In W. Banzhaf and F. E. Eeckman, eds.,
Evolution and Biocomputation: Computational Models of Evolution,
Lecture Notes in Computer Science Vol. 899 (Berlin: Springer, 1995), pp. 53-68.

Three Illustrations of Artificial Life's Working Hypothesis

Mark A. Bedau

Reed College, 3203 SE Woodstock Blvd., Portland OR 97202, USA
Email: mab@reed.edu

Abstract. Artificial life uses computer models to study the essential nature of the characteristic processes of complex adaptive systems—processes such as self-organization, adaptation, and evolution. Work in the field is guided by the working hypothesis that simple computer models can capture the essential nature of these processes. This hypothesis is illustrated by recent results with a simple population of computational agents whose sensorimotor functionality undergo open-ended adaptive evolution. These might illuminate three aspects of complex adaptive systems in general: punctuated equilibrium dynamics of diversity, a transition separating genetic order and disorder, and a law of adaptive evolutionary activity.

1 Artificial Life's Working Hypothesis

Artificial life studies computer models of the processes characteristic of complex adaptive systems—processes like self-organization, self-reproduction, adaptation, and evolution. Complex adaptive systems take many forms, each of which differs from the others in myriad ways. By abstracting away from the diverse details, artificial life hopes to reveal fundamental principles governing broad classes of complex adaptive systems. This hope rests on artificial life's working hypothesis that *simple computer models can capture the essential nature of complex adaptive systems* [1].

I propose to pursue artificial life's working hypothesis by applying a "thermodynamic" methodology [5, 6, 3, 4, 2, 7]. Recently it has been suggested that there is a close, intrinsic connection between the content of evolution and thermodynamics (e.g., Brooks and Wiley [8]). By contrast, I envisage the two fields as sharing the methodology of developing and investigating statistical macrovariables. Thermodynamics investigates macrovariables like temperature, pressure, and specific heat, and the fruits of this method include simple, basic laws and classifications (like the ideal gas law and the phase transition separating the solids and liquids). By analogy, the "thermodynamic" approach in artificial life seeks to identify statistical macrovariables that capture the distinctive features of complex adaptive systems. The most straightforward sign that this methodology is bearing fruit would be the demonstration that appropriate macrovariables can be used to frame simple, basic laws and classifications that apply to broad classes of complex adaptive systems.

This methodology involves formulating statistical macrovariables that are general enough to apply across a wide variety of systems, and then using these variables to search for underlying quantitative order unifying different systems. It is natural to begin this endevour with *simple models*, for macrovariables are easiest to formulate initially in simple models and simple models are easiest to study. Furthermore, simple models can reveal the essential nature of complex adaptive systems in general—at least, that is artificial life's working hypothesis.

This working hypothesis might be false, of course. It is at odds with the conclusions often drawn from the historicity, contingency, and variety of evolving biological systems (e.g., [25, 16]). One should bear in mind though that processes rife with historicity, complexity, and variety may well still fall under simple, basic laws and classifications, especially if these laws and classifications emerge through the application of statistical macrovariables. The "thermodynamic" methodology applied to simple computer models is a promising way to identify such laws and classifications, if they exist.

2 A Simple Model of Evolution

The model studied here is designed to be simple yet able to capture the essential features of an evolutionary process [27, 5, 6, 3, 4, 2, 7]. This model is motivated by the view that evolving life is typified by a population of agents whose continued existence depends on their sensorimotor functionality, i.e., their success at using local information to find and process the resources needed to survive and flourish. Thus, information processing and resource processing are the two internal processes that dominate agents' lives, and their primary goal—whether they know this or not—is to enhance their sensorimotor functionality by suitably coordinating these two internal processes. Since the requirements of sensorimotor functionality typically alter as the contingencies of evolution change, continued viability and vitality calls for sensorimotor functionality to adapt in an open-ended, autonomous fashion. The present model attempts to create agents with sensorimotor functionality that can undergo this open-ended, autonomous evolutionary adaptation.

The model consists of agents residing in a two-dimensional world, sensing their local environment, moving, and ingesting resources. All that exists in the world besides the agents are heaps of resources that are concentrated at particular locations, with levels decreasing with distance from a central location. The resource is refreshed periodically in time and randomly in space.

Agents interact with the resource field at each time step by extracting any found at their current site and storing it in their internal resource reservoir. Agents must continually replenish their internal resource supply to survive. Agents pay a resource tax just for living and a movement tax proportional to the distance traveled. If an agent's internal resource supply drops to zero, it dies and disappears from the world. On the other hand, an agent can remain alive indefinitely if it can continue to find sufficient resources.

