time, we have direct influence on the environment. These two interacting factors, direct experimenter control and indirect minimat impact, combine to create the actual environment in which the minimats live and evolve.

3.1 Basic environment structure parameters


There are an endless number of ways to parameterize the possible patterns of food distribution and growth across space and time in this 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-bearing locations in the world—that is, the percentage of locations in the world in which food can be present simultaneously. One can think of the world as having an underlying water supply on which all the food-plants depend; if the water supply is sufficiently large, food can grow in every location simultaneously (giving 100% density), but if the supply is insufficient, there may be only enough for half as many food-plants to grow at once (giving food in roughly every second location, for 50% density), or only a hundredth (giving about one food-plant in every 10 x 10 area, for 1% density), and so forth—see Figure 1 a for an example of a 5% food location density random environment, and 1 c for 40% density.

The maximum food location density thus controls the grossest aspect of spatial energy distribution in the environment. If food-bearing locations are distributed purely randomly in the world with a given density (as they are in many of the experiments described in this article), then the density also determines how far on average a minimat will have to travel from one fertile location before encountering another. Note that this parameter only affects the maximum *possible* density of food in the world, by controlling the density of food-bearing locations; as minimats eat the food in the world, the actual number of locations that currently contain food will go down, and so the effective food density that they experience will also go down (this becomes important in section 6.4.1).

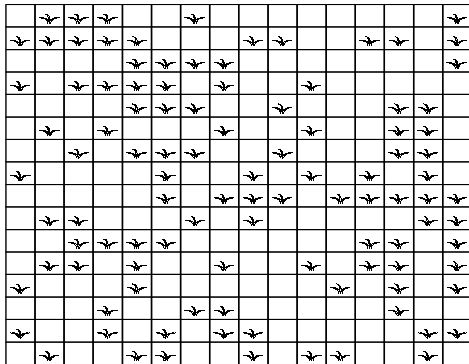
Food energy consumed in the world must be replenished to fuel the ongoing behaviors of the minimat population. The pattern over time with which food reappears in the world determines its temporal distribution. New food could be made to pop up anywhere in the world, but to keep things simple (and somewhat realistic), we allow food to regrow only in the set of fertile food-bearing locations determined when the world is first initialized. That is, the world begins with food-plants present in a set of locations whose prevalence is determined by the food location density parameter. Thereafter, new food can only regrow in one of those originally specified fertile locations. We specify a second environment structure parameter, the food regrowth delay, to control how long it takes food to regrow in a fertile location after its previous contents have been eaten. When a particular food-plant is eaten (consumed all at once by a lucky minimat, who swallows the entire energy contents of a particular location in one gulp), the now-foodless location will remain empty until the regrowth delay-specified number of time steps has passed. This can vary from zero—meaning no delay, so that food regrows immediately on the next time step—to a high number of time steps (maximally 32



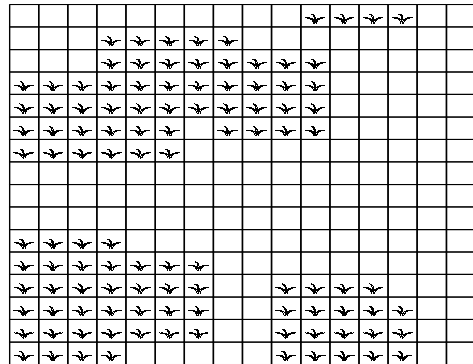
(a)



(b)



(c)



(d)

Figure 1: Example food distributions in a 16 x 16 subregion of the world for 5% maximum food location density with a clump scale of 1 (a) and 4 (b) and for 40% maximum food location density with a clump scale of 1 (c) and 4 (d) showing overlapping food clumps.

in our simulations here). The regrowth delay can be thought of as the number of time steps a location must lie fallow before it can support the growth of food-plants again. Whereas 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.

3.2 Controlling the amount of energy in the world

We now know where food can grow in the world, and when it can regrow—but we must also specify how *much* food will grow at any one time in a given location. In the minimat world, food-plants shoot up to some full “size”—amount of energy—as soon as they grow, and they stay that “size” until they are eaten. Thus any location in the world will only ever contain either no food energy or some maximum amount, nothing in between. The maximum amount of food per location can be specified directly, but we usually use a more indirect computation of this parameter. We specify the desired maximum energy “flux” in the world—that is, the greatest rate at which energy can reappear after consumption in the world per location per time step (where the average is defined across all world locations and all time). This is not the same as defining some constant amount of energy that rains down on the world at each time step. Instead, imagine a pasture in which grass can grow to some maximum height, so that there is some maximum amount of biomass, and hence food energy, that can be present. Now after the pasture has been grazed by a passing herd of antelope, it can regrow to its maximum energy level again in some amount of time. This maximum rate of regrowth, in energy per unit area per unit time, is the energy flux. Without grazing, that potential maximum flux cannot be realized—the pasture will just stay at its holding energy capacity until the next herbivore comes along.

The food energy flux parameter thus controls the maximum amount of energy that can flow through the minimat world. This flow only happens if there is an energy sink, though: the consumption of energy by the minimats through eating. Consumed energy will be replaced as quickly as it is eaten, up to the energy flux rate. As a consequence, the energy flux rate also controls the carrying capacity of the world—that is, the number of minimats that can survive simultaneously for an indefinite period. Minimats burn up a certain (fixed) amount of energy E per time step, and so by computing

$$C = flux * locations / E$$

we can find the maximum number (carrying capacity C) of minimats the world can support.

The energy flux is thus crucial for allowing us to make a level playing field for different minimat populations in differently structured environments. Our goal is to construct environments with various food location densities and regrowth delays, while holding the carrying capacity constant, so we can explore the effect of the first two spatial and temporal parameters alone. To do this, we need only vary one additional aspect of the environment: the “size” to which food-plants grow, that is, the amount of energy they each contain when they (re)appear. To see how the different environmental parameters are interrelated, first imagine that we

have constructed an environment in which food grows everywhere (100% density) and regrows instantly after being eaten (zero regrowth delay) to a level of 2.0 units of energy. Now we want to compare evolution in this first environment to what happens when we change the maximum food density to 50%. If we leave the maximum food amount per location at 2.0 units, then half of the locations will be able to grow 2.0 units of energy, and half will grow 0.0 units. Thus, in the second environment, we have the same maximum food amount, but half the maximum energy flux—only half as much energy is ever available for the population of minimats, and this can seriously affect their livelihood. To achieve equal energy flux, and equal carrying capacity, we must double the size of the food-plants (the amount of energy) that can grow at each location in the second environment.

The same thing applies if we have a food regrowth delay. Consider a third environment with 100% food density and a regrowth delay of 1 time step. In this environment, food appears only half as often (every second time step) as it does in the case of the first zero regrowth delay world. Therefore, to give this environment an equal energy flux, we again have to double the amount of food that grows in one place. In general, the energy flux can now be seen to correspond to:

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

and we can compute this food-plant size as follows:

$$food_plant_size = \frac{energy_flux * (1 + regrowth_delay)}{food_max_density}$$

3.3 Higher-order environment structure parameters

Just three of the preceding parameters (maximum food density, regrowth delay, and energy flux) are necessary to specify a complete spatiotemporal distribution of food in a particular environment. But such a distribution will be homogeneous across space and time, lacking higher-order structure. This is obviously an unrealistic portrayal of many natural environments and habitats, in which food and other resources are found instead in (often hierarchically) organized patterns (Bell, 1991). Resource items often occur in discrete patches separated by relatively empty (resourceless) space; the patches themselves are furthermore typically clustered into habitats separated by even greater distances. Much of the study of foraging behavior (Stephens & Krebs, 1986; Bell, 1991) relies on the notion of patchy environments that creatures must explore and exploit appropriately. To be able to model such environments more closely, we introduce one more parameter: the clump scale value.

This parameter allows us to add at least one level of hierarchical organization to the creatures’ environments, that of the resource patch. We stock the world with equal-sized, often overlapping, square patches of food-bearing locations. The clump scale specifies the length and width of these fertile square patches. Thus, for example, with a clump scale of 1, the environment consists of 1 x 1 “patches” of food (i.e., food grows in single locations—this is the same as having no patch structure at all), while for a clump scale of 4, the world is filled

with the appropriate number of (possibly overlapping) 4 x 4 patches needed to achieve the desired overall density of (individual) food-bearing locations. Figure 1 b shows an example environment with 5% food location density and a clump scale of 4, while Figure 1 d shows an environment with 40% food location density and clump scale 4.

