

Portland State University

PDXScholar

Systems Science Faculty Publications and
Presentations

Systems Science

1998

Effect of Environmental Structure on Evolutionary Adaptation

Jeffrey Alan Fletcher

Portland State University, jeff@pdx.edu

Mark A. Bedau

Reed College

Martin Zwick

Portland State University, zwick@pdx.edu

Follow this and additional works at: https://pdxscholar.library.pdx.edu/sysc_fac

Let us know how access to this document benefits you.

Citation Details

Fletcher, Jeffrey Alan; Bedau, Mark A.; and Zwick, Martin, "Effect of Environmental Structure on Evolutionary Adaptation" (1998). *Systems Science Faculty Publications and Presentations*. 142. https://pdxscholar.library.pdx.edu/sysc_fac/142

This Conference Proceeding is brought to you for free and open access. It has been accepted for inclusion in Systems Science Faculty Publications and Presentations by an authorized administrator of PDXScholar. Please contact us if we can make this document more accessible: pdxscholar@pdx.edu.

Effect of Environmental Structure on Evolutionary Adaptation

Jeffrey A. Fletcher, Mark A. Bedau, Martin Zwick

Systems Science Ph.D. Program, Portland State University, Portland, Oregon 97207-0751, jeff@sysc.pdx.edu
Department of Philosophy, Reed College, 3203 SE Woodstock Boulevard, Portland, Oregon 97202, mab@reed.edu
Systems Science Ph.D. Program, Portland State University, Portland, Oregon 97207-0751, zwick@sysc.pdx.edu

Abstract

This paper investigates how environmental structure, given the innate properties of a population, affects the degree to which this population can adapt to the environment. The model we explore involves simple agents in a 2-d world which can sense a local food distribution and, as specified by their genomes, move to a new location and ingest the food there. Adaptation in this model consists of improving the genomic sensorimotor mapping so as to maximally exploit the environmental resources. We vary environmental structure to see its specific effect on adaptive success. In our investigation, two properties of environmental structure, conditioned by the sensorimotor capacities of the agents, have emerged as significant factors in determining adaptive success: (1) the information content of the environment which quantifies the diversity of conditions sensed, and (2) the expected utility for optimal action. These correspond to the syntactic and pragmatic aspects of environmental information, respectively. We find that the ratio of expected utility to information content predicts adaptive success measured by population gain and information content alone predicts the fraction of ideal utility achieved. These quantitative methods and specific conclusions should aid in understanding the effects of environmental structure on evolutionary adaptation in a wide range of evolving systems, both artificial and natural.

Adaptation as a Function of Environmental Structure

An evolving system consists of a population of agents adapting their behavior to an environment through the process of natural selection. The difficulty of the adaptive challenge obviously depends upon the population, the environment, and the interaction between the two. In this paper, we adopt an environment-centered view, that is, we examine how environments vary in the adaptive challenge which they present. This orientation reflects a kind of figure/ground reversal. One commonly takes the environment as ground and the adapting population as figure. That is, one treats the adaptive challenge as fixed and examines the resulting dynamics of adaptation, e.g., as a function of different adaptive capabilities of the population. Here, we treat the population as relatively given and study how varying the environment affects the difficulty of the adaptive task to be solved. This reversal of focus is found in some other recent studies (e.g., Wilson,

1991; Littman, 1993; Todd and Wilson, 1993; Todd et al. 1994, Todd and Yanco, 1996, and Menczer and Belew, 1996) and it recalls the earlier work of Emery and Trist (1965) on the causal texture of environments of social organizations. These studies tend to pursue one of two projects: either providing an abstract categorization of environments, or gathering experimental evidence about how artificial agents actually adapt in different simulated environments. Here and in a previous paper (Fletcher, Zwick, and Bedau 1996) we pursue both projects simultaneously; we experimentally study how the adaptation of given (possibly sub-optimal) agents varies in response to environmental structure. Since our characterization of environmental structure is quantitative, we can seek evidence for general laws relating adaptive success and environmental structure.

We intentionally have made our model quite simple. In this way we can more easily develop quantitative methods and results which can then be applied to more complicated evolving systems. In particular, our model encompasses the following simplifications:

- simple environment
 1. agents do not affect the environmental structure, which is static
 2. agents randomly sample the entire environment
- simple agents
 1. simple internal representation of the environment (implicit in the genome)
 2. no genotype/phenotype difference (every gene encodes a response to a unique sensory condition)
 3. simple behavior: movement (and food ingestion)
 4. no temporal organization of behavior
- simple evolutionary process
 1. minimal inter-agent interactions
 2. no sexual recombination

The following complex features of agent-environment interaction were, however, retained in our model:

- differences between objective environment and an agent's sensory discrimination
- uncertain consequences of action (an agent's sensory horizon is smaller than its movement horizon)
- no explicit fitness function

The population in our model consists of sensorimotor agents. Each agent responds to limited sensory input from the environment with a single behavioral output specified by the agent's genome. The adaptive task consists of finding an output to associate with each possible input. The difficulty of the adaptive task, therefore, involves at least the following aspects of environmental structure:

- the quantity of sensory information, i.e., the variety of sensed environmental conditions with which behaviors must be associated (a “syntactic” aspect)
- the utility of the information, i.e., the benefit of adaptive behaviors over non-adaptive behaviors (a “pragmatic” aspect)

The first draws upon the information theory of Shannon and Weaver (1949); the second draws upon game theory (sometime referred to as decision theory) of von Neumann and Morgenstern (1944). Both aspects are needed to characterize the evolutionary challenge. We refer to them jointly as *environmental structure*.