An agent's movement is governed by its genetically hardwired sensorimotor strategy. A sensorimotor strategy is simply a map taking sensory data from a local neighborhood (the five site von Neumann neighborhood) to a vector indicating a magnitude and direction for movement:

$$S : (s_1, ..., s_5) \rightarrow \mathbf{v} = (r, \theta) \ . \tag{1}$$

A agent's sensory data has two bits of resolution for each site, allowing the agents to recognize four resource levels (minimal resources, somewhat more resources, much more resources, maximal resources). Its behavioral repertoire is also finite, with four bits of resolution for magnitude r (zero, one, ..., fifteen steps), and three bits for direction θ (north, northeast, east, ...). A unit step in the NE, SE, SW, or NW direction is defined as movement to the next diagonal site, so its magnitude is $\sqrt{2}$ times greater than a unit step in the N, E, S, or W direction. Each movement vector $\mathbf{v}$ thus produces a displacement (x, y) in a square space of possible spatial destinations from an agent's current location.

The graph of the strategy map S may be thought of as a look-up table with 2^{10} entries, each entry taking one of 2^7 possible values. This look-up table represents an agent's overall sensorimotor strategy. The entries are input-output pairs that link each sensory state (input) that an agent could possibly encounter with a specific behavior (output). The different entries in the look-up table represent genetic loci, and the movement vectors assigned to them represent alleles. Since agents have 1024 loci, each containing one out of a possible 128 alleles, the total number of different genotypes is 128^{1024}. Although finite, this space of genotypes allows for evolution in a huge space of genetic possibilities, which simulates the much larger number of possibilities in the biological world.

In order to investigate how adaptation affects the evolutionary dynamics of this model, I introduce a *behavioral noise* parameter, B_0, defined as the probability that an agent's behavior is chosen at random from the 2^7 possible behaviors, rather than determined by the agent's genetically encoded sensorimotor strategy. Thus, behavioral noise severs the link between genotype and phenotype. If $B_0 = 1$, then agents survive and reproduce differentially, and children inherit their parents' strategy elements (except for mutations), but the inherited strategies reflect only random genetic drift rather than the process of adaptation.

Sensorimotor strategies evolve over generations. An agent reproduces (asexually) when its internal resource supply crosses a threshold. The parent produces one child, which is given half of its parent's supply of resources. Parental allele values are inherited except when a point mutation at a locus gives a child a randomly chosen allele value. The mutation rate μ determines the probability with which individual locus mutate during reproduction. At the limit of $\mu = 1$, every allele value will mutate and thus each allele of child is chosen completely randomly.

It is important to note that selection and adaptation in the model are "intrinsic" or "indirect" in the sense that survival and reproduction are determined solely by the contingencies involved in each agent's finding and expending resources. No externally-specified fitness function governs the evolutionary dynamics [27, 5]. Good strategies for flourishing in this model would allow agents to acquire and manage resources efficiently. However, it is an open question which specific strategies *would* efficiently acquire and manage resources, and there might be no universally optimal strategy. A strategy's worth is relative to the environment; a strategy might be optimal in one environment and suboptimal in another. The environment of the present model consists of the fluctuating resource field and the competing strategies possessed by the agents in the population. Both of these environmental components change during the course of evolution. The strategies directly evolve, and the resource field indirectly changes because different populations of strategies affect it differently. For this reason, the model has the potential to show an open-ended evolutionary dynamic consisting of the perpetual creation of adaptive novelty.

This potential for an unpredictably shifting adaptive landscape is one reason the model resists treatment by the analytical methods used in traditional mathematical population genetics [9, 14, 15]. Not only are there thousands of loci and hundreds of alleles per locus, but the vicissitudes of natural selection indirectly cause unpredictable fluctuations in the finite population's size, age structure, and genotype distribution. In general, the only way to discern any underlying order in the model's behavior is through extensive computer simulation focussed on appropriate statistical macrovariables.

These complications notwithstanding, the model is an unabashedly abstract and idealized representation of a population of evolving agents, lacking many of the features often emphasized in the biological literature. For example, the environment lacks the spatial structure required for migration effects, there are no explicit interactions (such as predation) among organisms, there is no intron/exon distinction in the chromosome, and there is no "continuity" of mutation (mutated allele values are not "near" previous values). Nevertheless, my working hypothesis is that this model captures the fundamental features of complex adaptive systems, and is thus a useful model for investigating the essential aspects of more realistic systems.

3 Measurement of Population Diversity

Population diversity is one plausible statistical macrovariable for artificial life to investigate. But how might population diversity be measured? My proposal, very roughly, is to represent the population as a cloud of points in an abstract genetic space, and then define the population's diversity as the spread of that cloud. In the present model, an allele is a movement vector, a spatial displacement, and an agent's genotype is a set of spatial displacements. To capture the total population diversity, D, then, collect all the displacements of all agents in all environments into a cloud, and measure the spread or variance of that cloud.

We can divide this total diversity D into two components. First, collect the spatial displacements of each agent in the population in a given environment, i.e., the traits encoded across the population at a given locus, and calculate the spread of this locus's cloud. The average spread or variance of *all* such locus distributions is a population's within-locus diversity, W. Now, form another, second-order collection of the centroid each locus's cloud, i.e., a cloud of the "mean" displacement at each locus. The spread or variance of this second-order cloud is the population's between-locus diversity, B; it measures the diversity of the different mean population responses.