At the beginning of time, before the introduction of minimats to the world, the environment is stocked with a full complement of food so that there will be plenty to eat for the initial population. The maximum food density, regrowth delay, and energy flux parameters are used to calculate the maximum food-energy per location, and then locations are filled with that much food-energy according to the desired maximum density and clump scale. Once this initial configuration of food has been determined, the food-bearing locations are the only places where food can ever grow throughout the ensuing history of this world.

4 The Minimats

As we described in our original statement of this research program (Todd & Wilson, 1993), we are ultimately interested in the evolution of three components of behavior in response to different environment structures: action, sensation, and memory. Only the first of these, action, is actually necessary to allow an organism to behave—and, indeed, to behave in a way well-adapted to its particular environment (Todd et al., 1994). That is, minimats need not be able to sense nor remember anything about their world (or themselves) in order to survive and reproduce in it, provided the world is generous and benign enough. Simply performing different actions with certain probabilities can suffice. Of course, no real organisms adopt this strategy (nor would they want to); it is impossible to build a physical system that cannot be affected by its immediate environment in *some* way. But we begin with this otherwise-impossible thought experiment as the logically simplest case we can construct, again with the goal of establishing a baseline against which it will be much easier to show the adaptiveness of evolved sensors and memory systems in comparison.

4.1 Action types

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

Minimats burn up energy with every action they make, and they die if they ever run out of this energy. The only way that minimats 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 is the food-plants growing in the world. When an individual performs the eat action, it attempts to absorb all the food-energy in its current location. Plants can only be consumed whole in these worlds; thus, if there are multiple minimats all trying to eat at one particular food-bearing location, only one of the minimats

will actually succeed. This successful minimat alone gets all the food-energy in that location, while all the others get no new energy. The winning minimat—the one that gets to eat—is chosen randomly from among all the contenders attempting to eat in a location. This minimat’s internal energy store is increased by the amount of food-energy that was available in the location. At the same time, the location’s food-energy is set to zero, ready to sprout again after the environmentally specified regrowth delay.

Minimats might not necessarily need to move about in their environment 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 there their whole lives. But we also allow minimats the option of moving, to open up a much wider and more realistic set of behavioral possibilities. Minimats can move one location at a time in one of the four orthogonal directions in the square grid (i.e., to the neighboring location to the north, south, east, or west of their current spot). Movement is always 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 choice of movement actually first resets the current heading so that the creature is facing the direction it is going to move (i.e., 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 important implications for the minimats’ behavior, as we will see in section 6.

The ultimate indication that particular behavior patterns have been well-suited to a given environment is the presence of minimats with those behaviors populating that environment after a suitable length of time. We could arrive at this state simply by filling the world with randomly generated minimats and seeing who remains (with what behaviors) after the selective filtering process of death weeds out the maladaptive ones. But besides not being a particularly realistic portrayal of evolution, this method could miss well-adapted patterns of behavior that happened to be absent from the initial population. Instead, we need the power of evolutionary search through small adjustments to existing behavior patterns, and for this, we need replication with modification. Hence, we allow the minimats to reproduce by asexual splitting.

When a minimat splits, it creates a possibly slightly modified copy of itself and divides its current energy between itself and its new “offspring.” The newborn is placed in the same location as the original splitting minimat (though with a randomly-assigned heading), and they go on to lead separate lives from that instant. The rare modifications in the newborn consist of changes in its probabilities of choosing the various actions in the minimat’s behavioral repertoire (see the next section on action selection). The mutation rate—the chance that any given action probability will be replaced with a new probability value—is 5% throughout the runs described here. This results in an individual who will behave slightly differently from the parent it split off from, allowing evolutionary search of the space of possible behavior patterns.

Finally, we allow minimats to “choose” to do nothing in a particular time step, just sit and wait for the moment to pass. In this way, minimats can “travel” through time in the same way they can traverse space. We hoped that including this action would allow minimats to capitalize on the temporal density of food in their world, in the same way that movement

operates on spatial density, but we have not yet seen evidence for this.

Every action a minimat performs has energetic side effects determined by the bioenergetic characteristics of the world (see Todd & Wilson, 1993; other approaches to energy use are presented by Menczer & Belew, 1994, 1996a, and 1996b in this issue, and Lund & Parisi, 1995, 1996). Actions all use up a certain amount of energy; throughout the simulations reported here, every action uses a constant 0.5 units of energy when performed (so minimats always use up 0.5 units of energy in every time step they are alive). Beyond these action-costs, some actions can further increase or decrease an individual's energy. Successful eating increases 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 minimats when they eat, and flows back out of the minimats by being burned up through the actions they perform. Because minimats only die when they run out of energy, no energy-filled carcasses ever end up in the world—dead minimats just disappear.

4.2 Action selection

Deciding how to “do the right thing,” while an open challenge for most simulated agent research (Maes, 1994), is not a very involved process for minimats. Minimats have neither sensors nor memories or internal states to distinguish any location from any other, nor any time step from the rest of eternity. As a consequence, there is not much they can use to decide which actions to perform in a given location at a given instant. In fact, all that minimats *can* use is the roll of a die—they select among all their 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 minimat, that determine the makeup of its behavioral repertoire, and which it uses to “decide” what to do next. Evolution shapes these probabilities over time so that minimats will behave in a way that is more or less well adapted to the environment they live in. (We also allow an action to be “turned off” directly by evolution, changing its probability to zero in one step by making that action inactive; if the action is mutated back to being active in a later minimat, its old probability level is restored, allowing a sort of genetic memory.) For instance, it is generally a bad idea in *any* environment for a minimat to split more often than it eats: every time a minimat splits it loses half of its energy (as well as paying the splitting action energy cost), and without eating often enough to replenish that lost energy, the minimat (and all its similarly low-energy offspring) would soon perish. Hence, we are unlikely to see a population of evolved minimats with an eating probability lower than their splitting probability.

Each minimat thus has an internal list of seven probabilities for selecting among the seven possible actions. These probabilities remain constant throughout a minimat's lifetime (minimats neither learn nor change developmentally). When a minimat splits, it is these probabilities that are modified slightly in the newborn offspring.

At every time step, each minimat chooses one action at random according to its own internal probability distribution and then performs that action. These probability distributions are therefore all we can look at to assess the behavioral repertoire of the minimats 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, the efforts of other researchers trying to make sense of the behaviors of a population of sensing, remembering, and acting neural-network-based simulated creatures are fraught with much greater difficulties. And as we will see in section 6, 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 What We Did

To collect most of the data reported in this article, we ran populations of minimats for 15,000 time steps and computed statistics for the final population at the end of each run. Our earlier studies had shown that the evolutionary dynamics of changing behaviors, population size, and environmental impact (e.g., changing food density as a result of minimat eating) had all settled down well before 15,000 steps. For each environment structure, we averaged five successful runs and computed the standard deviation to get an idea of the degree of variation of evolved behavioral solutions across runs. However, because of occasional extinctions, not all runs are successful (see section 6.1), and so we ran as many simulations at each parameter setting as were necessary to produce the five that went to completion with 15,000 steps.

Each run was begun with a population of 500 minimats, evenly dispersed across the environment. Food was initially stocked at its maximum specified density, in energy amounts specified by the energy flux equation. Each of the first-generation minimats had a randomly assigned behavior probability distribution, a random heading, and 20 units of energy to start out life with. 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 compete with each other for domination. But new behaviors also appear as time proceeds and the minimats split, further fueling the evolutionary competition throughout each run.

The first thing that happens as the initial random minimats begin to act in any environment is a rapid population explosion, from the starting point of 500 individuals up to somewhere in the neighborhood of 1500 individuals within the first few time steps. This is because the initial population contains a number of minimats with very high probabilities of splitting, so that is what they do. But because reproducing is basically *all* some of these oversexed minimats do, they soon run out of energy and die, and the population crashes to around 300 more self-controlled minimats. From that point on, the successful minimats begin to make their presence felt, as they responsibly move about their world, eat, and only very occasionally have offspring. The population size grows steadily over the next several hundred time steps before leveling off. This initial population explosion followed by a rapid collapse and slower rebuilding is seen in the majority of runs. The

final stable population size that is reached, though, and the behaviors that evolve within that population, depend greatly on the exact structure of the environment.

We focus here on understanding the evolutionary effect of environments created with maximum food location density between 20% and 100%, where the patterns are clearest. (Note that we describe density in percent in the text, but label it from 0.0 to 1.0 in the graphs.) Similarly, we begin our discussion of results by considering regrowth delays between 0 and 4 time steps, which correspond to 100% temporal density (food regrows at every time step) down to 20% (food every fifth time step). Thus we have covered matching ranges of spatial and temporal density. We proceed to consider a wider range of regrowth delays and then look at the effects of clump scale on evolved behavior.

The minimats and their environment were simulated on a 4096-processor 7 MHz Connection Machine 2 (CM-2) using the *Lisp data-parallel language. We achieved a computation rate of about five time steps per second for a population of up to 16,000 minimats.

