

# Dynamics of the Environment for Adaptation in Static Resource Models

Mark A. Bedau

Reed College, 3203 SE Woodstock Blvd., Portland OR 97202, USA  
mab@reed.edu, <http://www.reed.edu/~mab>

**Abstract.** We measure the environment that is relevant to a population's adaptation as the information-theoretic uncertainty of the distribution of local environmental states that the adapting population experiences. Then we observe the dynamics of this quantity in simple models of sensory-motor evolution, in which an evolving population of agents live, reproduce, and die in a two-dimensional world while competing for resources. Although the distribution of resources is static, the agents' evolution creates a dynamic environment for adaptation.

## 1 The Environment for Adaptation

The process of adaptation is driven by the environment, but it also shapes the environment. If we focus on the environment for adaptation, that is, the aspect of the environment that is relevant to adaptive evolution, we would tend to agree with Lewontin [7] that organisms construct their environment:

The environment is not a structure imposed on living beings from the outside but is in fact a creation of those beings. The environment is not an autonomous process but a reflection of the biology of the species.

But to think clearly about how organisms and their environment affect one another, we must first clarify what the environment for adaptation is. How can it be measured? How is it affected by the behavior of the evolving population? This paper takes a first step at answering these questions.

Part of an answer to these questions involves distinguishing two kinds of local environmental states: those that exist in the environment and so could conceivably be relevant to the agents' adaptation, and those that the agents experience and so that affect their adaptation. To appreciate how evolution is both cause and effect of the environment, one must focus on the latter. In this paper we propose how to measure the environment for adaptation, and we observe its dynamics in evolving systems with fixed resource distributions—what we call *static resource models*. We see that even simple systems with static resource distributions have a dynamic environment for adaptation.

Classifying and quantifying the environment for adaptation is the subject of a variety of recent work in artificial life. The relevance of this work includes both

its implications for theoretical insights into the evolutionary process and its practical consequences for engineering and controlling evolving systems. Wilson [14] and Littman [8] provide an abstract categorization of environments, with Wilson focusing on the degree of non-determinism in an environment and Littman characterizing the simplest agent that could optimally exploit an environment. But neither method provides a dynamic quantitative measure of the environment the population experiences. Other recent work experimentally investigates agents adapting in different environments. Todd and Wilson [12] introduce an experimental framework for investigating how adaptation varies in response to different kinds of environments, and Todd et al. [13] demonstrate different adaptations in different kinds of environments. In neither case, though, is environmental structure actually classified or measured. The present project grows out of previous work on how adaptability depends on quantitative measures of environmental structure [2, 4, 5]. What is novel here is the focus on the environment actually experienced, which allows the measures to be dynamic where the previous ones were static. While Lewontin [7] verbally stresses the centrality of the environment for adaptation, being able to quantify and observe it provides the empirical footing needed to respond to skeptics [6].

## 2 Adaptation in Packard’s Static Resource Models

All of our empirical observations are from computer simulations of a certain model—originated by Norman Packard [9, 3]—that is designed to be a very simple model of the evolution of sensory-motor strategies. Packard’s model consists of agents sensing the resources in their local environment, moving as a function of what they sense, ingesting the resources they find, and reproducing or dying as a function of their internal resource levels. The model’s spatial structure is a grid of sites with periodic boundary conditions, i.e., a toroidal lattice. Resources are immediately replenished at a site whenever they are consumed. The agents constantly extract resources and expend them by living and reproducing. Agents ingest all of the resources (if any) found at their current location and store them internally. Agents expend resources at each time step by “paying” (constant) “existence taxes” and “movement taxes” (variable, proportional to distance moved). If an agent’s internal resource supply drops to zero it dies.