More formally, I define total diversity as the mean squared deviation between the average movement of the whole population, averaged over all agents and over all environmental conditions, and the individual movements of particular agents subject to particular conditions, i.e.,

$$D = \frac{1}{IJ} \sum_{i=1}^{I} \sum_{j=1}^{J} [(x_{ij} - \bar{x}^{IJ})^2 + (y_{ij} - \bar{y}^{IJ})^2] \ , \tag{2}$$

where I is the number of agents i, J is the number of environmental conditions (or, in the present model, loci) j, (x_{ij}, y_{ij}) is the movement vector of agent i subject to input j, and $\bar{x}^{IJ} = \frac{1}{IJ} \sum_{i=1}^{I} \sum_{j=1}^{J} x_{ij}$ (similarly for $\bar{y}^{IJ}$). So, $(\bar{x}^{IJ}, \bar{y}^{IJ})$ is the (x, y) displacement of the population averaged over all agents i and loci (environments) j. Then, the within- and between-locus components of the total diversity are defined as follows:

$$W = \frac{1}{IJ} \sum_{i=1}^{I} \sum_{j=1}^{J} [(x_{ij} - \bar{x_j}^{I})^2 + (y_{ij} - \bar{y_j}^{I})^2] \ , \tag{3}$$

$$B = \frac{1}{J} \sum_{j=1}^{J} [(\bar{x_j}^{I} - \bar{x}^{IJ})^2 + (\bar{y_j}^{I} - \bar{y}^{IJ})^2] \ , \tag{4}$$

where $\bar{x_j}^{I} = \frac{1}{I} \sum_{i=1}^{I} x_{ij}$ (and similarly for $\bar{y_j}^{I}$). So, $(\bar{x_j}^{I}, \bar{y_j}^{I})$ is the (x, y) displacement of the population in locus (environment) j averaged over all agents i. (Further formal analysis of diversity and its components is developed elsewhere [6, 3, 4].) From the analysis of variance [20], we know that the total diversity is the sum of the within- and between-locus components, $D = W + B$.

The relative size of D, W, and B reflects a population's genetic structure, as two extreme kinds of populations can illustrate. First, consider a population consisting of "random agents," in the sense that each agent's alleles are chosen randomly from the set of possible alleles, different agent's alleles being chosen independently. In this case, the distribution across the population at any given locus will be a huge cloud covering the whole set of possible spatial displacements, so the population's within-locus diversity W will be quite large. Since the centroid of each of these huge clouds will be virtually the same point—the center of the space of possible behavioral displacements—the distribution of these centers of gravity will be quite tight, and so the between-locus diversity will be nearly zero, $B \approx 0$. The population's total diversity will approximately equal the within-locus diversity, $D \approx W$.

A second extreme case is a population consisting of "quasi-clonal" (nearly genetically identical) agents that act differently in different environments. In this case, the within-locus diversity is nearly zero, $W \approx 0$, since the average spread of the cloud of behavioral displacements at each environment-locus is minimal. On the other hand, since the average behaviors in different environments are quite different, the between-locus diversity is large and equal to the total diversity, $D \approx B$. In this way, the relations among D, W, and B clearly distinguish the quasi-clonal and random agent populations.

3.1 Punctuated Equilibria

One of the most controversial topics in recent evolutionary biology has been the existence, cause, and implications of punctuated equilibria [13, 17, 10, 23, 26]. Artificial life systems might shed some new light on this controversy, since they often display punctuated equilibria in quantities like species concentration and average fitness (e.g, [19, 21, 28]). Yet the causes of these punctuated dynamics remain uncertain. Ecological complications such as host–parasite interactions or genetic complications such as extensive epistasis are typically thought to be implicated, and it is almost universally assumed that adaptation plays an essential role. My observations question whether any of these factors are essential.

I measured diversity in a series of simulations in which mutation rate and the presence or absence of adaptation were varied, while all other parameters of the model, including the size of the world and the resource environment, were held constant. Alleles were assigned to the founder population randomly, with displacement direction chosen from the eight compass directions and distance in steps chosen from zero, one and two. Thus, in the founder population, the total diversity was relatively low, $D = 2.5$, and virtually all of the total diversity was in the within-locus component, $D \approx W$ and $B \approx 0$.

Diversity dynamics in the present model routinely display clear punctuated equilibria when the mutation rate is suitably low. Figure 1 shows the typical dynamics of diversity for simulations in which $\mu = 10^{-5}$. Diversity remains largely static for significant periods of time, but every now and then diversity is punctuated by very rapid changes. The resulting picture is characterized by relatively flat plateaus separated by abrupt cliffs. (Figure 1 shows the within- and between-locus diversity components, W and B. The interesting diversity punctuations occur with respect to B. B approximates D since W is very low in these simulations and $D = W + B$, so the punctuations also occur with respect to D.)