6 What Happened

6.1 Overall measures of evolved populations

One of the most obvious indications that different environment structures have an effect on minimat evolution would simply be systematic changes in ultimate population size as environmental parameters change. In Figure 2, we see that this is exactly what happens. In general, as the maximum food density increases, so too does the population size attainable in that environment. It is easier to find food, the minimat's central task, because there are more food-bearing locations (though each food-plant contains less energy, to maintain a constant energy flux). Similarly, as food regrowth delay increases from 0 to 4, again so does the population size. Things get a bit messy for low food densities (under 20% food-bearing locations), where these general trends seem to reverse. In addition, there is a large standard deviation for 90% food location density (also seen in several of the later graphs), which arises because of the underlying variation in evolved patterns of behavior between different populations. (Generally, large error-bars mean there were a variety of different evolved strategies in a particular environment.) But overall the pattern is clear: Changes in food density and regrowth delay bring about clear differences in evolved population size.

The large population sizes with increasing food density of Figure 2, though, mask one of the first mysteries of minimat evolution. As we indicated in the previous section, the data shown in these figures is averaged across *successful* runs, where the population made it to 15,000 time steps. But in different environments, there are varying degrees of evolutionary success. As food density increases, there is actually a *decreasing* rate of successful runs—that is, higher food densities lead to a greater number of early extinctions across runs. In the “toughest” environment we looked at, with a food density of 80% and no regrowth delay, 64% of the populations (9 of the 14 runs we had to do to get 5 successful ones) quickly go extinct—this despite food being present in almost every location, and regrowing as soon as it is eaten! This trend largely

disappears, however, at the highest food density, 100%, and as regrowth delay increases, when almost all runs successfully reach 15,000 steps.

We do not yet know why many seemingly benign, plentiful environments present so much difficulty to the evolving populations, nor why a small change in food density or regrowth delay can have such a large impact on evolutionary success. The evolved behavioral patterns do not show corresponding jumps at these environmental parameter settings and so offer us no clues. One possibility is that there is an important interaction between the amount of energy per food-plant that can be eaten and the energetic costs of the minimat's moment-to-moment actions. If each food-plant eaten does not give the average minimat enough energy to live long enough to find the next piece of food, we would expect frequent extinctions in these environments. We will vary the energy costs of actions (or the energy flux) in different environments to check this possible cause of environmental difficulty for minimat evolution.

6.2 General patterns of evolved behaviors

6.2.1 Minimats spend most of their time eating

Our main interest, of course, is not how many minimats there are in a population, nor how successful their different populations are, but rather what they *do* (cf. Fletcher et al., 1996, who focus on evolutionary “success”). How do changes in environmental structure affect the behaviors that evolve? In Figure 3, we see the answers. As we saw for population size, the patterns are different for very low food location densities, but above density 20%, we see movement alone (summed across all four movement directions) increasing with greater food density, and eating, splitting, and sitting decreasing. Greater regrowth delays also make for more movement and less eating. Overall, minimats in this range of environments spend the vast majority of their time (over 80%) just trying to eat whatever food might be in their current location and only move occasionally (2-18%) to other possibly greener pastures. They sit still infrequently—it seems better at least to try to eat, if one is not moving on, than not even to take the chance by just sitting, because in either case the same amount of energy (0.5 units) is lost on this time step. (But then why should minimats *ever* sit rather than eat? This action might remain in the population due to mutation, with insufficient selection pressure to decrease its prevalence further. But this does not explain its systematic changes across environments, so some mystery remains.) And minimats reproduce only rarely, because of the energetic costliness of this action; on average (for higher food densities), only about one minimat in 200 will split on a given time step.

6.2.2 Moving makes sense at high food densities

The majority of excitement in a minimat's life comes from eating and sometimes blindly stumbling to a neighboring location, and it is these actions that we have endeavored most to understand. At high food location densities, the likelihood of moving onto a neighboring spot that can grow food is also high; therefore, movement is a low-risk strategy; anywhere else is just about as likely to be a good grazing spot as where




Figure 2: Population size vs. maximum food location density for four different regrowth delays, showing that the number of minimats in the world goes up as food becomes more common. Note that the large error bars for 90% food location density (here and in later figures) indicate a greater variety of evolved behaviors at that parameter value.

the minimat is right now. Furthermore, moving can provide a valuable payoff: A minimat can escape the food-scramble competition in its current location. Granted, it will probably end up in a new, similarly competitive situation; but the alternative, never to move, will cause minimats and their offspring to end up piling higher and higher on one spot, parents and children starving one another in their struggle for the limited local resources. Thus, those minimats that do move provide a better chance of survival for themselves and their offspring, and this behavior pattern will tend to spread through the population.

6.2.3 Moving seems risky if food is sparse

But what of the minimats stuck in sparse-food environments, where it may take several successive movements to get from one fertile food-bearing location to another? In these cases, rarely moving seems like a worse strategy than never moving at all: Consider, for example, a minimat born on a fertile location in a 20% food density world, who moves once every 20 to 30 time steps. Once this minimat moves off of its birthplace, it will (if moving randomly) probably take about four more movements to find another fertile spot. Given its low movement rate, this could take 100 foodless time steps, and by that point, the minimat could well have starved to death. So for such sparse environments, it may make more sense to stay in one location by never moving. Why then do we find a population average movement rate of 2 to 8% in 20% density environments, and even 5% movement in the very sparse 5% density environments?

Because averages obscure the real story. There is no “average minimat” with 2.3 offspring and a Volvo (only driven 5% of the time). In fact, some minimats never move—and some never reproduce—at all. As we indicated in section 1.2, each minimat population is composed of distinct subpopulations

with quite different behavior patterns, and to understand how environments affect minimat behavior, we must consider the subpopulations themselves.

6.2.4 Most minimats do not move if food is sparse

So to answer our question of why there is still a low average movement rate at low food densities, we need to see if all minimats in the population are moving equally. The answer is a resounding “no”—most minimats never move once in their entire lifetime. In Figure 4 a, we show the percentage of minimats in the population that have any active move actions—that is, that can ever move—as a function of food density and regrowth delay. Over 80% of the minimats in low food density environments are completely immobile, spending their whole lives eating in the same location (and sometimes splitting or sitting), just as we suspected they should two paragraphs ago. As food density (and regrowth delay) increases, so too do the number of movers in the population, until at 100% food density nearly all minimats are moving, again as we expected.

6.2.5 Moving minimats move about the same amount

A subpopulation of mobile minimats is thus responsible for all of the “average” movement we find in Figure 3 b, while the rest sit idly by. How much are these movers themselves actually moving? Surprisingly, this is roughly constant across environments, as we see in Figure 4 b. At very low food densities, those minimats that move at all, move quite a bit (20-40%), which matches our earlier suspicions about the folly of low movement rates in sparse environments. Movement rates of 10 to 20% change rather little across the higher food densities (those above 20%), though. Rather, it seems that evolution changes the mix of a small set of subpopulations in many different environments, while keeping the behaviors