In terms of these aspects, an adaptive task is difficult if the environment sends many messages requiring an adaptive response, or if they have little utility. The syntactic aspect is central to Ashby's (1956) conceptualization of adaptation, according to which environmental variety poses a problem to which behavioral variety is the response. Agents also experience the second aspect of environmental structure directly as they gain the resources yielded by particular responses to particular sensory inputs. In our previous paper (1996) we began to explore these issues. Here we extend that work by quantitatively measuring both aspects of environmental structure and showing their relationship to our quantitative measures of adaptive success.

Modeling Adaptation in Diverse Environments

Our observations are from computer simulations of adaptation in a series of constructed environments. The model consists of many agents that sense their local environment, move as a function of what they sense, and ingest what resources they find where they move. This model is a modification of those previously studied by Bedau and Packard (1992), Bedau, Ronneburg and Zwick (1992), Bedau (1994), Bedau and Bahm (1994), Bedau (1995), Bedau, Giger and Zwick (1995), and Fletcher, Zwick, and Bedau (1996). All of these models are extensions of one originally developed by Packard (1989).

Agent and Environment Interactions

The world is a grid of 128 x 128 sites with periodic boundary conditions, i.e., a toroidal lattice. All that exists in the world besides the agents is a resource field, which is spread over the lattice of sites. The resource level at a given site is set at a value chosen from the interval $[0-R]$,

where R is the maximum resource level (chosen arbitrarily as 255). In the framework of Emery and Trist (1965), our model is a type-II (“placid, clustered”) rather than type-III (“disturbed, reactive”) environment, because the principal consideration is location rather than actual or potential inter-agent interactions.

Here we consider only *static* resource fields, i.e., fields in which resources are immediately replenished whenever they are consumed, so that the spatiotemporal resource distribution, i.e. structure, is constant. In static resource models the population has no effect on the distribution of resources. Nevertheless, since the agents constantly extract resources and expend them by living and reproducing, the agents function as the system's resource sinks and the whole system is dissipative.

Adaptation is resource driven since the agents need a steady supply of resources in order to survive and reproduce. Agents interact with the resource field at each time step by ingesting all of the resources (if any) found at their current location and storing it in their internal resource reservoir. Agents must continually replenish this reservoir to survive for they are assessed a constant resource tax at each time step. If an agent's internal resource supply drops to zero, it dies and disappears from the world. As a practical expedient for speeding up the simulation, each agent also runs a small risk, proportional to population size, of randomly dying.

Each agent moves each time step as dictated by its genetically encoded sensorimotor map: a table of behavior rules of the form: IF (environment j sensed) THEN (do behavior k). Only one agent can reside at a given site at a given time, so an agent randomly walks to the first unoccupied site near its destination if its sensorimotor map sends it to a site which is already occupied. (Population sizes range from about 2% to 10% of the number of sites in the world, so at the larger population sizes these collisions will occur with a non-negligible frequency.) An agent receives sensory information about the resources (but not the other agents) in its von Neumann neighborhood of the five sites above, below, to the left, to the right, and at its present location. An agent can discriminate only four resource levels (evenly distributed over the $[0-R]$ range of objective resource levels) at each site in its von Neumann neighborhood. Thus, each sensory state j corresponds to one of $4^5 = 1024$ different detectable local environments. Each behavior k is a jump vector between zero and fifteen sites in any one of the eight compass directions (north, northeast, east, etc.). The behavioral repertoire of these agents thus consists of $8 \times 16 = 128$ different possible behaviors. This sensorimotor map, consisting of a movement genetically hardwired for each detectable environmental condition, is the agent's “genotype.” These genotypes are extremely simple, amounting to nothing more than a lookup table of 1024 sensorimotor rules. On the other hand, the space in which adaptation occurs is vast, consisting of 128^{1024} distinct possible genotypes. (As the next section explains, in some environments some von Neumann neighborhoods do not

exist and so the corresponding sensorimotor rules cannot ever be used; this lowers the number of *effectively* different genotypes in these environments.)

An agent reproduces (asexually) 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 sensorimotor map, except that mutations may replace the behaviors associated with some sensory states with randomly chosen behaviors. The mutation rate parameter determines the probability of a mutation at a single locus, i.e., the probability that the behavior associated with a given sensory state changes. At the extreme case in which the mutation rate is set to one, a child's entire sensorimotor map is chosen at random.

Sensorimotor strategies evolve over generations. A given simulation starts with randomly distributed agents containing randomly chosen sensorimotor strategies. The model contains no *a priori* fitness function (Packard 1989), so the population's size and genetic constitution fluctuates with the contingencies of extracting resources. Agents with maladaptive strategies tend to find few resources and thus to die, taking their sensorimotor genes with them; by contrast, agents with adaptive strategies tend to find sufficient resources to reproduce, spreading their sensorimotor strategies (with mutations) through the population. The basic components of our model have some similarities to the LEE model studied by Menczer and Belew (1996) including: varying the adaptive challenge by varying the patterns in a resource grid, movement in the grid as the adaptable behavior, asexual reproduction, and no explicit fitness function.

During each time step in the simulation, each agent follows this sequence of events: it senses its present von Neumann neighborhood, moves to the new location dictated by its sensorimotor map, consumes any resources found at its new location, and then goes to a new location chosen at random from the entire lattice of sites. This repositioning constantly scatters the population over the entire environment, exposing it to the entire range of detectable environmental conditions. Since the resource field is static, the set of detectable environmental conditions remains fixed throughout a given simulation. Agents never have the opportunity to put together unbroken sequences of behaviors, since each behavior is followed by a random relocation. And since all agents are taxed equally, rather than being taxed according to distance moved, all that matters to an agent in a given detectable local environment is to jump to the site most likely to contain the most resources. Thus, the adaptive challenge the agents face is to make the best possible single move given specific sensory information about the local environment. Adaptation occurs through multiple instances of these one-step challenge-and-response trials.