Each agent moves each time step as dictated by its genetically encoded sensory-motor map: a table of behavior rules of the form IF (environment  $j$  sensed) THEN (do behavior  $k$ ). An agent receives sensory information about the existence of resources (but not the other agents) in the von Neumann neighborhood of five sites centered on its present location in the lattice. Thus, each sensory state  $j$  corresponds to one of  $2^5 = 32$  different detectable local environments. Each behavior  $k$  is a jump vector between one and fifteen sites in any one of the eight compass directions, or it is a random walk to the first unoccupied site, so an agent’s genome continually instructs it to move somewhere. Thus, an agent’s genotype, i.e., its sensory-motor map, is just a lookup table of sensory-motor rules. But the space in which adaptation occurs is fairly large, consisting

of  $121^{32} \approx 10^{66}$  distinct possible genotypes. In a resource field missing some of von Neumann neighborhoods, the number of *effectively* different genotypes will be smaller (about  $10^{29}$  for the resource fields studied here).

An agent reproduces (asexually, without recombination) if its resource reservoir exceeds a certain threshold. The parent produces one child, which starts life with half of its parent's resource supply. The child also inherits its parent's sensory-motor map, except that point mutations can replace the behaviors linked to some sensory states with randomly chosen behaviors (there is roughly one mutation per reproduction event, on average). A time step in the simulation cycles through the entire population and has each agent, in turn, complete the following sequence of events: sense its present von Neumann neighborhood, move to the new location dictated by its sensory-motor map (if that site is already occupied, it randomly walks to the first unoccupied site), consume any resources found at its new location, expend resources to cover existence and movement taxes, and then, if its resource reservoir is high enough or empty, either reproduce or die. A given simulation starts with randomly distributed agents containing randomly chosen sensory-motor strategies. The model contains no *a priori* fitness function, as Packard [9] has emphasized. Agents with maladaptive strategies find few resources and thus to die, taking their sensory-motor genes with them; by contrast, agents with adaptive strategies tend to find sufficient resources to reproduce, spreading their strategies (with mutations) through the population.

Here we restrict resources to blocks: a square grid of lattice sites which all have the same fixed quantity of resources. In a resource distribution formed of blocks, every site that is not inside a block is part of a resource-free desert. We focus on two distributions, one with many small blocks randomly sprinkled across space, the other with one large block. If the resource distribution contains just one block, then all evolution concerns strategies for exploiting the resources on the block. But if the resource distribution contains many blocks scattered in space, then each block is an evolutionary island supporting its own evolutionary development and subject to migration from nearby islands. So a behavioral innovation that originates on one island can hop from island to island and eventually colonize the entire archipelago. Resource blocks come in different sizes (widths). If the block is small enough (width  $\leq 3$ ) then the agent's sensory information always unambiguously indicates the agent's exact location on the block (NW corner, middle, N edge, etc.). As the width increases above 3, so does the ambiguity of the agents' sensory information [4]. Agents cannot always tell exactly where they are, for the sites inside the block all look alike, as do the sites at its edge. Here we studied two size blocks: one ( $3 \times 3$ ) that is unambiguous, and another ( $30 \times 30$ ) with rampant ambiguity.

Agents on a given  $3 \times 3$  block usually all follow the same strategy, for agents following heterogeneous strategies tend to collide and be bumped into the resource desert. The strategy observed are cycles jumping through a sequence of sites on the block. The simplest cycles (period 2) consist of jumping back and forth between two sites. Since a  $3 \times 3$  block contains 9 distinct sites, it can support at most a period 9 strategy. A period  $n$  strategy has room for at most

$n - 1$  agents (one agent in the cycle must move first and the space to which it is jumping must be unoccupied). Thus, longer period strategies can support larger populations because they can exploit more of the energetic resources on a block. All agents reproduce at the same rate, so a block with a larger population will produce offspring at a higher rate. Thus, blocks with larger period strategies will exert greater migration pressure and thus will have a selective advantage. So evolution in an archipelago of tiny  $3 \times 3$  resource islands will exhibit one main kind of adaptive event: lengthening the period of an existing strategy. Agents on one large  $30 \times 30$  block tend to exhibit a fluidly changing ecosystem of coexisting strategies. These strategies fall into two main categories: Edge strategies and the Random strategy. The Random strategy relies on one central gene: if inside the middle of a block (i.e., if sensing the neighborhood with resources at all five sites), do a random walk to the first unoccupied site. (Recall that one of the possible genetically encoded behaviors  $k$  is a random walk.) Although unsuccessful on a  $3 \times 3$  block, a  $30 \times 30$  block is large enough for the Random strategy to succeed. An Edge strategy consists of moving in a straight line when inside the block and jumping back into the middle of the block when detecting the edge. A proliferating Edge strategy will fill a region of the block with a perpetually rolling population.  $30 \times 30$  blocks have three main kinds of adaptive events: lengthening an Edge strategy's jump back, discovering a new Edge strategy compatible with the strategy ecology, and discovering the Random strategy.