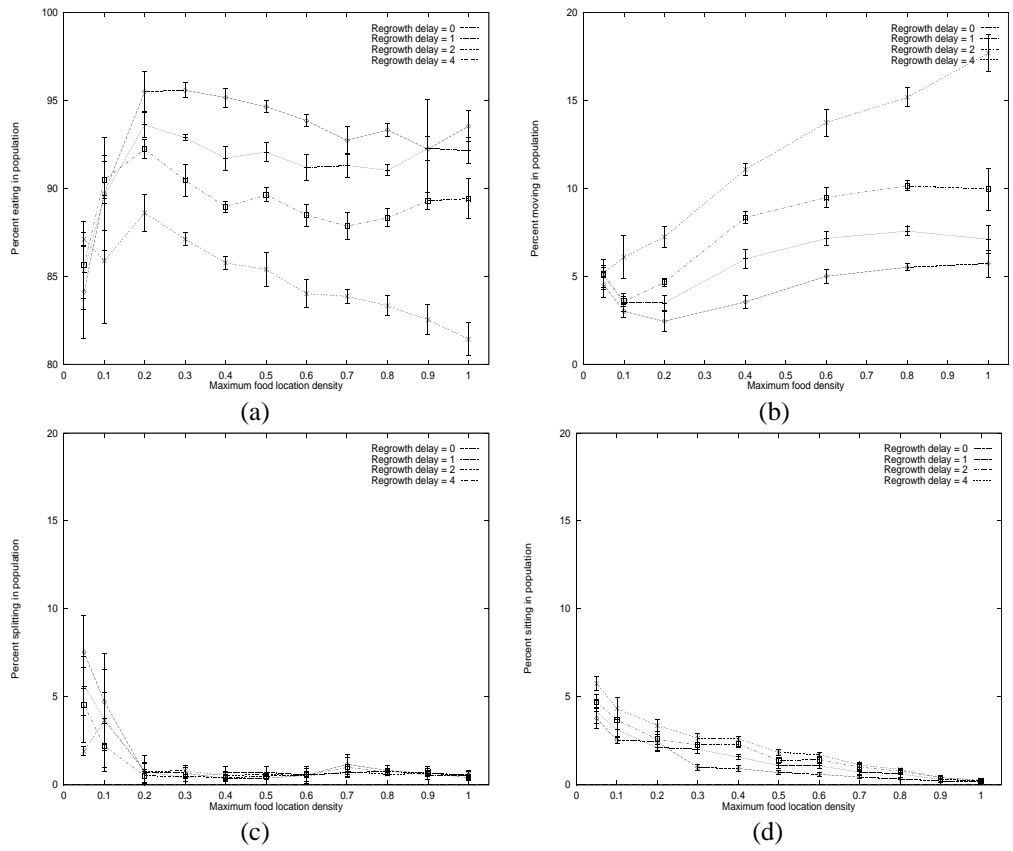




Figure 3: The percentage of time that minimats eat (a), move (b), split (c), and sit (d) vs. maximum food location density, showing that evolved minimat behavior changes in regular ways with changes in environment structure. Note that the y-axes are all the same scale but have different ranges, with eating being the most common action by far, followed by moving.



(a)



(b)

Figure 4: (a) The percentage of mobile minimats (Movers) in the population vs. the maximum food location density, showing that movement increases overall in the population as food becomes more common, because the number of Movers increases. (b) The percentage of time that Movers actually move vs. the maximum food location density, showing that Movers spend about the same amount of time moving across a wide range of food location densities.

within each of those subpopulations largely constant. Thus, we need to consider not only the minimat behaviors adapted to different environments, but also the mix of these behaviors across environments and the evolutionary stable strategies these mixes represent. (We save this game-theoretic analysis for the future.)

The importance of looking at the mix of roughly behaviorally constant subpopulations is reinforced when we consider the rates of reproduction in the population. Figure 5 a shows the percentage of potential reproducers (dubbed “splitters”) in the population (those minimats with nonzero reproduction rates), which rises for low and high food densities. Figure 5 b shows the percentage of time the potential reproducers actually spend splitting, which again stays more constant across environments (again once food density is above 20%), even as the reproducing/nonreproducing subpopulation mix changes. (Higher splitting rates at the low food location densities are probably possible because each food-plant eaten provides a large surge of energy, which the minimat can use to reproduce many times without running out.)

6.2.6 Many minimats do not reproduce

The relationships between moving and reproducing can be seen by considering four different subpopulations: minimats who both move and reproduce, or those who move without reproducing, or reproduce without moving, or neither move nor reproduce. In Figure 6, we show the number of individuals (rather than population percentages) in each of these four subpopulations supported by different food location densities (all for regrowth delay 0—similar patterns hold for other delays). The number of movers (both reproducing and nonreproducing) increases over tenfold across the food density range, while the number of nonmovers declines. (When combined, of course, the four subpopulations yield the population sizes shown in Fig. 2.)

Interestingly, the number of minimats that do not reproduce is greater for both movers and nonmovers alike, across all food densities. These minimats never pay the high energy cost of splitting in two and so are more likely to amass greater amounts of energy and stay further away from starvation than their more prolific compatriots. Thus we do not yet seem to have overcome the unnatural presence of potentially immortal minimats we noted early in our studies (see Todd, 1993). But (nearly) all populations do have at least some fertile members, as shown in Figure 5 a, so the problem we earlier feared, of evolution coming to a halt through lack of reproduction, is not a concern here.

6.3 Movement effects of food location density

Now let us return again to our question of movement rates in sparse food environments. We saw that most minimats (80-90%) do not move at all in such environments, and those few that do, move rather often (20-40%). This lessens the time spent between fertile locations, and so helps avoid starvation in empty desert regions. But of course, it also lessens the time spent on any fertile locations before moving off again. Overall, movement still seems to be a risky proposition in these environments—why should *any* portion of the population leave a food-bearing location? Early on in the course of evolution,

such movement could help wandering minimats come upon fertile locations that were not already filled with sessile eaters. But after enough time steps, for instance, the 15,000 time-step case we are considering here, the population will have settled to a steady state dominated by the nonmoving minimats, crowded onto the scattered fertile locations (as many as eight minimats per location at 5% food location density, and three per location at 20% food location density). In such a situation, random movement seems particularly hopeless: Even if a mobile minimat does find another food-bearing location, it will already be swarming with immobile competitors.

6.3.1 Minimats do not move randomly


Faced with this puzzle, we decided to check our assumptions. Random movement implies that each minimat would move in each of the four possible directions approximately equally often. We counted up the number of directions each mobile minimat could move in, expecting an average close to 4. To our surprise, the average came out very close to 1—most minimats can only move in a single direction. Moreover, this evolutionarily restricted, nonrandom movement is very quickly and strongly selected for (at least at low food densities), as Figure 7 shows—within 2000 time steps of this typical run (10% food location density, regrowth delay 4), nearly 80% of the movers travel with just one movement type.

So our assumption of random movement had been wrong. What were these minimats doing instead? There are four possibilities: If they only ever move in direction 0—straight ahead—they will move in a straight line across their toroidal world, wrapping back to where they started after covering 64 locations. If they move solely in direction 1—turning right and taking a step ahead—they will describe a four-location loop, ending up back where they started after executing four move-actions. This is the same for direction 3, except that the four-move loop winds to the left instead. Finally, if a minimat moves repeatedly in direction 2—turning around 180 degrees and then stepping forward—it will simply loop back and forth over two locations.


6.3.2 Minimats mostly move in small loops

Now it remains only to find out *which* of these types of loops—whether of length 64, 4, or 2—our moving minimats are circumnavigating. In Figure 8, we have plotted just this, across different food densities, holding regrowth delay constant at 0 time steps². (We will explore the effects of regrowth delay in the next section.) Truly random movement—those minimats who move in more than one direction—is very low across most food densities, only rising at the very highest densities, when moving anywhere is likely to lead to food. Moving in a 64-location loop, which crosses the entire world, also remains an uncommon strategy until the higher food densities. (For very low food location densities, as we found earlier, the story is more complicated.) But with 2- and 4-location loops, things get more interesting.

²Note that we have added together the minimats moving only in direction 1 and those moving only in direction 3, because in both cases they are traveling on 4-location loops. The direction of these loops does not make a behavioral difference. Throughout the runs, movements 1 and 3 are about equally common.



(a)



(b)

Figure 5: (a) The percentage of reproducing minimats (Splitters) in the population vs. the maximum food location density. (b) The percentage of time that Splitters actually split vs. the maximum food location density, showing that Splitters spend about the same amount of time splitting across a wide range of food location densities.




Figure 6: The number of minimats in each of four behaviorally distinct subpopulations, differentiated by their moving and reproducing ability or lack thereof, vs. maximum food location density. Moving minimats increase with increasing food location density, and avoiding reproduction remains a common strategy across all environments.




Figure 7: Percentage of mobile minimats (Movers) that have one, two, three, or four active movement actions (directions of movement) over time for one entire example run of 15,000 cycles (with maximum food location density 10% and regrowth delay 4). Minimats that move in only one direction are rapidly selected for.




Figure 8: Percentage of mobile minimats (Movers) who move in loops of lengths 2, 4, or 64 or without looping vs. maximum food location density. Short loops predominate at low food densities to allow minimats to return to particular food-bearing locations. (Error bars for the rates of nonlooping movements are omitted for clarity; at high food densities the large error bars indicate a variety of evolved behavioral strategies.)

Moving off of a fertile location can make some sense, even at low food densities, if one could be guaranteed of being able to return to that fertile location in the near future. This is exactly what 2- and 4-location loops allow—after the requisite number of movements, the minimat will be back where it started. Notice though that we do not know exactly how long it will take for the minimat to return to any given location. Because all actions are chosen probabilistically, it may take some time before the minimat executes the two or four moves necessary, interspersed with its other (more common) eating and (less common) splitting and sitting actions. (We can compute the approximate time to complete a loop, though, which will become useful when we consider regrowth delay in the next section.) But sooner or later, the minimat will get back to any food-bearing location it left, and in the meantime, it will have visited one (for 2-move loops) or three (for 4-move loops) other locations. And this is where the benefit of looping in sparse environments comes in: Mobile minimats will have a chance of finding other food-bearing locations that their immobile cousins will not.