It is notable that these punctuated equilibria occur in such a simple model. None of the ecological or genetical complications usually thought to be implicated are explicitly present in the model. For example, the model allows no explicit ecological interactions like those between host and parasite and the genetic structure has no epistasis. It is true that the model *could* support the emergence of implicit sub-populations that follow competing or cooperating resource-finding strategies. If such sub-populations were to exist, they would produce a substantial within-locus diversity W, for the average trait at given loci would differ between the sub-populations. The slightly positive values of within-locus diversity W in the simulation with adaptation (Fig. 1, top) is too low to be consistent with significantly different sub-populations. The simulations without adaptation

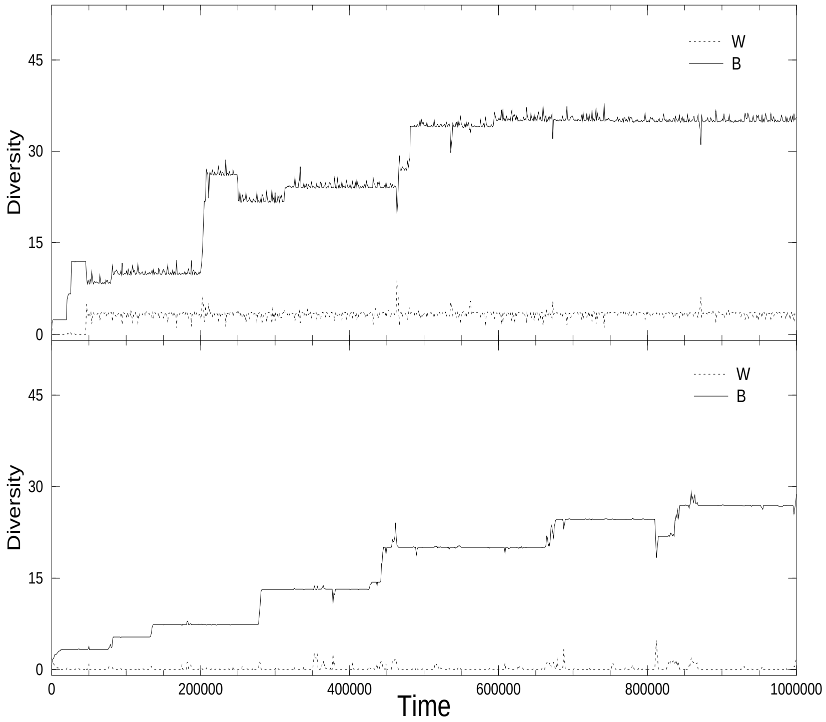

Fig. 1. Punctuated equilibria in diversity dynamics from the first 1,000,000 time steps of two typical low-mutation simulations ($\mu = 10^{-5}$). Adaptation above ($B_0 = 0$), no adaptation below ($B_0 = 1$). Time series for the two diversity components, W and B, are shown. The founder populations in these simulations have fairly low diversity, so punctuations initially tend to increase diversity, as shown here. On longer time scales, punctuations are equally likely to decrease and increase diversity.

(Fig. 1, bottom) show W values virtually equal to zero, which means that the population is virtually clonal and so has no sub-populations. Thus, although interactions between sub-populations might sometimes contribute to punctuations in some of the simulations, in general sub-populations play no fundamental role in the punctuated equilibria we observe in this model.

The most striking aspect of these punctuations is their presence even when adaptation is absent. Although punctuated equilibria in the absence of adaptation occur only when the mutation rate μ is suitably low, the effect is quite robust. Therefore, the presumption that punctuated equilibria must reflect the operation of adaptation is simply wrong. If punctuated equilibria are observed in the presence of adaptation, without additional evidence one cannot assume that adaptation plays any important role in their genesis. Evidently, there is an intrinsic tendency for evolving systems absent adaptation—that is, stochastically branching, trait-transmitting processes—to produce punctuated diversity dynamics, provided the branching rate is suitably poised.

3.2 Transition Separating Genetic Order and Disorder

Punctuated diversity dynamics fit into a broader pattern suggesting that evolving systems can be classified into two qualitatively different categories. I measured total diversity D and its within-locus W and between-locus B components in a series of pairs of adaptation/no-adaptation simulations, smoothly varying the mutation rate μ (on a log scale). The resulting diversity data reveal a transition separating two qualitatively different kinds of genetic systems.

One indication of this transition comes from the qualitative nature of the observed diversity dynamics. As noted in the previous section, when μ is low diversity dynamics typically consist of punctuated equilibria, the frequency of which is proportional to the mutation rate. On the other hand, when μ is high the diversity dynamics exhibit noisy fluctuations around a stable equilibrium value. The amplitude of these fluctuations is inversely proportional to the mutation rate.

The relationship between the total diversity and its two components clearly indicates the two different kinds of genetic systems and the transition between them. When the mutation rate is low, the total diversity is well approximated by the between-locus diversity, $D \approx B$. This shows that low mutation systems consist of the sort of "quasi-clonal" population mentioned in Sect. 3. On the other hand, when the mutation rate is high, the total diversity is well approximated by the within-locus diversity, $D \approx W$. Thus, high mutation systems consist of the sort of "random agent" population also mentioned in Sect. 3.