### 3 A Measure of the Environment for Adaptation

We are interested in how best to measure the environment that is relevant to a population's adaptation. For the sake of concreteness, we will develop and apply our methods in the context of the static resource models described above. The quantities we define and observe fit within a family of measures that has been developed and studied in similar contexts previously [2, 4, 5].

The basic idea behind our measure of the environment for adaptation is quite simple. The first step is to define a partition of relevant local environments. Of all the environmental states that exist or could be defined in an evolving system, first focus on those that are *local* to the individual agents in the system in the sense that different agents can be in different local environmental states at the same time. Next limit attention to just those that are *relevant* to the survival or reproducibility of the agents, i.e., those that can be expected to be relevant to the population's adaptation. This includes only those states that can affect or be affected by the agents. Finally, consider a *partition* of local relevant environments—a set of states that is disjoint and exhaustive in the sense that at each time each individual agent is unambiguously in exactly one of the states. For example, in Packard's model studied here, a natural partition of local relevant environments is the set of different von Neumann neighborhoods that the agents can distinguish (the presence or absence of resources in each of the five detectable neighboring sites). At each time each agent is in exactly one of these local states. The states affect the agents by being their sensory

input, and the agents affect the states through their behavior [2]. The states are relevant to the agents' adaptation, for the agents' adaptive task is to evolve a genome that associates each those states with an appropriate behavioral rule in a sensory-motor strategy and natural selection shapes only the rules for those von Neumann neighborhoods that are actually experienced by the population.

During the course of evolution, the agents experience states in the partition of local relevant environments with different frequencies. The frequency (probability) distribution of the states is a straightforward reflection of the environment that is currently relevant for the population's adaptation. That is, let  $\{e_i\}$  be the partition of local relevant environments, and let  $P^E(e_i)$  be the frequency with which agents in the population experience the  $i^{\text{th}}$  local state  $e_i$  (during some time window) in the global environment  $E$ . Then the probability distribution  $P^E$  reflects the nature of the environment that is *currently relevant* to the population's adaptation. It is important to note that the distribution  $P^E$  is defined relative to the partition  $\{e_i\}$ . Different partitions will yield different probability distributions, so the usefulness and interest in the distribution  $P^E$  hinges on the choice of partition  $\{e_i\}$ .<sup>1</sup> In Packard's model where  $\{e_i\}$  is the von Neumann neighborhoods,  $P^E$  shows the frequencies with which the population visits different neighborhoods.

The final step in our measure of the environment for adaptation is to quantify the variety in  $P^E$ , and information-theoretic uncertainty (Shannon entropy) [11] is the natural tool to use:

$$H(P^E) = - \sum_i P^E(e_i) \log_2 P^E(e_i). \quad (1)$$

$H(P^E)$  reflects two aspects of  $P^E$ : its width (number of different neighborhood indices  $i$ ) and flatness (constancy of the value of  $P^E(e_i)$  for different  $i$ ). The wider and flatter  $P^E$  is, the more uncertainty there is about which neighborhood the agents will experience and the higher  $H(P^E)$  will be.

Even in static resource models,  $P^E$  and so  $H(P^E)$  can change, for  $P^E$  depends on the adapting population's behavior. As the population's behavior evolves, the population may experience different local environments, so  $P^E$  and thus  $H(P^E)$  may be dynamic. The dynamics of  $H(P^E)$  simultaneously reflects both cause and effect of the adapting population. That is,  $H(P^E)$  measures both the environment *created by* the population's *past* adaptation and the environment *relevant for* the population's *future* adaptation, as reflected in the local environmental states the population experiences. The population's adaptation to its environment changes the frequency with which it experiences different local environments, and this creates a self-organizing dynamic in  $H(P^E)$ . (Parisi,

---

<sup>1</sup> Since the distribution  $P^E$  pools information about the local relevant environments experienced by all the agents, it in effect *averages* this information across the population. If different subsets of agents experience substantially different subsets of local environments (e.g., if they exist in different niches), the averaging in  $P^E$  will obscure this difference. To counteract this, one could collect different probability distributions for different subsets of agents.