Varying Environmental Structure

We want to study adaptation in a variety of environments that differ only in their environmental structure. At the

same time, to make population size a measure of adaptability that can be meaningfully compared across the different environments, we want all of these environments to have the same total quantity of resources. If we let R be the maximal possible resource level at a site (in the present simulation $R = 255$), we can achieve this goal by engineering the environments so that the average resource level at a site is $R / 2$. (Although a site can have any of 256 different objective resource levels, recall that the agents can discriminate only four resource levels.) The following suite of environments meets these desiderata:

1. Flat: Each site in this environment has a resource level set to $R / 2$.
2. Random: Resource levels in this environment are chosen at random with equal probability from the interval $[0-R]$, thus ensuring that the average level is $R / 2$.
3. Sine waves: Resource are assigned by two sine waves, one along the x -axis and the other along the y -axis. The amplitude of these waves is scaled in such a way that when both are maximal and overlapping the site has the maximum resource level, when both are minimal the site has no resources, and the average resource level is $R / 2$. The frequencies of the two sine waves can be varied independently and are expressed in the number of sine-wave periods which cover the x - or y -axes.
4. Substituting Flat or Random levels in Sine waves. In these environments the sine wave-generated resource level is substituted at randomly chosen sites with either constant or random values. Since the constant resource level is set equal to $R / 2$, and the random resource levels are chosen with equal probability from the interval $[0-R]$, the average resource level per site remains $R / 2$ regardless of the density of sites. The density of substituted sites is a model parameter.

In a previous paper (Fletcher, Zwick, and Bedau 1996) we provided several figures illustrating the various patterns generated in our suite of environments. We also discussed how these environments apply to Wilson's (1991) and Littman's (1993) environment classification schemes.

Quantitative Measures

To study how adaptability depends on environmental structure, we define separate measures of environmental structure and adaptive success. We then observe how adaptive success (our dependent variable) responds when we manipulate environmental structure (our independent variable). The measures we propose illuminate how adaptation and environmental structure interact.

Two Aspects of Environmental Structure

Adaptation is sensitive to those aspects of environmental

structure that the agents perceive and act upon. One such aspect is the variety of the environmental conditions which the agents can discriminate; a second is the utility provided by the environment for adapting to these environmental conditions. These two aspects correspond to the syntactic information content in the environment and pragmatic value of the information, respectively.

Information. A natural way to quantify the former is with the information-theoretic uncertainty or Shannon entropy (Shannon and Weaver 1949) of the distribution of detectable local conditions:

$$H(E) = -\sum_i F_E(v_i) \log_2 F_E(v_i)$$

where v_i is the i^{th} detectable local environmental condition (in this case, a distinct resource distribution in the von Neumann neighborhood), and $F_E(v_i)$ is the frequency of occurrence, across all sites in environment E , of v_i .

$H(E)$ measures the information content of the environmental conditions that the agents can detect, i.e., the reduction in uncertainty about v when an agent detects a local environmental condition. This measure is a particular way of integrating two aspects of the distribution $F_E(v)$: its width (number of different v) and flatness (constancy of $F_E(v_i)$). Everything else being equal, the wider or the flatter $F_E(v)$ is, the more uncertain an agent will be about which neighborhood it will detect, the more information an agent will get when it does detect its neighborhood, and the higher $H(E)$ will be. We can equivalently refer to $H(E)$ as the detectable environment's uncertainty, Shannon entropy, or information content.

Since the environments studied here all have static resource distributions, in every case $H(E)$ is constant over time. $H(E)$ would change in environments with dynamic resource distributions and thus would apply to a wide variety of environments in addition to those studied here.

Utility. To measure the pragmatic differences among environments, we calculate what the expected utility would be for a perfectly adapted population in each environment. We measure this ideal expected utility, $U^*(E)$, as how much resources on average each agent would receive per time step in a perfectly adapted population in excess over what the average agent would receive in a randomly behaving, non-adapted, population.

Like $H(E)$, $U^*(E)$ is a property of the environment, given the innate capacities of the agents, and it can be calculated *a priori*—before any simulation is run. For each distinct von Neumann neighborhood in the environment, the utility (above the average utility of random action) of all moves from each instance of the von Neumann neighborhood is tallied. The highest tally gives the best average expected utility for this neighborhood. The average of all the best expected utilities, weighted by the frequency of each neighborhood type in the environment, is $U^*(E)$. This would be the result of successful application of a Maximum Expected Value strategy in a game against nature (Von Neumann and Morgenstern 1944).

Given that the objective resource levels are the same in all our environments, it might seem that the ideal utility, $U^*(E)$, should be equal for all environments. There are two reasons why this is not the case:

1. The limited jump range of agents (15 sites) makes it impossible in some environments for agents to jump to a maximal resource hill from some environmental locations.
2. Even if an environment always contains a close by resource hill, agents can only have one behavior mapping per distinct von Neumann neighborhood type. Therefore, if the best action in different instances of that neighborhood is different, even a perfectly adapted agent could not move to a maximum benefit site from each environmental location.