6.4 Action effects of regrowth delay

We have now seen how the most obvious aspect of a minimat's world—the density of food locations—affects the evolution of minimat behaviors, including prevalence and type of looping movements. We also saw a hint in the initial graphs of how the second environmental parameter, regrowth delay, affects evolved behavior, and how the two parameters interact. But our initial range of regrowth delays, as presented in Figure 3, was too restricted to get a complete picture of its behavioral impact. We now consider a wider range of regrowth delays and see whether or not they have an independent effect on evolved minimat behavior.

6.4.1 Regrowth delay affects evolved behavior

In Figure 9, we show the evolved average mix of minimat behaviors in environments with regrowth delay ranging from 0 to 32 (in nonequal jumps), all for a fixed food location density of 40% (a midrange density). As the time taken for food to reappear goes up, the amount of time minimats spend moving also goes up, while their eating rates go down. Splitting and sitting rates remain very low throughout, changing little. This pattern of increasing movement and decreasing eating mirrors the changes we saw with increasing food location density in Figure 3. So is regrowth delay having the same effect on evolved behavior as food density? Can this temporal variable simply be mapped onto the spatial one, as far as minimats are concerned? In our quest to characterize the environmental structure that matters to minimats, can we get away with using just a single parameter?

The answers to these questions could of course go either way. Temporal structure seems, to us memory-endowed observers, to be an important element of any environment; but to minimats, lacking memory and sensation of their environment, it could well be that taking advantage of temporal structure is a luxury they are forced to do without. Certainly regrowth delay has an effect on evolved minimat behavior, as we saw in Figure 9. But if regrowth delay has, first, an effect on food density, and it is these food density differences that are actually influencing minimat evolution, then we would be justified in saying that regrowth delay as such is not the parameter affecting evolved behavior (see section 1.1). This would mean we could talk about all the minimat environments we have considered so far, purely in terms of their effective food density, when analyzing them for their expected behavioral impacts.

It is the case that regrowth delay has a direct effect on the




Figure 9: Percentage of time that minimats eat, move, sit, and split, for food regrowth delays from 0 to 32, with maximum food location density 40%. These patterns are similar to those in Figure 3, but for different reasons (see text).

current density of food in the world. With increasing delays, food takes longer to reappear in a given grazed location, and so at any given time step, there will be fewer locations containing food—and hence, actual food density will be lower (see section 3.1). (Of course, this only occurs as a result of minimat eating, making this a clear instance of the minimat population’s impact on its own environment.) This can clearly be seen in Figure 10, where we plot the actual food density experienced by minimats against increasing regrowth delay in an environment with 40% maximum food location density. With zero regrowth delay, food regrows instantly to its maximum 40% density, but as delay increases, less food has regrown at any given time step and the effective experienced density in the environment decreases.

6.4.2 Regrowth delay has a different effect from food location density

But before we conclude that regrowth delay does nothing more than change the food density that influences minimat behavior, we have to see whether or not the behavioral effects of regrowth delay vary with its effects on density in the same way that behavior varies with density in Figure 3. In the latter, eating went down and movement went up as density increased. Here, density increases with lower regrowth delays, so if the two parameters have matching effects, we should see less eating at lower regrowth delays, and more moving. But this is exactly the opposite of what Figure 9 shows. Instead, as regrowth delay goes *up* and food density goes *down*, we get increasing movement and decreasing eating. So the effect that regrowth delay has on food density does *not* account for its effect on evolved behaviors—in fact, just the opposite. We must conclude that regrowth delay is indeed an independent parameter of environment structure that makes a difference to evolving minimats.

We can understand the general trends in evolved behaviors

for different regrowth delays (Fig. 9) as follows: Eating is of course still the most important thing to do. But as regrowth delays increase, moving makes more and more survival sense as well. If food has been grazed from a minimat’s current location and will not appear there again for some time, it is not a particularly good idea to continue to sit there and just try to eat, gnawing at the empty ground. (This is different from the situation of instant food regrowth in locations with many competing minimats. There, it is a reasonable strategy, adopted by the many immobile minimats, to sit in one location and just try to eat at every time step, because there is always a chance that they will be the lucky one to get the food that time. With long regrowth delays, there is no chance for *anyone* to get food from a grazed location for some time.) Moving on in search of other fertile locations can be more rewarding, and so the average rate of movement increases with increasing regrowth delay. In fact, as we saw in Figure 4, any single mobile minimat in one of these environments moves about the same amount as any other—roughly 30% of the time—but the number of mobile minimats themselves increases with regrowth delay, accounting for the growing average movement rate in Figure 9.

6.4.3 Minimats can loop to “time” their environment

Moving away from a slowly regrowing food location rather than uselessly waiting around for its reblossoming can be energetically rewarding because it allows the minimat to find other food-bearing locations. But it still means leaving behind a known food-bearing location for possible starvation in the unknown, unless, as we saw for sparse food environments, there is a guarantee of returning to the original regrowing location. And so here too again looping behavior proves useful. If a minimat can graze a food-bearing location, leave it in the hopes of finding other food-bearing locations, and then return to the first location just as its regrowth delay is up




Figure 10: Actual average food location density vs. food regrowth delay for maximum food location density 40%. The actual food location density goes down with increasing delay because it takes longer for food eaten by the minimats to regrow in the world.

and food reappears to be eaten, the minimat will be in ideal sync with the temporal structure of the environment and will be able to exploit maximally the available resources.

Of course, no minimat can time things this precisely, lacking any of the internal state necessary for such controlled behavior. But they can come close, in a probabilistic sense. Recall that the mobile minimats in the environments of Figure 9 spend roughly 30% of their time moving, and so about 70% of the time eating. This means that a 2-move loop will take about 6 time steps, and a 4-move loop will take about 12. It is most adaptive to loop in such a way that the minimat will not return to a given location *before* food could have regrown there, but rather after, when it has already reappeared and is again available for eating. Thus we would predict that 2-move loops taking 6 time steps will be most beneficial in environments with regrowth delays up to 6, and 4-move-loops will be most beneficial for regrowth delays above 6, up to 12 or so.

In Figure 11, showing the prevalence of the different loop types across increasing regrowth delays, we see that this is largely the case: The 2-move loops are used most at low delays, while 4-move loops are used more at higher delays, and the two types cross just at a delay of six time steps. The proportion of 4-move loops does not fall off very quickly as delay gets even longer, but eventually nonlooping movement (and some 64-move loops) start to increase when the delays get too long for looping to accommodate.

6.5 Action effects of clump scale

So far we have seen that evolved minimat behavior is sensitive to both the overall density of food locations in the environment and the delay with which food regrows. These are, in some sense, the lowest-order spatial and temporal parameters with which an environment could be characterized: They re-

fect mere prevalence of resources in space and time, without capturing any aspect of higher-level *organization* or structure in space and time. We next sought to investigate whether minimat behavior *could* evolve to capitalize on the kinds of higher-level environmental structure that is so obvious to organisms (such as ourselves) that can compare environmental conditions across space and time. We introduced an environmental structure parameter that controls the clumpiness or patchiness with which food grows in the world, as this is one of the most commonly seen and analyzed higher-level environmental structures (see section 3.3). We did not hold out much hope that our minimats could evolve to take advantage of patch structure in their world, given their sensory and cognitive limitations. We instead expected that here would be the first example of environment structure that does *not* matter in the minimat world.

We tested minimat populations in environments with clump scales 1, 2, 3, 4, 5, 7, and 10, all with 5% overall food location density and zero regrowth delay (see Fig. 1 for examples of these environments with clump scales 1 and 4). The effect of clump scale on average evolved behaviors is shown in Figure 12. As food in the world appears in bigger and fewer patches, the rate of eating decreases (though slightly), movement increases, and splitting and sitting decrease marginally. These trends mirror those for increasing food density (Fig. 3), just as the patterns for increasing regrowth delay did (Fig. 9). Now, as we did for regrowth delay, we must ask if clump scale affects evolved behaviors in some new way, or merely through its effects on food density? Is it an independent environmental structure parameter that matters, or not?

6.5.1 Clump scale does not affect overall food density

The answer at first seems clear: As we changed clump scale, we held the overall food location density constant (by de-




Figure 11: Percentage of mobile minimats (Movers) who move in loops of lengths 2, 4, or 64 or without looping for food regrowth delays from 0 to 32, with maximum food location density 40%. Short loops predominate at short delays to allow minimats to “time” their environment and return to food-bearing locations after food has regrown. (Error bars for the rates of nonlooping movements are omitted for clarity.)