The way in which the transition between these quasi-clonal and random populations depends on mutation rate can be made vivid by plotting the *component diversity*, i.e., the extent to which the total diversity D is dominated by neither W nor B but has a large contribution from each. The component diversity can be defined as the proportion of the area of a square of side D is covered by a rectangle with sides $2W$ and $2B$:

$$C = \frac{4WB}{D^2} \quad . \tag{5}$$

(The factor of 4 scales C so that $0 \leq C \leq 1$.) I noted above that W will be near zero in a quasi-clonal population, and B will be near zero in a random population. Thus, the component diversity C will be near zero in both of these two kinds of populations. The component diversity C can approach one only if neither diversity component dominates the total diversity, which would entail that the population is neither quasi-clonal nor random.

Figure 2 shows the time average of the component diversity C as a function of the mutation rate, for systems both with and without adaptation. A transition between two qualitatively different genetic systems is clearly indicated. Notice that C is close to zero if the mutation rate is either high or low, and C approaches its maximal value of one at intermediate mutation rates, roughly, $10^{-3} \leq \mu \leq$

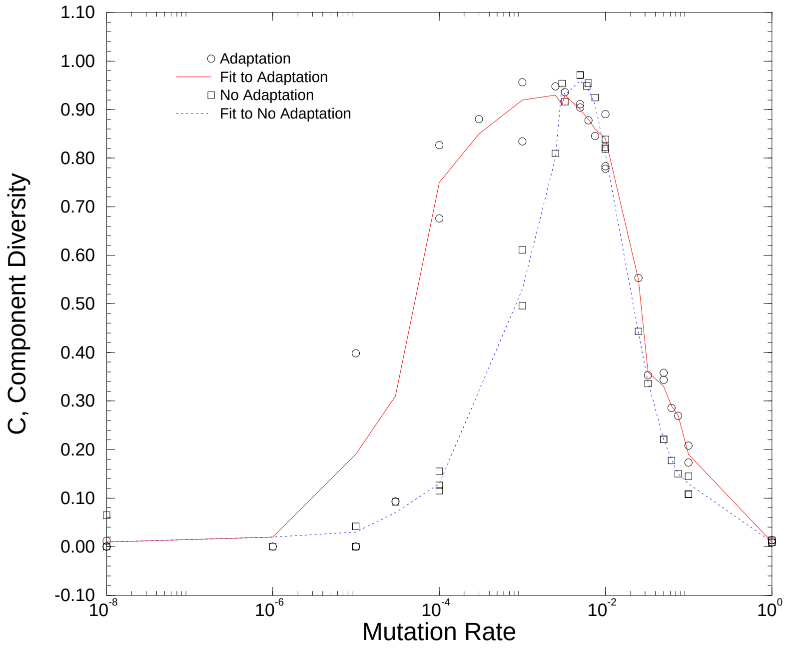

Fig. 2. Transition in diversity dynamics, reflected by the time average of the component diversity, C, as a function of mutation rate (shown on a log scale to improve resolution). The transition separates two regions of qualitatively different behavior. Systems with low mutation rate μ are genetically "ordered"—the genetic structure of each agent in general is highly correlated with those of the other agents. High μ systems are genetically "disordered"—the genetic structure of each agent in the population is uncorrelated with those of the other agents. (The leftmost data points represent not $\mu = 10^{-8}$ but $\mu = 0$.)

10^{-2}. It is striking that this transition exists whether or not the agents' genetic strategies are adapting during the course of evolution. Even if all genes are merely drifting because of the operation of behavioral noise, we still see the two qualitatively different genetic systems and the transition between them. (In fact, the transition seems to be sharper without adaptation. Further details about the diversity dynamics and the effects of mutation and adaptation are described elsewhere [6, 3, 4].)

Figure 2 paints a picture of an abstract space of evolving systems with two qualitatively distinct regions dividing the mutation spectrum. Low mutation systems are genetically "ordered," consisting of a population of genetically identical (or, nearly identical) agents—a quasi-clonal population. Different loci encode different traits, and from time to time this more or less static distribution of traits across loci abruptly shifts, causing punctuations in the prevailing genetic stasis. By contrast, high mutation systems are genetically "disordered," consisting of a population of genetically dissimilar agents, each of which has a random collection of alleles—a random population. Over time, the gene pool is a continually fluctuating random distribution. These ordered and disordered regions are separated by a transitional region. (Whether this transitional region itself contains further structure is a topic of ongoing work.)