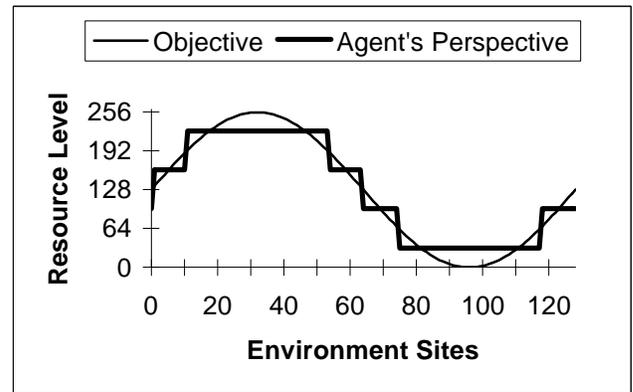


Figure 1 Side view of the 1 x 1 sine-wave environment in a 128 x 128 toroidal lattice of sites, showing both the objective resource field and the agents' perspective of it. Note that, although the objective resource level at a site can have 1 of 256 possible values, the agents can distinguish only 4 resource levels.

Figure 1 shows a cross-section of the 1 x 1 environment and illustrates both cases above. An agent sitting at position 80 jumping 15 sites to the left makes the best move possible, but still does not reach the resource maximum located at position 32. In this same environment, an agent sitting at position 5 would sense a resource plateau—the same resource level in each of the 5 sites that makes up the von Neumann neighborhood centered on its position. The ideal behavior from position 5 would be a jump 15 to the *right*, but at position 60 the agent senses the same von Neumann neighborhood and the ideal behavior is a jump 15 to the *left*.

The agent's coarse sensory discrimination and limited jump range make the $U^*(E)$ value for the 1 x 1 environment only 11.5. On the other hand, some high frequency, regularly patterned environments always have a maximum resource site unambiguously located within an agent's jump range, e.g. in the 64 x 64 sine-wave environment, the $U^*(E)$ value is 127.3 (very close to the

maximal value of $R/2$ or 127.5).

Note that $U^*(E)$ does not capture environmental utility *completely*. It says nothing about the distribution of benefits for sub-optimal behaviors. Nevertheless, it gives a rough measure of the differences in expected utility among a wide variety of environments.

To develop a feel for aspects of the detectable environmental structure measured by $H(E)$ and $U^*(E)$ consider our suite of environments:

1. If E is the flat environment, all local environmental conditions *are* identical, so they all *look* identical to the agents in the population. Thus, $H(\text{flat}) = 0$. Also, $U^*(\text{flat}) = 0$ since there is no structure to adapt to.
2. If E is the random environment, all detectable environments occur with (approximately) equal frequency, which makes $H(\text{random})$ close to its maximal value, which is \log_2 of the number of different v . Since the agents in our model can detect two bits of information about resource levels at each site in their von Neumann neighborhood, there are $4^5 = 2^{10}$ detectable environmental conditions, so $H(\text{random}) \approx 10$. (In the random environments we generated, typically $H(\text{random}) = 9.95$.) In a random environment, the best behavior varies among different instances of each von Neumann neighborhood. The $U^*(\text{random})$ value is about 85.
3. Sine-wave environments vary in the x and y frequency of the sine waves, and the number and frequency of detectable neighborhoods varies with these frequencies. Thus, $F_E(v)$ can have a variety of shapes, and both $H(E)$ and $U^*(E)$ can take a variety of values, as shown in the table below:

Environment	$H(E)$	$U^*(E)$
1 x 1	2.65	11.5
4 x 4	3.99	73.3
64 x 64	2.00	127.3
34 x 42	7.09	119.1

4. If some fraction of the sites in a sine-wave environment are replaced with flat or random resource levels, $H(E)$ and $U^*(E)$ values can vary quite a bit. Low density of replaced sites tend to make $F_E(v)$ slightly flatter, which makes $H(E)$ slightly higher, regardless of whether the resource levels in the new sites are flat or random. As the density of replaced sites approaches one however, depending on whether the substituted levels are flat or random, $F_E(v)$ approaches the shape of $F_{\text{flat}}(v)$ or $F_{\text{random}}(v)$, so $H(E)$ approaches the value of $H(\text{flat})$ or $H(\text{random})$ and $U^*(E)$ approaches $U^*(\text{flat})$ or $U^*(\text{random})$.

Finally, we wish to reiterate that both $H(E)$ and $U^*(E)$ do not simply reflect the objective properties (i.e., the resource field) of the environment; they reflect this field *as perceived* by agents of the population. In this respect, it is like the ways in which Wilson (1991) and Littman (1993)

characterize environments.

Measures of Adaptation

We have developed two different measures of adaptive success: an external one that is based on the gain in population size due to adaptation, and a more internal view of adaptation based on the degree to which agents achieve the ideal expected utility.

External Measure of Adaptation. The model we study here is resource driven, and a population's size reflects its ability to locate the resources found in the environment. Although, in all the environments we studied, objective resource levels were roughly equal, we cannot assume in general that observed population size *by itself* is an accurate indication of the degree to which adaptation has taken place. Given the resources available in the environment and given the agents' existence taxes, even non-adapting randomly behaving agents might still survive by accidentally "bumping into" resources. To factor this out, we compare equilibrium population size in a given environment with the equilibrium size of a "reference" population in exactly the same kind of environment. The reference population has exactly the same set of internal features (sensory and behavioral capacities, existence tax, etc.) as the observed population, *except* that its behaviors are always chosen at random instead of being based on sensory input. We denote this reference population size $P_R(E)$, while $P(E)$ denotes the actual equilibrium population size of population P in environment E. Thus $A_P(E)$, the adaptive success of population P in environment E is the proportion of increase above the reference population size:

$$A_P(E) = \frac{P(E) - P_R(E)}{P_R(E)}$$

Internal Measure of Adaptation. $A_P(E)$ gives us a way to compare the adaptive success of our standard population among different environments, but we can also look more closely at the internal causes of the observed population differences. That is, we can also express adaptive success in terms of $U(E)$: how much utility agents receive compared to non-adapting agents. As we will see, $U(E)$ and $A_P(E)$ are highly correlated.