Figure 12: Percentage of time that minimats eat, move, split, and sit for clump scales from 1 to 10, with maximum food location density 5% and food regrowth delay 0. Clump scale affects behavior similarly to food location density (Fig. 3).

creasing the number of bigger clumps in the environment accordingly). So clump scale is not affecting overall food density, and its effects must be independent of that density. But overall food density in a clumpy world is not something that most minimats can ever assess. As we saw in previous environments, most minimats that move do so in small loops, and that is also the case in clumpy environments. So most minimats only ever see a very small portion of the total world (this is true not only for the mobile looping minimats, but of course also for those that never move at all). Their concern then is not overall worldwide food location density, but only the food density in their own tiny pathway. When the environment is not patchy (clump scale 1), containing no higher-order structure, then the statistics of food density in a small local area such as a 4-location loop do match the density statistics of the whole world (e.g., if overall food density is 75%, then on average, 75% of the locations in a 4-location loop will contain food, too). But for clumpier environments, this relationship breaks down, and the world that minimats experience may not reflect the characteristics of the world as a whole. This leads to complications, as we will see.

6.5.2 But clump scale does affect experienced food density

Consider a minimat living out its life by eating and making 2-location loops. If this minimat lives in a nonclumpy world with 5% food density, then it has a $2 * 5\% * 95\% = 9.5\%$ chance of living in a 2-location loop with one fertile, food-bearing location and one nonfertile location in it. It also has a $5\% * 5\% = .25\%$ chance of living in a 2-location loop with both locations being fertile. The rest of the time, it would find itself in a pair of locations that are both nonfertile, and because it would then never find food, it would quickly starve and disappear from the population. Thus the only cases we really care about for 2-location loopers—the only ones where these minimats could possibly stick around long enough to be analyzed in our system—are those with one or both locations bearing food. If we renormalize for just those two conditions, we find that these (nonstarving) minimats in the 5% food world will live in 2-location loops with one food-bearing location $9.5/9.75 = 97\%$ of the time, and will have both locations capable of growing food just 3% of the time.

In contrast, consider our 2-location-looping minimat in a clumpy world with clump scale 10 (and still overall food density 5%), where the world is filled with 10×10 squares of food-growing fields. How often will it live in a 2-location loop with only one food-bearing location? This will only happen if the minimat's loop happens to lie half on and half off one of the food patches, and this will happen roughly 36% of the time (because there are 36 edge locations in a 10×10 square, but some squares might overlap—again we ignore all cases where the minimat would live in totally empty space, because it would quickly die there). The other 64% of the time, the minimat would find itself in the interior of a food patch, and it would thus live on a 2-location loop with food in *both* locations. Thus, to a minimat living in a patchy environment, the world looks very different from the randomly distributed food case. In particular, food appears much denser, precisely because food *is* much denser anywhere that a minimat can possibly survive. So to compare patchy

and nonpatchy (nonclumpy) environments, we must compare how they appear to behaving minimats, not how they appear to omniscient environment designers with a global perspective.

We have performed this kind of comparison from the perspective of 2-location looping minimats in Table 1. First, we have listed the percentage of 2-location loops that have one or two food locations in them (ignoring empty loops), for nonpatchy (clump scale 1) environments with varying overall maximum food location densities of 5 to 100%. Next, we have performed this same analysis for patchy environments, for clump scales between 1 and 10, all for overall maximum food location density of 5%. Finally, we have shown the closest match between each of the patchy environments and the nonpatchy environments, mapping from clump scale to corresponding experienced apparent food location density (from the minimat's perspective). We see that as clump scale increases, so too does the experienced food density, which rapidly approaches asymptote. Thus 2-location looping minimats living in patchy worlds will “feel” that they are living in worlds with higher overall food density than is actually the case. (This analysis also holds for the 4-location loopers, though it is slightly more complex for them.)

So the mapping from clump scale to experienced food density for loopers and nonmovers alike is clear: Increasing clump scale corresponds to increasing experienced food density. If we then compose this relationship with the relationship between food density and behavior in Figure 3, we get the same pattern of results as in Figure 12: Increasing clump scale and correspondingly the food density corresponds to decreasing rates of eating and increasing rates of movement. By this test, changing clump scale is not doing anything to affect evolved behaviors other than changing the experienced food density. This is further supported by the percentages of different types of looping behavior across clump scales, shown in Figure 13. The pattern there matches that for the loopers at different food densities shown in Figure 8. Thus, on the basis of these pieces of evidence, we would seem to be able to eliminate patchiness as an independent structural parameter from our description of minimat environments.

6.5.3 Clumpy worlds look different to different loopers

But there is one more piece of data to consider before finalizing this conclusion. Look again at Figure 8. The type of movement that becomes most prominent there at high food densities is moving straight ahead (movement 0), that is, circumnavigating the world in a 64-location loop. But in Figure 13, this type of movement stays low across high clump scales, even though the experienced food density (for 2- and 4-location loopers) is also high. The reason for this lack of forward-moving individuals in clumpy worlds is apparent when we take the minimat's-eye perspective again: While a 2- or 4-location looping minimat only “sees” the prevalence of food in its tiny local loop neighborhood, a forward-moving minimat on a 64-location loop would “experience” a much wider swath of the world and so would encounter food at a rate closer to the true overall food location density³. Because

³Note though that patchiness does also affect the food density experienced by forward-moving minimats. This is because all nonstarving ones will cross at least one patch, and so will encounter a stretch of food at least *clump scale* locations wide. As a consequence, the experienced food location density seen

Percentage of non-empty 2-step loops with 1 or 2 food-bearing locations vs. clump scale:

	<i>Clump scale</i>									
	1	2	3	4	5	7	10	15	20	25
1:	97	65	48	38	32	24	18	12	8	7
2:	3	35	52	62	68	76	82	88	92	93

Percentage of non-empty 2-step loops with 1 or 2 food-bearing locations vs. maximum food location density:

	<i>Maximum food location density</i>										
	5%	10%	20%	30%	40%	50%	60%	70%	80%	90%	100%
1:	98	94	89	82	75	66	57	45	33	18	0
2:	2	6	11	18	25	34	43	55	67	82	100

Nearest matches based on these 2-step loop statistics, showing the experienced food density as a function of clump scale:

<i>Clump scale:</i>	1	2	3	4	5	7	10	15	20	25
<i>Exp. Density:</i>	5%	50%	70%	80%	80%	90%	90%	90%	100%	100%

Table 1: Comparisons of the effects of clump scale and maximum food location density on experienced food location density.




Figure 13: Percentage of mobile minimats (Movers) who move in loops of lengths 2, 4, or 64 or without looping for clump scales from 1 to 10, with maximum food location density 5% and food regrowth delay 0. Bigger food clumps do not lead to more long loops, in contrast to increased food location density (Fig. 8), indicating that these two parameters have different behavioral effects.

the maximum food location density used in Figure 13 is 5%, any (nonstarving) forward-moving minimats would experience many fewer food locations than its (nonstarving) 2- or 4-location looping counterparts. Thus this big-loop behavioral strategy is a losing proposition in patchy environments, and minimats instead evolve to stay local and ignore global statistics.

6.5.4 So clump scale is an independent environmental parameter

We must therefore now conclude that clump scale *does* make a difference to minimats. To immobile individuals, and to those traversing 2- and 4-location loops, clump scale is reducible to changes in the experienced apparent food location density. But to 64-location loopers, changing the clump scale does *not* change the experienced food density nearly as much. Thus, because clump scale has a differential effect on differently behaving minimats, it represents a unique selective force in their evolution, and its effect cannot be reduced to that of food density (nor regrowth delay) alone.

This is a surprising result, that a higher-order environmental structure such as patchiness should matter to creatures without sensors or memory to keep track of this structure. But it comes about because, in some sense, the behaviors themselves of our minimats have some higher-order structure: Minimats can be immobile or move only locally, or they can travel across the entire world. Thus there are two broad spatial scales at which minimats can evolve to behave. Changing clump scale has different effects at these two different spatial scales, greatly increasing the experienced food location density at the small scale, while altering it only insignificantly at the large scale. Changing overall food density in a random, unclumped environment, on the other hand, has the same overall effect at both of these two spatial scales. So the structure inherent in minimat behavioral possibilities will be matched to the structural changes induced by these two environmental parameters in different ways. (This also implies that if minimat behavior did *not* allow large-scale spatial travel, then environmental patchiness would not have a different effect from apparent food density on evolved behavior.)