Nolfi, and Cecconi [10] and Bedau [2] study this issue.) The dynamics of  $H(P^E)$  show how adaptation changes the environment for subsequent adaptation.


To get a feel for how  $P^E$  and  $H(P^E)$  reflect the environment for adaptation, assume the environment consists of an archipelago of randomly scattered  $3 \times 3$  resource blocks in Packard's model. Since prolonged stay in the resource desert is lethal, the population will overwhelmingly experience some subset of the nine neighborhoods on the block. For example, if the populations on the different resource islands are all exhibiting the same period two cyclic strategy, then the bulk in  $P^{3 \times 3}$  is evenly divided between the two neighborhoods visited (agents bumped off out of the cycle will visit a few other local environments before perishing). As evolution increases the length of the strategy followed by the agents, the population will encounter new environmental states and  $P^{3 \times 3}$  will stay flat but become wider. Thus the environment for adaptation,  $H(P^{3 \times 3})$ , will tend to increase. The magnitude of this trend is bounded by the number of different resource sites on the tiny block; with nine distinct resource sites,  $H(P^{3 \times 3}) \leq \log_2 9 \approx 3.17$ .

Now consider an environment containing only one large  $30 \times 30$  block. Whether the block is populated by an ecology of coexisting Edge strategies or dominated by the Random strategy, most of the population will be inside the block, so the bulk of  $P^{30 \times 30}$  will reflect that one neighborhood with most of the rest shared among various edge neighborhoods (plus a small fraction for the agents bumped into the desert). The dynamics of  $H(P^{30 \times 30})$  can be varied. Lengthening the inside jump of an Edge strategy increases the relative frequency with which the population encounters the neighborhood inside the block, and this makes  $P^{30 \times 30}$  less flat and thus shrinks  $H(P^{30 \times 30})$ . Discovering a new Edge strategy raises the number of sites experienced in the ecology of strategies, so  $P^{30 \times 30}$  becomes wider and  $H(P^{30 \times 30})$  rises. Discovering the Random strategy typically increases the frequency with which the population encounters the inside the block, so  $P^{30 \times 30}$  narrows and  $H(P^{30 \times 30})$  falls. Thus, the environment for adaptation on a  $30 \times 30$  block should both rise and fall, but the eventual evolution of the Random strategy will ultimately drive it to a low value.

## 4 Observations of the Environment for Adaptation

We observed  $H(P^E)$  in 100  $3 \times 3$  blocks scattered randomly in space, and in one  $30 \times 30$  block. Both resource distributions pump resources into the environment at the same rate, so they can support comparable maximum population sizes. In this resource-driven and space-limited model, population size is a good reflection of the dynamics of adaptation (fitness). Except for varying the resource distribution, all the parameters in the models observed were the same.

Evolution on an archipelago of  $3 \times 3$  blocks always exhibits a pattern of period-lengthening adaptations. The population starts with a small period cycle on one of the blocks. This strategy migrates across the archipelago, until it has colonized virtually all the blocks. Eventually the appropriate mutations will create a longer-period cycle on one of the blocks and the agent with this longer cycle



**Fig. 1.** Dynamics of population size and environment for adaptation,  $H(P^E)$ , in two typical simulations. The top two graphs are from an archipelago of 100  $3 \times 3$  blocks. The bottom two graphs are from one  $30 \times 30$  block.

will sometimes supplant the agents on that block with the shorter cycle, so the block is now populated by a tiny subpopulation with the longer cycle. Children from this subpopulation migrate to other blocks and supplant the shorter-cycle agents there, until they eventually colonize the entire archipelago. This sequence is repeated again and again with ever longer-period cycles, until the cycle length gets close to its maximum possible value (nine).