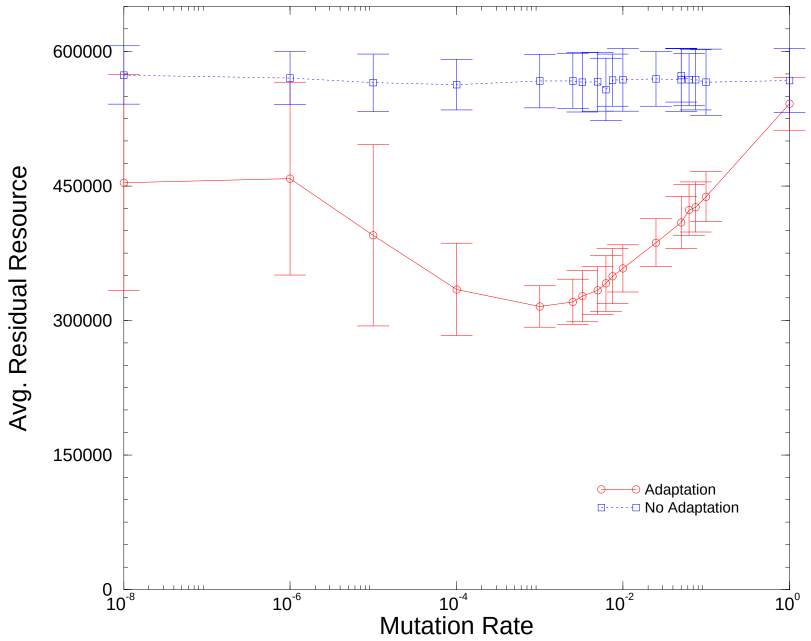

Fig. 3. Time averages of the amount of uningested resource in the world as mutation rate μ is varied (shown on a log scale to improve resolution). In one set of simulations adaptation operates normally; in the other set of simulations adaptation is prevented with behavioral noise $B_0 = 1$. The "bars" surrounding each point indicate the standard deviation of the time series of resource values. (The leftmost data points represent not $\mu = 10^{-8}$ but $\mu = 0$.)

Figure 3 shows that this transition separating genetic order and disorder has a striking connection with population fitness. Since my model is resource-driven, the population's overall fitness is reflected by its efficiency at extracting the available resources from the environment. (Exactly the same amount of resources were pumped into all simulations.) A crude (inverse) measure of this resource-extraction efficiency is the amount of residual (uningested) resource present in the world. The time average of residual resource is plotted against mutation rate, in Fig. 3. When the dependence of residual resource is compared with the diversity transition shown in Fig. 2, we can see that maximal resource-extraction efficiency occurs when the mutation rate is at or slightly below the transition (a region that one might describe as near "the edge of disorder"). As the mutation rate rises significantly into the region in which systems are disordered, resource-extraction efficiency falls off dramatically. (There is some indication that fitness also falls off if the mutation rate is well into the region of ordered systems, but this is unclear since it is difficult to gather clean statistics at very low mutation rates.) Although the transition between genetic order and disorder exists whether or not adaptation happens, effective adaptation is evidently optimal around this edge of disorder.

This effect might reflect a balance between two competing demands of evolutionary learning. On the one hand, the need to remember what has been learned requires a sufficiently low mutation rate; on the other hand, the need to explore novel possibilities requires a sufficiently high mutation rate. Optimal evolutionary learning, then, requires a mutation rate that appropriately balances these competing needs. This optimally poised mutation rate appears to coincide with the region around the edge of disorder.

4 Measurement of Adaptive Evolutionary Activity

A fundamental feature of any complex adaptive system is its adaptive evolutionary dynamics. But how might this property be measured? I think that we should conceive of adaptive evolutionary activity as the creation through the evolutionary process of sensorimotor functionality, i.e., of sensorimotor traits that are beneficial to the agents that possess them and that persist in the population because of this benefit. But how might this process be measured, especially when we might not know which traits have *any* functionality, and, if they do, what kind and how much? The difficulty—some would say impossibility—of answering this question was stressed in a classic paper by Gould and Lewontin [18] which subsequently generated a flood of critical debate (e.g., [12, 11, 29, 22, 24]).

I propose that we can address this issue by measuring the extent to which a trait is well-tested by natural selection. Every time as agent uses one of its sensorimotor traits, natural selection has an opportunity to provide some feedback about the trait's benefit or cost. If the trait persists in the lineage through repeated use and, in particular, accumulates more usage than would be expected *a priori*, then we have evidence that it is persisting because of its beneficial effects. Measuring a trait's adaptive significance in this way, then, involves measuring the extent to which its use exceeds *a priori* expectations.

In the context of the present model, sensorimotor traits are alleles. To measure the "raw" usage of an allele, assign a usage variable $u^t_{i_s}$ to the s^{th} allele of the $i^{\rm th}$ agent. An allele's usage variable is set to zero when the allele first enters the population through mutation (or at the very beginning of the simulation). Then, usage is incremented every time an allele is actually used, i.e., when the agent receives the sensory input genetically linked with the $s^{\rm th}$ locus and the behavior encoded by the $s^{\rm th}$ allele is thereby triggered:

$$u_{is}^{t+1} = \begin{cases} u_{is}^t + 1 & \text{if } i \text{ uses the } s^{\rm th} \text{ allele at } t \\ u_{is}^t & \text{otherwise} \end{cases} . \tag{6}$$

Recall that, if $B_0 > 0$, behavioral noise can prevent the $s^{\rm th}$ allele from actually producing i's behavior at t; in this case, i would *not* use the $s^{\rm th}$ allele even after receiving the sensory input that normally triggers its use. If $B_0 = 1$ then $u_{is}^t = 0$ for all i, s, and t.