From this point, once we have abandoned the notion that *no* higher-order environmental structure will matter to minimats, it quickly becomes clear that some kinds of higher-order structure will make a difference in minimat evolution, and some kinds will not. For instance, if the environment were structured fractally, changes in overall food density would have the *same* effect on experienced density at all spatial scales, by definition. So we would not expect fractal organization to make a difference to minimats; their evolution will be “blind” to this structure. On the other hand, it is easy to construct a highly structured environment with food placed in a regular pattern that will strongly select certain behavior patterns, but in a different way from clump scale. In the environment shown in Figure 14, nonstarving 64-location loopers and 2-location loopers would do well—experiencing 50%

by 64-location loopers goes up fairly linearly with clump scale—but is still much lower than that experienced by 2- and 4-location loopers. For instance, at clump scale 7, when 2-location loopers are experiencing an apparent 90% food location density, 64-location loopers experience only an apparent 14% density.

food density—because of their linear movement patterns. In contrast, 4-location loopers would fare poorly—experiencing 25% food density—because of their two-dimensional movement pattern. Hence we could devise another structure parameter based on regular food spacing that would also be important for characterizing minimat environments. Clearly our search for the complete set of important environment structure parameters in minimat worlds is not yet over.

7 What We Learned

We have come full circle: We argued initially that we have to take the organism’s perspective in terms of their behavioral capabilities to figure out what their general behavioral strategies might be, before we can study how the environment will shape those particular strategies. But now we have found that we cannot even define the environment per se irrespective of the organism’s perspective—we have to consider their behaviors to know what their environment is. Contrary to Simon’s statement quoted at the beginning of Section 1.2, it is not sufficient to know the organism’s behavioral capacities, sensory abilities, and goals; we must know how they actually put these together into the behaviors performed during their lifetimes, to construct their own individually-experienced environments. Thus, just like our moving minimats, we are in a loop: Evolved behaviors are affected by the environment, but behaviors determine that environment and its selection effects as well. We must consider both environment and behavior simultaneously in an interlocked fashion.

This turns our approach to characterizing evolved behavior in terms of environment structure on its head. The environment is not a fixed entity, predetermined and mandated for the minimats (or any organism) to respond to; depending on the behaviorally generated viewpoint, the environment to which an organism or population must adapt can assume a variety of sizes, shapes, and statistical properties. Thus, we should not begin by designing environments with global parameters and then see what behaviors evolve in response. We need to look first—or at least simultaneously—at the behaviors that are possible and see how different behavioral strategies will constrain the experienced environment. We must then compare these behavioral strategies in terms of how well-adapted they are to their *particular* self-constructed environments, not how well adapted they are to the global environment or that constructed by some other behavior pattern.

We are arguing here that organisms construct their own environment through what they “choose” to experience as a consequence of their behaviors. Note that it is also true that organisms construct their environment in a more direct fashion as well, by affecting the actual distribution of resources and other organisms around themselves. We see clear evidence of this in the minimat world in Figure 10, showing how food density decreases through the effect of minimats eating. Bedau (1996, in this issue) discusses the importance of organisms “pushing around” other organisms and thereby in some sense constructing their “social” environment, and Kirsh (1996, in this issue) compares a variety of methods by which organisms can change their environment to be more accommodating to their abilities. The issue of environment construction is central to Godfrey-Smith’s discussion of mind-




Figure 14: A regular pattern of food-bearing locations that will select for minimats moving in 2- and 64-location loops, who will experience fertile locations in 50% of the locations they traverse. Minimats moving in 4-location loops will only experience 25% fertile locations, and so will be selected against. This contrasts with both clumpy environments (Fig. 13) and environments with randomly located food (Fig. 8).

environment interactions (1996, in this issue), leading him to distinguish between narrow constructionist viewpoints where the organism causes direct physical change in the external environment, and a broader definition where changes in the organism's sensory and mental capacities make new things in the environment take on importance for the organism. But even this broader definition of organism-mediated environment construction, akin to Simon's (1956) discussed earlier, does not fully reveal the importance of considering indirect, behaviorally selected environment construction as a key step in analyzing behavior-environment interactions.

The fact that even sensory-less, memory-less minimats can construct their environments through behavioral selection only strengthens the importance of considering this type of environment-behavior feedback. If minimats, who can experience nothing, can still select the environment that they must adapt to, then organisms that *can* sense their world will have all the more occasion to choose their own environment structures to deal with (see, e.g., the "Just Say No" principle of Kirsh, 1996, in this issue). This indicates that a fruitful approach for designers of artificial agents in particular environments (and analyzers of real animals in their environments) is first to consider the range of appropriate behavioral strategies, then to analyze the experienced environments that each strategy would select, and finally to design sensors that will add the most useful information about those environments to guide the agent's behavior (or for real animals, analyze their sensory abilities in light of their behaviors and selected environments). For instance, adding expensive, full-blown vision to an artificial agent (or a primitive real organism) might not give it much more adaptive advantage than a simple range-finder (cf. Smithers, 1994, on the advantages of simple sensors for simple robot behaviors, and Arbib & Cobas, 1991, on the simple sensory needs of prey-capturing frogs and robots).

So is our whole research program of looking for environmental effects on evolved behaviors now hopelessly undermined? Any changes we make to what we perceive as the environment might not actually be changes in what the minimats (or other organisms under study) select as their environment, so how can we tell what is going on? In fact, the behavior-environment selection-adaptation loop might actually make our job easier: This feedback process might make behavior *less* sensitive to changes in environment, so that we have less to look for. As we saw in Figures refigure4 b and 5 b, particular patterns of minimat moving and splitting seem to be rather insensitive to a wide range of environmental conditions. Slightly closer to home, paleontologist Donald Prothero has said, "We've oversold the idea that animals, especially land mammals, are responsive to environmental change. Animals seem to be remarkably resistant to a lot more change than we thought" (quoted in Kerr, 1996, p. 432). Instead, the behavior-environment feedback loop may create a set of attractor states into which behavioral strategies tend to fall, and from which it is difficult to dislodge them, even as the environment changes.

This is the next step in our environment structure research program: to analyze behavior-environment feedback in terms of its effects on behavioral evolution. We hope to find evidence for behavioral fixed points and basins of attraction that yield a predictable kind of "inertia" to the evolutionary process of adapting behaviors to particular environments. To do so, we will once again take the perspective of the minimat (or animat) in its environment, and then, when in the minimat world, do as the minimats do, to discover the nature of the environments they select for themselves, and how those self-constructed environments exert selective pressures on their behaviors in turn.

Acknowledgments

Much of this research was performed in the resource-dense, minimal-delay environment of Stewart Wilson's Adaptive Animat Research Group (AARG) at the Rowland Institute for Science, Cambridge, MA. Our thanks to him, and to Anil Somayaji, Niko Kriegeskorte, Greg Werner, Oded Maron, Anita and Imogen Todd, and the members of the Center for Adaptive Behavior and Cognition for help with this research and manuscript. And in memory of soggy-fibers—we were proud to be the last users of one of the first thinking machines to be created.