Consider two environments that result in the same $A_P(E)$ and also result in the same $U(E)$ of 25 resource units per agent per time step. If $U^*(E)$ of one environment is 100 and $U^*(E)$ of the second is 50, then the population in the first environment achieved 25% of the ideal expected utility and in the second the population achieved 50% of the expected utility. The behaviors, and therefore the agents' genomes, are closer to the ideal in the second environment. Since the $A_P(E)$ values were the same, the $A_P(E)$ measure does not capture this adaptation difference. We capture this more internal view of adaptive success with the ratio of $U(E)$ to $U^*(E)$:

$$A_U(E) = \frac{U(E)}{U^*(E)}$$

We use the subscript U to distinguish adaptive success measured using the utility ratios from $A_p(E)$ which is our measure of adaptation based on population size.

Measures of Diversity in Adaptive Behavior

Finally, we can study another internal aspect of adaptation by measuring the change in Shannon entropy of the population's alleles (weighted by gene usage) as adaptation takes place. Three different measures are relevant. First, we can simply calculate the total Shannon entropy (diversity) of alleles used by a population. We periodically sample the population for a small time interval and calculate the Shannon entropy of all alleles used during this interval. We designate this by $H(L)$ where L stands for alleles. Second, we can bin this same data by gene and calculate the diversity of alleles within each gene (von Neumann neighborhood), and then average these results (weighted by gene usage) to calculate an overall within-gene allelic diversity. This we designate as $H(L|G)$: the diversity of allele, L , given the gene (or von Neumann neighborhood), G . Because there are $128 = 2^7$ different alleles, the maximum $H(L)$ or $H(L|G)$ value is 7. Third, the difference

$$I(L : G) = H(L) - H(L | G)$$

is the mutual information which can also be considered the "between-gene" diversity. In our present simple environments, variation increases $H(L|G)$ while selection reduces it; successful adaptation is reflected in an increase in $I(L : G)$. For more details on these information-theoretic measures, see Bedau, Zwick, and Bahm (1995).

NB: For notational simplicity, we now will drop the argument E for measures U , U^* , A_p , and A_U . H without an argument will always refer to environmental information, but we will write $H(E)$ explicitly where we need to differentiate it from other Shannon entropy measures.

Environmental Structure and Adaptation

We studied adaptation in a total of 70 different distinct environments with environmental information content, H , values ranging from 0.00 to 9.95 (10 is maximum) and ideal expected utility, U^* , values ranging from 0.00 to 127.25 (127.50 is maximum). For all of these runs we used a mutation rate of 0.001 and allowed population size to reach equilibrium. Our task is to understand the relationship between our independent environmental variables H and U^* , and our dependent variable, adaptive success, measured using either A_p or A_U .

Observations Using A_p

We first note a very high correlation between A_p values and the actual average utility, U . Figure 2 shows this relationship, which holds very well across the whole range of environments tested. This gives us confidence that average resource consumption is directly proportional to

population size gain, and therefore factors such as the chance of random death do not significantly affect population size results.

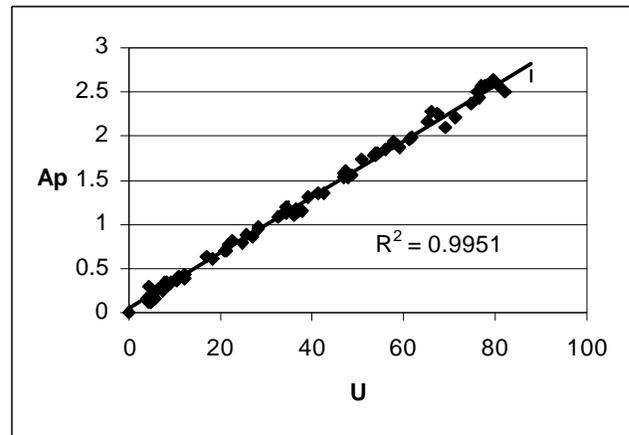


Figure 2 Adaptive success (measured using population) as a function of actual equilibrium utility per agent per time step. The relationship is proportional and shows that other factors besides U do not significantly contribute to A_p .

Next we examine how A_p depends on our environmental parameters. We expect that it will be inversely dependent on H and directly dependent on the ideal utility, U^* . In other words, increased adaptive success will be associated with less uncertainty of sensory inputs to adapt to, and higher utility for adapting. Figure 3 shows the relationship between A_p and U^* . We can see that A_p tends to increase with U^* , but this relationship is weak. This "fuzziness" may be partially due to our other environmental parameter H .

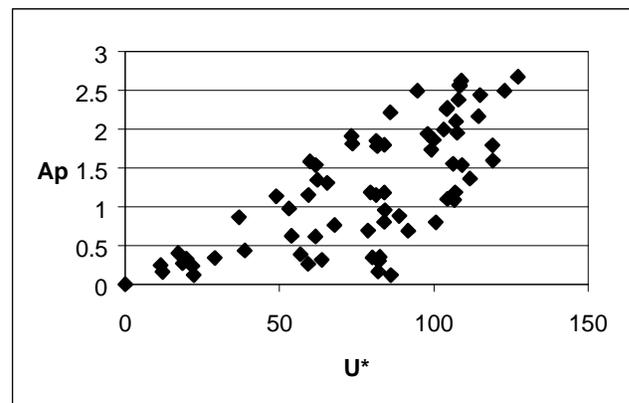


Figure 3 Adaptive success (measured using population) as a function of average ideal utility per agent per time step. This shows a weak relationship.