The top two graphs in Figure 1 show time series of population level and the environment for adaptation,  $H(P^E)$ , in a typical simulation in an archipelago of  $3 \times 3$  blocks. The course of evolution has four epochs. Epoch I starts with a period-2 cyclic strategy on a single block. As this strategy reproduces and starts to migrate to other blocks, the population rises, until about all the blocks are colonized. Epoch II starts at about time 75,000 when a period-lengthening mutation enables a period-3 cycle to become established on one of the blocks. The 3-cycle then invades the rest of the archipelago. Epoch III starts at around time 150,000 when a coordinated combination of mutations create a 5-cycle on one block, which then invades the rest of the archipelago. Another period-lengthening mutation on one block initiates Epoch IV with a 6-cycle at around time 200,000. This strategy then colonizes the rest of the archipelago. These period-lengthening adaptations each cause a rise in  $H(P^E)$ , as the more complicated strategies make the experienced environment more complex.

Evolution on the  $30 \times 30$  block tends to exhibit its own distinctive pattern of adaptations. One or two edge strategies in the initial random seed population will live on the block, growing to fill the space available in the niche on the block used by those strategies. From time to time new strategies arise by mutation, but most die off quickly. Some find an unoccupied niche on the block and coexist with their ancestral strategies. Others compete for niche space with ancestral strategies and supplant their ancestors. In this way, the block supports a changing ecosystem of subpopulations surviving through different strategies in different regions on the block. Eventually, though, the Random strategy always gets a foothold in the population and drives all Edge strategies extinct.

The bottom two graphs in Figure 1 show data from a typical illustration of a co-evolving succession of Edge strategies ending with domination by the Random strategy. The course of evolution again has four epochs. In Epoch I, one strategy (call it B11) jumps off the bottom edge a half-dozen steps into the block. The population of agents following this strategy grows until at about time 4000 it fills the space along the entire bottom edge with a population size pushing 80. Except for momentary spikes (caused when a crop of children get bumped off the block, thus briefly experiencing the desert neighborhood before dying),  $H(P^E)$  is basically stable during this epoch. The entire population experiences two neighborhoods: the inside of the block and the bottom inside edge. Epoch II starts around time 6000 when a second strategy (R11) starts to coexist with B11 by occupying the block's right edge. This strategy lives only right along the edge, so only a handful of agents can occupy the niche created by this strategy.  $H(P^E)$  shows a slight rise as a subpopulation experiences two new neighborhoods (right inside edge and bottom right corner). Epoch III starts around time 9000 when a



third strategy (L11) shares the block with B11 and R11 by exploiting the hitherto unused left edge with a small jump into the block. L11 grows to fill its niche along the left edge, with a subpopulation of a couple dozen.  $H(P^E)$  rises again during this epoch as a new neighborhood (inside left edge) gets experienced more and more frequently. Epoch IV starts around 13000 when the Random strategy arises inside the block and it quickly sweeps across the block, knocking B11, R11, and L11 out of their niches and filling virtually all the space on the block.  $H(P^E)$  drops dramatically because the neighborhoods experienced by the population shows much less variety. As the population grows, more agents experience an edge or the desert, so  $H(P^E)$  rises slightly.

## 5 Conclusions

Two main conclusions follow from this work, one methodological and the other substantive. The methodological conclusion is that we now have a dynamic and quantitative reflection of the environment that the evolving population actually experiences—the environment for adaptation. This measure of the environment for adaptation is quite general and feasible, and it can straightforwardly be applied to many other evolving systems. Its drawback is that it depends on an *a priori* partition of the environmental states relevant to the population’s adaptive evolution. Such a partition is easy in the simple systems studied here but more complex systems can present difficulties, especially if qualitatively new kinds of environmental states emerge unpredictably in the course of evolution.

The substantive conclusion is that the environment for adaptation is *dynamic* even in simple static resource models like those studied here. Sometimes a population evolves a more complicated behavioral strategy which enables it to experience and exploit more of the environment’s complexity. Subsequent evolution then takes place in the context of this more complicated environment for adaptation. In other cases the population benefits by evolving a simpler behavioral strategy and thus simplifying its environment. An agent’s environment for adaptation includes the rest of the population with which the agent interacts. In static resource models, agent–agent interactions like collisions can be a significant part of the adaptive landscape, so the shape of the landscape depends on the collection of behavioral strategies in the population. Space is itself a resource [1], and available space is dynamic since it depends on the changing spatial location of the population with which agents interact. The dynamics of the environment for adaptation strongly depend on the actual course that evolution happens to take. Repeating a simulation with exactly the same resource distribution (data not shown here) can yield qualitatively different environmental dynamics, which shows the “fluidity” of these environmental dynamics. These conclusions pertain in the first instance just to the particular systems studied here, but they are likely to hold quite generally.