References

- Ackley, D., & Littman, M. (1992). Interactions between learning and evolution. In C. G. Langton, C. Taylor, J. D. Farmer, & S. Rasmussen (Eds.), *Artificial life II* (pp. 487-509). Reading, MA: Addison-Wesley.
- Ackley, D., & Littman, M. (1994). Altruism in the evolution of communication. In R. A. Brooks & P. Maes (Eds.), *Artificial Life IV: Proceedings of the Fourth International Workshop on the Synthesis and Simulation of Living Systems* (pp. 40-48). Cambridge, MA: MIT Press/Bradford Books.
- Arbib, M. A., & Cobas, A. (1991). Schemas for prey-catching in frog and toad. In J.-A. Meyer & S. W. Wilson (Eds.), *From animals to animats: Proceedings of the First International Conference on Simulation of Adaptive Behavior* (pp. 142-151). Cambridge, MA: MIT Press/Bradford Books.
- Barbour, M. G., Burk, J. H., & Pitts, W. D. (1987). *Terrestrial plant ecology*, (2nd ed.). Menlo Park, CA: Benjamin/Cummings.
- Bedau, M. A. (1996). The extent to which organisms construct their environments. *Adaptive Behavior*, 4(3-4) (this issue).
- Bedau, M. A., & Packard, N. H. (1992). Measurement of evolutionary activity, teleology, and life. In C. G. Langton, C. Taylor, J. D. Farmer, & S. Rasmussen (Eds.), *Artificial life II* (pp. 431-461). Reading, MA: Addison-Wesley.
- Bell, W. J. (1991). *Searching behavior: The behavioural ecology of finding resources*. London: Chapman and Hall.
- Blythe, P. W., Miller, G. F., & Todd, P. M. (1996). Human simulation of adaptive behavior: Interactive studies of pursuit, evasion, courtship, fighting, and play. In M. Mataric, J.-A. Meyer, J. Pollack, & S. W. Wilson (Eds.), *From animals to animats 4: Proceedings of the Fourth International Conference on Simulation of Adaptive Behavior*. Cambridge, MA: MIT Press/Bradford Books.
- Cecconi, F., & Parisi, D. (1993). Neural networks with motivational units. In J.-A. Meyer, H. L. Roitblat, & 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.
- Cliff, D., & Bullock, S. (1993). Adding "foveal vision" to Wilson's animat. *Adaptive Behavior*, 2(1), 49-72.
- Cliff, D., Harvey, I., & Husbands, P. (1993). Explorations in evolutionary robotics. *Adaptive Behavior*, 2(1), 73-110.
- Cliff, D., & Ross, S. (1994). Adding temporary memory to ZCS. *Adaptive Behavior*, 3(2), 101-150.
- Collins, R. J., & Jefferson, D. R. (1992). AntFarm: Towards simulated evolution. In C. G. Langton, C. Taylor, J. D. Farmer, & S. Rasmussen (Eds.), *Artificial life II* (pp. 579-601). Reading, MA: Addison-Wesley.
- Erdmann, M. E., & Mason, M. T. (1988). An exploration of sensorless manipulation. *IEEE Journal of Robotics and Automation*, 4, 369-379.
- Fletcher, J. A., Zwick, M., & Bedau, M. A. (1996). Dependence of adaptability on environmental structure in a simple evolutionary model. *Adaptive Behavior*, 4(3-4) (this issue).
- Floreano, D. (1993a). Emergence of nest-based foraging strategies in ecosystems of neural networks. In J.-A. Meyer, H. L. Roitblat, & 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, May). Patterns of interactions in shared environments. In *Pre-proceedings of the Second European Conference on Artificial Life*, Brussels, Belgium.
- Godfrey-Smith, P. (1996). Precipitous of *Complexity and the function of mind in nature*. *Adaptive Behavior*, 4(3-4) (this issue).
- Gracias, N., Pereira, H., Lima, J. A., & Rosa, A. (1996, May). Gaia: An artificial life environment for ecological systems simulation. In *ALIFE V oral presentations* (preproceedings of the Artificial Life V workshop, Nara, Japan, pp. 105-112).
- Grant, W. (1992, June). *Artificial life worlds as discovery environments for learning*. Talk presented at the Artificial Life III Workshop, Santa Fe, NM.
- Hartvigsen, G., & Starmer, W. T. (1995). Plant-herbivore coevolution in a spatially and genetically explicit model. *Artificial Life*, 2(2), 239-259.
- Hendriks-Jansen, H. (1996). *Catching ourselves in the act: Situated activity, interactive emergence, evolution, and human thought*. Cambridge, MA: MIT Press/Bradford Books.
- Jefferson, D., Collins, R., Cooper, C., Dyer, M., Flowers, M., Korf, R., Taylor, C., & Wang, A. (1992). Evolution as a theme in artificial life: The Genesys/Tracker system. In C. G. Langton, C. Taylor, J. D. Farmer, & S. Rasmussen (Eds.), *Artificial life II* (pp. 549-578). Reading, MA: Addison-Wesley.
- Kerr, R. A. (1996). New mammal data challenge evolutionary pulse theory. *Science*, 273, 431-432.
- Kirsh, D. (1996). Adapting the environment instead of oneself. *Adaptive Behavior*, 4(3-4) (this issue).
- Koza, J. R. (1992). Genetic evolution and co-evolution of computer programs. In C. G. Langton, C. Taylor, J. D. Farmer, & 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, & S. W. Wilson (Eds.), *From animals to animats 2: Proceedings of the Second International Conference on Simulation of Adaptive Behavior* (pp. 262-270). Cambridge, MA: MIT Press/Bradford Books.
- Littman, M. L. (1994). Memoryless policies: Theoretical limitations and practical results. In D. Cliff, P. Husbands, J.-A. Meyer, & S. W. Wilson (Eds.), *From animals to animats 3: Proceedings of the Third International Conference on Simulation of Adaptive Behavior* (pp. 238-245). Cambridge, MA:

MIT Press/Bradford Books.

Lund, H. H., & Parisi, D. (1995). Pre-adaptation in populations of neural networks evolving in a changing environment. *Artificial Life*, 2(2), 179-197.

Lund, H. H., & Parisi, D. (1996, May). Generalist and specialist behavior due to individual energy extracting abilities. In *ALIFE V oral presentations* (preproceedings of the Artificial Life V workshop, Nara, Japan, pp. 294-304).

Maes, P. (1994). Modeling adaptive autonomous agents. *Artificial Life*, 1(1-2), 135-162.

Maley, C. (1996, May). Simulating invasions in a spatially structured ecosystem. In *ALIFE V poster presentations* (preproceedings of the Artificial Life V workshop, Nara, Japan, pp. 205-211).

Mason, M. T. (1993). Kicking the sensing habit. *AI Magazine*, 1, 58-59.

Menczer, F., & Belew, R. K. (1994). Evolving sensors in environments of controlled complexity. In R. A. Brooks & P. Maes (Eds.), *Artificial Life IV: Proceedings of the Fourth International Workshop on the Synthesis and Simulation of Living Systems* (pp. 210-221). Cambridge, MA: MIT Press/Bradford Books.

Menczer, F., & Belew, R. K. (1996a). Latent energy environments. In R. K. Belew & M. Mitchell (Eds.), *Adaptive individuals in evolving populations: Models and algorithms*. Reading, MA: Addison-Wesley.

Menczer, F., & Belew, R. K. (1996b). From complex environments to complex behaviors. *Adaptive Behavior*, 4(3-4) (this issue).

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., & Cecconi, F. (1992). Learning, behavior, and evolution. In F. J. Varela & 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.

Resnick, M. (1994). *Turtles, termites, and traffic jams: Explorations in massively parallel microworlds*. Cambridge, MA: MIT Press/Bradford Books.

Shepard, R. N. (1987). Evolution of a mesh between principles of the mind and regularities of the world. In J. Dupre (Ed.), *The latest on the best: Essays on evolution and optimality* (pp. 251-275). Cambridge, MA: MIT Press/Bradford Books.

Simon, H. A. (1956). Rational choice and the structure of the environment. *Psychological Review*, 63(2), 129-138.

Smithers, T. (1994). On why better robots make it harder. In D. Cliff, P. Husbands, J.-A. Meyer, & S. W. Wilson (Eds.), *From animals to animats 3: Proceedings of the Third International Conference on Simulation of Adaptive Behavior* (pp. 64-72). Cambridge, MA: MIT Press/Bradford Books.

Stephens, D. W., & Krebs, J. R. (1986). *Foraging theory*. Princeton, NJ: Princeton University Press.

Taylor, C. E., Jefferson, D. R., Turner, S. R., & 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, May). Artificial death. In *Preproceedings of the Second European Conference on Artificial Life*, Brussels, Belgium.

Todd, P. M., & Miller, G. F. (1991a). Exploring adaptive agency II: Simulating the evolution of associative learning. In J.-A. Meyer & 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., & Miller, G. F. (1991b). Exploring adaptive agency III: Simulating the evolution of habituation and sensitization. In H.-P. Schwefel & R. Maenner (Eds.), *Proceedings of the First International Conference on Parallel Problem Solving from Nature* (pp. 307-313). Berlin: Springer-Verlag.

Todd, P. M., & Wilson, S. W. (1993). Environment structure and adaptive behavior from the ground up. In J.-A. Meyer, H. L. Roitblat, & 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.

Todd, P. M., Wilson, S. W., Somayaji, A. B., & Yanco, H. A. (1994). The blind breeding the blind: Adaptive behavior without looking. In D. Cliff, P. Husbands, J.-A. Meyer, & S. W. Wilson (Eds.), *From animals to animats 3: Proceedings of the Third International Conference on Simulation of Adaptive Behavior* (pp. 228-237). Cambridge, MA: MIT Press/Bradford Books.

Toquenaga, Y., Kajitani, I., & Hoshino, T. (1994). Egrets of a feather flock together. In R. A. Brooks & P. Maes (Eds.), *Artificial Life IV: Proceedings of the Fourth International Workshop on the Synthesis and Simulation of Living Systems* (pp. 140-151). Cambridge, MA: MIT Press/Bradford Books.

Werner, G. M., & Dyer, M. G. (1992). Evolution of communication in artificial organisms. In C. G. Langton, C. Taylor, J. D. Farmer, & S. Rasmussen (Eds.), *Artificial life II* (pp. 659-687). Reading, MA: Addison-Wesley.

Werner, G. M., & Dyer, M. G. (1993). Evolution of herding behavior in artificial animals. In J.-A. Meyer, H. L. Roitblat, & 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 & 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.

Wilson, S. W. (1994). ZCS: A zeroth level classifier system. *Evolutionary Computation*, 2(1), 1-18.

Yamauchi, B. M., & Beer, R. D. (1994). Sequential behavior and learning in evolved dynamical neural networks. *Adaptive Behavior*, 2(3), 219-246.

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.