Not all raw usage indicates an allele's adaptive significance, however, since harmful alleles accumulate usage. (In fact, it is only when harmful alleles are used that natural selection can eliminate them.) To determine an allele's proven adaptive value, we need to screen off that usage that might not signify the allele's adaptive value. One way to do this is to measure the *duration* during which an allele lineage could accumulate usage in the absence of adaptation, and then count an allele's usage as having adaptive significance only if the allele's age exceeds this duration.

More precisely, one can measure the extent of the adaptive evolutionary activity underlying all traits in a given simulation of the model—with some particular setting of model parameters (resource field, resource taxes, size of world, etc.)—as follows. Let the age A_{is} of the $s^{\rm th}$ allele of the $i^{\rm th}$ agent be defined as the number of time steps since that allele was originally introduced into i's genetic lineage by a mutation at the $s^{\rm th}$ locus. Then measure the age distribution of all alleles of all agents at the model parameter settings of interest *except that behavioral noise is fully turned on,* $B_0 = 1$. Adaptation cannot affect the allele age distribution when behavioral noise is always present—genotype and phenotype are unconnected—so the allele age distribution reflects only genetic drift.

Given this measured distribution of ages A_{is}, define the *drift duration*, t_μ, as the shortest duration which is less than A_{is} for all s and i. We can be quite confident that, in the simulation of interest, no allele can survive in the population for longer than the drift duration *if* the allele's presence is due to chance alone. To calculate the "net" usage $\tilde{u}_{is}$ of the $s^{\rm th}$ allele of the i agent, we modify Eq. 6 by adding the constraint that the allele's age must exceed the drift duration t_μ:

$$\tilde{u}_{is}^{t+1} = \begin{cases} \tilde{u}_{is}^t + 1 & \text{if } i \text{ uses the } s^{\rm th} \text{ allele at } t \text{ and } A_{is} \geq t_\mu \\ \tilde{u}_{is}^t & \text{otherwise} \end{cases} . \tag{7}$$

Finally, adaptive evolutionary activity A^t is simply the sum of the net usage:

$$A^t = \sum_{i,s} \tilde{u}_{is}^t \ . \tag{8}$$

4.1 A Law of Adaptive Evolutionary Activity

The drift duration t_μ was measured in a series of simulations across the mutation spectrum. (Limited computational resources prevented measurement of t_μ for $\mu \leq 10^{-3}$.) All model parameters were set exactly as in the simulations discussed in Sec. 3.1 and Sec. 3.2 above. Then the time average $A = \langle A^t \rangle_t$ of evolutionary activity was measured across the mutation spectrum, for various values of behavioral noise, $0 \leq B_0 \leq .25$.

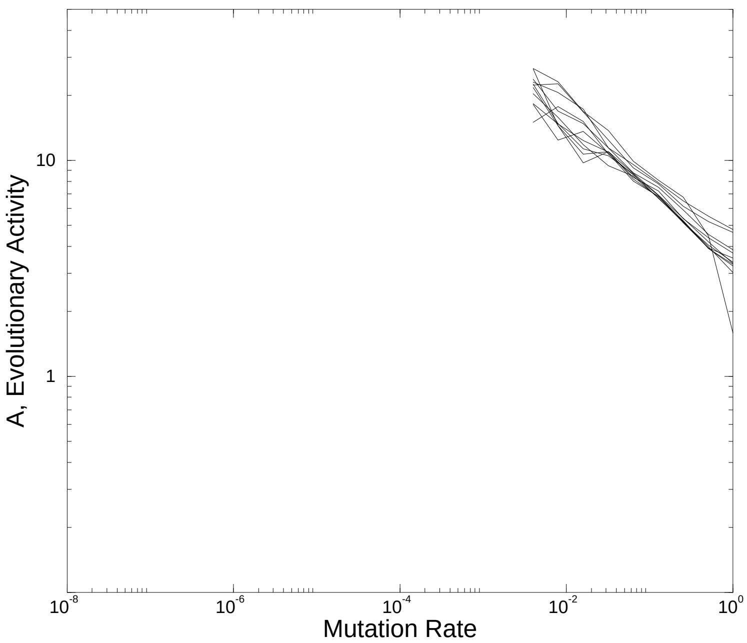

Fig. 4. Average evolutionary activity A as a function of mutation rate for several values of behavioral noise, $0 \leq B_0 \leq .25$. To facilitate comparison with Fig. 2 and Fig. 3, the same mutation rate scale is used. Due to the computational resources necessary for the calculation of the drift duration t_μ when $\mu \leq 10^{-3}$, evolutionary activity has not yet been measured at lower mutation rates.

Figure 4 shows how A was observed to depend on the mutation rate μ. We see that, within the range of mutation rates sampled, evolutionary activity approximately follows a power law:

$$A = \mu^{\alpha} \ , \tag{9}$$

with $\alpha \approx -2.3 \pm 0.3$. Notice that the dependence of adaptive evolutionary activity A on the mutation rate corresponds very closely with the dependence of resource-extraction efficiency on mutation rate depicted in Fig. 3.