Figure 4 plots the relationship between A_p and $1/H$. Again, there is an indication of the expected relationship, but we observe two notable exceptions. First, there is a series of points that show a low A_p across the whole range

of I/H values. These tend to be environments with very low U^* values (typically less than 30). Second, there is series of environments where A_p is flat for high A_p values. We discuss this leveling off of adaptive success for environments with low H values (high I/H) in the next section.

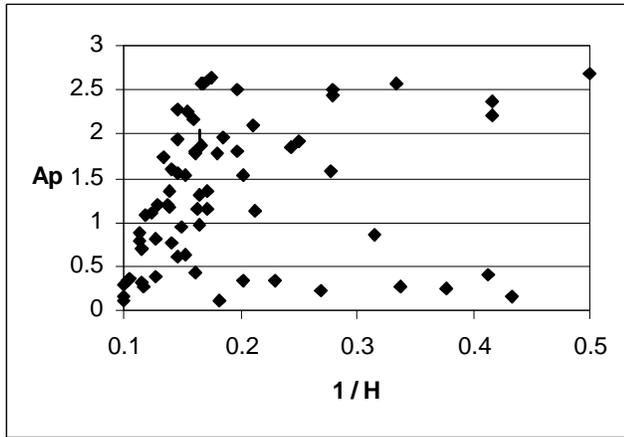
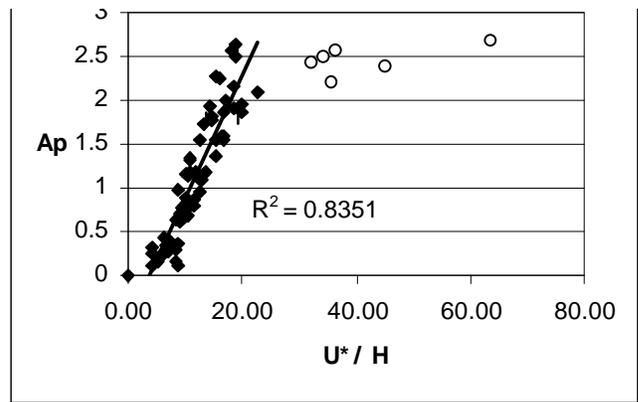


Figure 4 Adaptive success (measured using population) as a function of the inverse of environmental information content. This shows a weak relationship with notable exceptions for low and high A_p values.

Figure 5 Adaptive success (measured using population) as a function of ideal utility divided by environmental information content (utility per bit of environmental information). This shows a significant relationship for U^*/H less than 30. A linear fitting of all data gives an R^2 value of 0.5846.

In Figure 5 we combine our two measures of environmental structure by dividing U^* by H . The relationship between this combined measure and A_p is strikingly improved over the relationship with A_p of either individually. There are six points that do not fit well into the linear relationship. There appears to be a hard upper limit to A_p of about 2.5 where neither raising the utility nor lowering the uncertainty of sensory inputs raises A_p . We discuss possible reasons for this in the next section. We should note that U^*/H is only one of many possible ways to model our data. We have used other models, such as exponential and polynomial models, to fit the entire range of data, but since the number of points in the non-linear range is small, we restrict ourselves here to the linear model. The rationale for a linear dependence of U^*/H is compelling. It is a measure of the utility of perfect adaptation per bit of environmental sensory information to be adapted to. Across all the environments we studied with U^*/H less than 30, this ratio predicts how well the population will adapt (as measured by population size gain). For U^*/H less than 30 there is a linear relationship with A_p ($R^2 = 0.8351$). For U^*/H greater than 30, A_p is maximal at roughly 2.5.



We have also obtained reasonable fits of the data using simple linear regressions of A_p against U^* and H (or I/H), but the dependence of A_p on the composite U^*/H is more compact and readily interpretable.

Observations Using A_u

Although population size is a traditional way of measuring adaptive success, A_u has the advantage of having a hard upper limit defined by U^* . Also, it gives us an internal view of how behaviors (genomes) are changing as adaptation takes place. Figure 6 plots A_u against the ratio U^*/H . Although the general trends are the same as seen in Figure 5, the relationship is rather weaker.

Figure 6 Adaptive success (measured using achieved fraction of ideal utility) as a function of ideal utility divided by environmental information content (utility per bit of environmental information). This relationship using this measure of adaptive success is much less significant than the one using A_p (shown in Figure 5). A linear fitting of all data gives an R^2 value of 0.3088.

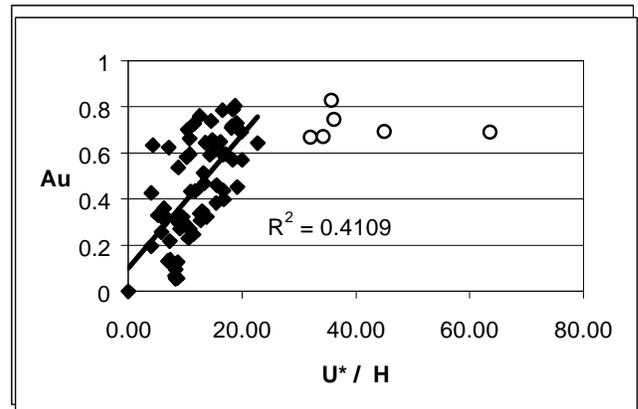


Figure 7 Adaptive success (measured using achieved fraction of ideal utility) as a function of the inverse of environmental information content. This shows a linear relationship with two notable exceptions: 1) very low U^* values (less than 30) shown with empty triangle; and 2) H values less than 5 (I/H values less than 0.2) shown with empty circles. This shows that for environments with moderate to high U^* values and H values greater than 5, the degree of utility achieved is inversely proportional to

the diversity of sensory inputs. For H values less than 5 ($1/H$ greater than 0.2), A_U is fairly flat at about 60-80%. A linear fitting of all data gives an R^2 value of 0.3007.