**Acknowledgements.** For helpful discussion or comments on the manuscript, thanks to Peter Godfrey-Smith, Norman Packard, Peter Todd, Marty Zwick,

and two anonymous reviewers. For assistance with this research, thanks to Matt Giger and John Hopson. Thanks to the Santa Fe Institute for hospitality and computational resources that supported some of this work.

## References

1. Abrams, P. (1992). Resource. In E. F. Keller and E. A. Lloyd, eds., *Keywords in Evolutionary Biology* (pp. 282–285). Cambridge, Mass: Harvard University Press.
2. Bedau, M. A. (1994). The Evolution of Sensorimotor Functionality. In P. Gaussier and J. -D. Nicoud, eds., *From Perception to Action* (pp. 134-145). New York: IEEE Press.
3. Bedau, M. A., Packard, N. H. (1992). Measurement of Evolutionary Activity, Teleology, and Life. In C. G. Langton, C. E. Taylor, J. D. Farmer, S. Rasmussen, eds., *Artificial Life II* (pp. 4431–461). Redwood City, Calif.: Addison-Wesley.
4. Fletcher, J. A., Zwick, M., Bedau, M. A. (1996). Dependence of Adaptability on Environment Structure in a Simple Evolutionary Model. *Adaptive Behavior* **4**, 283–315.
5. Fletcher, J. A., Bedau, M. A., Zwick, M. (1998). Effect of Environmental Structure on Evolutionary Adaptation. In C. Adami, R. K. Belew, H. Kitano, and C. E. Taylor, eds., *Artificial Life VI* (pp. 189–198). Cambridge, MA: MIT Press.
6. Godfrey-Smith, Peter. (1996). *Complexity and the Function of Mind in Nature*. Cambridge: Cambridge University Press.
7. Lewontin, R. C. (1983). The Organism as the Subject and Object of Evolution. *Scientia* **118**, 63–82. Reprinted in R. Levins and R. Lewontin, *The Dialectical Biologist* (pp. 85–106). Cambridge, MA: Harvard University Press.
8. Littman, M. L. (1993). An Optimization-Based Categorization of Reinforcement Learning Environments. In J. A. Meyer, H. L. Roiblat, and S. W. Wilson (Eds.), *From Animals to Animats 2* (pp. 262–270). Cambridge, MA: Bradford/MIT Press.
9. Packard, N. H. (1989). Intrinsic Adaptation in a Simple Model for Evolution. In C. G. Langton, ed., *Artificial Life* (pp. 141–155). Redwood City, Calif.: Addison-Wesley.
10. Parisi, D., Nolfi, S., Cecconi, F. (1992) Learning, Behavior, and Evolution. In F. Varela and P. Bourgine, eds., *Towards a Practice of Autonomous Systems* (pp. 207–216). Cambridge, Mass.: Bradford/MIT Press.
11. Shannon, C. E., Weaver, W. (1949). *The Mathematical Theory of Communication*. Urbana, Ill.: Univ. of Illinois Press.
12. Todd, P. M., Wilson, S. W. (1993). Environment Structure and Adaptive Behavior From the Ground Up. In J. A. Meyer, H. L. Roiblat, and S. W. Wilson (Eds.), *From Animals to Animats 2* (pp. 11–20). Cambridge, MA: Bradford/MIT Press.
13. 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, and S. W. Wilson (Eds.), *From Animals to Animats 3* (pp. 228–237). Cambridge, MA: Bradford/MIT Press.
14. Wilson, S. W. (1991). The Animat Path to AI. In J. -A. Meyer and S. W. Wilson (Eds.), *From Animals to Animats* (pp. 15–21). Cambridge, MA: Bradford/MIT Press.