Figure 7 shows A_U plotted against $1/H$. This shows a relationship between the syntactic information content of environments, H , and the degree to which ideal utility is achieved. This relationship appears to be approximately linear with two notable exceptions. First, the points shown with empty triangles represent environments with very low U^* values (less than 30). For these environments, the degree of ideal utility achieved, A_U , appears to be not well correlated with $1/H$ values. The second exception occurs for H less than 5 ($1/H$ greater than 0.2). These points are shown with empty circles. In this region A_U appears to level off at about 60-80%. In other words, even as the adaptive task becomes easier (less uncertainty of sensory inputs), the degree of utility achieved through adaptation does not appear to improve appreciably. We would not expect a population to ever reach 100% perfect adaptation. The mutation rate alone would keep this from happening. Additionally, as mentioned before, as environments become more crowded, agents are more likely to land on each other and be diverted to another nearby site. In this case, even if agents' genomes were ideal, the utility gain would not be the ideal value since they were being "bumped" from their target site. Our measure depends on the actual utility and we would expect U to level off below the ideal, U^* , at least in part due to the reasons above. For the environments studied here and a mutation rate of 0.001, A_U plateaus at around 60-80% and lowering H (raising $1/H$) does not improve A_U . This plateau appears to begin for $1/H$ greater than 0.2 (or H less than 5). There was no discernable relationship for U^* vs. A_U —not shown.

Relationship Between A_P and A_U

We can tie our two measures of adaptive success together by deriving the relationship depicted in Figure 5 (which uses A_P) from the relationship depicted in Figure 7 (which uses A_U). The four steps below are intended to illustrate a hypothesized relationship, rather than to prove a mathematical one.

1. $1/H \propto A_U$ [relationship shown in Figure 7]
2. $1/H \propto U/U^*$ [definition of A_U]
3. $U^*/H \propto U$ [multiply both sides by U^*]
4. $U^*/H \propto A_P$ [$U \propto A_P$ by Figure 2]

We have, thus, $A_P \propto (A_U) \cdot (U^*)$. A_P measures the increase in population size due to adaptation and is dependent on the utility of adapting per bit of environmental sensory information to be adapted to (except for high values of this ratio where A_P is maximal). A_U measures the fraction of ideal utility achieved due to adaptation and indirectly the degree to which ideal behaviors are achieved. A_U is inversely dependent on the uncertainty of sensory inputs (with two exceptions: (1)

where expected utility is quite low—in which case A_U is variable, and (2) where the diversity of sensory inputs is low, in which case A_U is maximal).

In retrospect, if we had constructed our environments with constant U^* rather than constant total resources, we would expect A_U and A_P to show the same dependence on environmental information.

Ashby's Law of Requisite Variety

Another indication that a population is adapting to an environment is that the variety of alleles across different genes matches the variety of sensory inputs and at the same time the variety of alleles for any particular gene is small. In other words, full adaptation to the static environments we are studying calls for the existence of a unique allele for each sensory input—that allele being the best behavioral response to the particular environmental condition.

At the start of a run, the distribution of alleles across all genes in the population is random and thus the overall allele uncertainty is maximal at $H(L) \approx 7$. Within each particular gene, the uncertainty, $H(L/G)$, over the population is similarly random and maximal, i.e., is also close to 7. By contrast, the between-gene diversity, which is the mutual information, $I(L;G)$, between allele and environmental condition is near zero.

Figure 8 illustrates what then happens as the population becomes well adapted. The between-gene diversity of the alleles, $I(L;G)$, which represents also the tightness of constraint between alleles and environmental conditions, approaches the uncertainty of the environment, $H(E)$, which equals 4. At the same time, the within-gene diversity, $H(L/G)$, drops to near zero, i.e., there is no allelic diversity not coupled to environmental diversity.

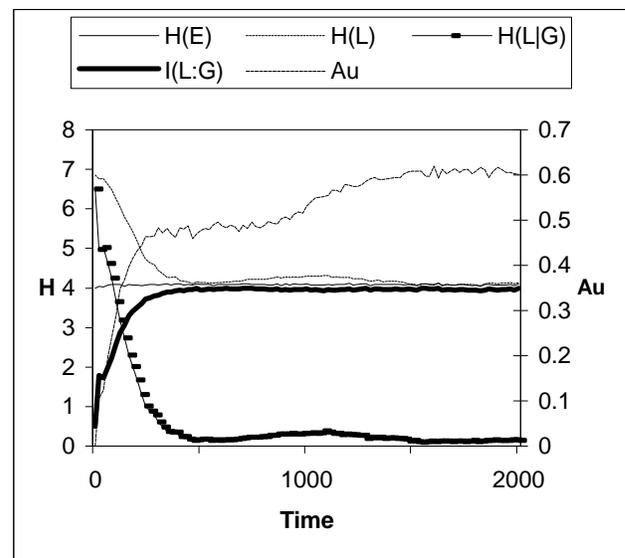


Figure 8 Time series showing the allelic diversity changes as adaptation takes place. A_U is also shown on the same

graph. This shows that as adaptation takes place the overall allelic diversity, $H(L)$, and the mutual information, between alleles and genes, $I(L:G)$, matches that of the environment, in this case 4.0. At the same time the diversity of alleles within genes, $H(L/G)$, drops to almost zero.

Both of these changes exemplify Ashby's "Law of Requisite Variety" (1956). Ashby's Law states that for optimal regulation a system needs to have a variety of responses to match the variety of environmental conditions it encounters, and second that this variety should not be mere randomness. This second point specifies that the "regulator of the system", here the genomic sensorimotor mapping, should be deterministic and not stochastic: for a particular environmental condition, there should be in the population only *one* response, namely the optimal one, not a mixture of responses. In Ashby's language, the uncertainty of the regulator state, given the disturbance, should be zero. This applies to the very simple evolving system studied here. For more complex evolutionary contexts (where resource levels are dynamic and inter-agent interactions are significant), there may well be advantages for non-zero $H(L/G)$. Both of these conditions are satisfied as adaptation approaches its maximal value: $H(L)$ matches $H(E)$ and $H(L/G)$ is near zero. It is interesting also to observe the slight bump in $H(L/G)$ at approximately $Time = 1000$. Here a temporary slight stochasticity of allelic response reflects the introduction through mutation of new and improved alleles, which as they spread, generate an increase in adaptive behaviors, A_u .

Conclusions

Our observations support two kinds of conclusions:

1. methodological conclusions about how to quantify major aspects of environmental challenge and adaptive success
2. substantive conclusions, based on experimental data, about how variations in these aspects of environmental challenge influence the degree to which populations in them adapt

Our measures of environmental structure have been applied to this simple evolutionary model, but we have defined them in general terms so that they can be applied across a wide range of evolving systems. Concepts such as the diversity of sensory inputs, H , and the utility for adaptive behaviors, U^* , are relevant to both artificial and natural systems where natural selection occurs. In addition, measuring external adaptive success by comparing population size to a non-adapting population in

the same environment, A_p , and internal adaptive success by measuring the degree to which ideal utility is achieved, A_u , can also be applied to many other adapting systems. We have also demonstrated how, for artificial systems where genome information is readily available, the diversity of alleles both across genes and within genes can illuminate the internal workings of the adaptive process.

On the substantive side, although there are certainly facets of environmental structure and adaptive success not captured by our measures, we have clear and testable indications about how environmental structure influences adaptation. For a wide range of environments, adaptive success depends upon both the syntactic (information-theoretic) and pragmatic (game-theoretic) aspects of environmental structure. In our work these two aspects are effectively integrated as utility per bit of sensory information. We expect that this measure will be useful in other studies of evolving systems. At the very least, both aspects of environmental structure will still need to be considered. When adaptive success is measured "internally" as the fraction of ideal utility gained, the utility aspect is encompassed implicitly. Adaptive success then depends simply (but less accurately) on sensory information content alone. Lastly, we have demonstrated explicitly Ashby's Law of Requisite Variety by showing that evolutionary adaptation is accomplished by the genomic representation of environmental information.

Acknowledgements

We would like to thank three anonymous reviewers for their detailed and thoughtful suggestions.

References

- Ashby, W. R. (1956). *An Introduction to Cybernetics*. London: Chapman & Hall.
- Bedau, M. A. (1994). The Evolution of Sensorimotor Functionality. In P. Gaussier and J. D. Nicoud (Eds.), *From Perception to Action*. New York: IEEE Press.
- Bedau, M. A. (1995). Three Illustrations of Artificial Life's Working Hypothesis. In W. Banzhaf and F. Eeckman (Eds.), *Evolution and Biocomputation—Computational Models of Evolution*. Berlin: Springer.
- Bedau, M. A., Bahm, A. (1994). Bifurcation Structure in Diversity Dynamics. In R. Brooks and P. Maes, (Eds.), *Artificial Life IV*. Cambridge, MA: Bradford/MIT Press.
- Bedau, M. A., Giger, M., Zwick, M. (1995). Adaptive

- Diversity Dynamics in Static Resource Models. *Advances in Systems Science and Applications*, 1, 1-6.
- 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*. Redwood City, CA: Addison-Wesley.
- Bedau, M. A., Ronneburg, F., Zwick, M. (1992). Dynamics of Diversity in an Evolving Population. In R. Manner and B. Manderick (Eds.), *Parallel Problem Solving from Nature, 2*. Amsterdam: North-Holland.
- Bedau, M. A., Zwick, M., Bahm, A. (1995). Variance and Uncertainty Measures of Population Diversity Dynamics. *Advances in Systems Science and Applications*, Special Issue I, 1-000.
- Emery, F. E., Trist, E. L. (1965). The Causal Texture of Organizational Environments. *Human Relations*, 18, 21-32.
- 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), 275-307.
- 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*. Cambridge, MA: Bradford/MIT Press.
- Menczer, F., Belew, R. K. (1996). From Complex Environments to Complex Behaviors. *Adaptive Behavior*, 4(3-4), 309-355.
- Packard, N. H. (1989). Intrinsic Adaptation in a Simple Model for Evolution. In C. G. Langton (Ed.), *Artificial Life*. Redwood City, CA: Addison-Wesley.
- Shannon, C. E., Weaver, W. (1949). *The Mathematical Theory of Communication*. Urbana, IL: University of Illinois Press.
- 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*. Cambridge, MA: Bradford/MIT Press.
- 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*. Cambridge, MA: Bradford/MIT Press.
- Todd, P. M., Yanco, H. A. (1996). Environmental Effects on Minimal Behaviors in the Minimat World. *Adaptive Behavior*, 4(3-4), 365-413.
- Von Neumann, J., and Morgenstern, O. (1944). *Theory of Games and Economic Behavior*. Princeton, N.J.: Princeton University Press.
- Wilson, S. W. (1991). The Animat Path to AI. In J. A. Meyer and S. W. Wilson (Eds.), *From Animals to Animats*. Cambridge, MA: Bradford/MIT Press.