It is notable that the approximate power law behavior of A in Fig. 4 holds up at a dozen different (relatively low) values of behavioral noise. This suggests that the law of adaptive evolutionary activity in Eq. 9 is fairly robust. An open question (requiring significant computational resources to answer) is how A will change when μ passes through and below the transition separating genetic order and disorder shown in Fig. 2. This question is especially intriguing given the adaptive significance of the transition revealed in when Fig. 2 is overlayed with Fig. 3.

5 The Status of Artificial Life's Working Hypothesis

The three results discussed here—punctuated equilibria in diversity dynamics, the transition separating genetic order and disorder, and the empirical law of adaptive evolutionary activity—illustrate the possible fruits of artificial life's working hypothesis that simple computer models can capture the essential nature of complex adaptive systems. I say *possible* fruits because it is not clear that these three effects *are* part of the essential nature of complex adaptive systems in general. Still, the results in the present model are sufficiently compelling for us to seriously entertain the hypotheses that these punctuation, transition, and power law effects have some significant universal application.

These three *specific* hypotheses about punctuation, transition, and adaptation must be sharply distinguished from the general *working* hypothesis that underlies this whole line of research in artificial life. The specific hypotheses are candidates for confirmation or disconfirmation in the short run, but the working hypothesis is not. In the short run, the working hypothesis is to be judged by whether it generates fruitful lines of research.

When held to this standard, the results presented above give the working hypothesis some provisionally credibility. The punctuation, transition, and adaptation results found in the present simple model will prompt the search for evidence for similar effects in other complex adaptive systems, both artificial and natural, and this in turn wil prompt the development of maximally general formulations of macrovariables like D, W, B, and A. These are exciting and promising lines of research.

In the long run, working hypotheses often *can* be effectively confirmed or disconfirmed. Artificial life's working hypothesis will win confirmation if enough of the specific hypothesis (like punctuation, transition, and adaptation) it spawns prove to be compelling. Whether this is so is an empirical matter, one which the "thermodynamic" methodology illustrated in this paper is well suited to address.

But how plausible are the three specific hypotheses about punctuation, transition, and adaptation? Are punctuated equilibrium diversity dynamics, a transition separating genetic order and disorder, and a power law dependence of evolutionary activity on mutation rate part of the essential nature of some significant class of complex adaptive systems? These questions remain open. But there is a straightforward empirical method by which we can pursue their answers. The hypotheses are eminently testable. Testing such hypothesis in a wide variety of artificial and natural systems is my vision of artificial life as-it-could-be.

References

1. M. A. Bedau, 1992, “Philosophical Aspects of Artificial Life,” in F. J. Varela and P. Bourgine, *Towards a Practice of Autonomous Systems*, Bradford/MIT Press, Cambridge, MA.
2. M. A. Bedau, 1994, “The Evolution of Sensorimotor Functionality,” in P. Gaussier and J.-D. Nicoud, eds., eds., *From Perception to Action*, IEEE Computer Society Press, Los Alamitos, CA.
3. M. A. Bedau and A. Bahm, 1993, “Order and Chaos in the Evolution of Diversity,” in *Proceedings of the Second European Conference on Artificial Life*, Brussels, Belgium.
4. M. A. Bedau and A. Bahm, 1994, “Bifurcation Structure in Diversity Dynamics,” in R. Brooks and P. Maes, eds., *Artificial Life IV*, Bradford/MIT Press, Cambridge, MA.
5. M. A. Bedau, and N. H. Packard, 1991, “Measurement of Evolutionary Activity, Teleology, and Life,” in C. G. Langton, C. E. Taylor, J. D. Farmer, and S. Rasmussen, eds., *Artificial Life II*, SFI Studies in the Sciences of Complexity, Vol. X., Addison-Wesley, Redwood City, CA.
6. M. A. Bedau, F. Ronneburg, and M. Zwick, 1992, “Dynamics of Diversity in an Evolving Population,” in R. Männer and B. Manderick, eds., *Parallel Problem Solving from Nature, 2*, New York, Elsevier.
7. M. A. Bedau and R. Seymour, 1994, “Adaptation of Mutation Rates in a Simple Model of Evolution,” in R. Stonier and X. H. Yu, eds., *Complex Systems—Mechanisms of Adaptation*, IOS Press, Amsterdam.
8. D. R. Brooks and E. O. Wiley, 1988, *Evolution as Entropy*, second edition, Chicago University Press, Chicago.
9. J. F. Crow and M. Kimura, 1970, *An Introduction to Population Genetics Theory*, Harper and Row, New York.
10. R. Dawkins, 1982, *The Extended Phenotype*, Oxford University Press, New York.
11. R. Dawkins, 1983, “Adaptationism Was Always Predictive and Needs No Defense,” *Behavioral and Brain Sciences*, 6, 360–61.
12. D. C. Dennett, 1983, “Intentional Systems in Cognitive Ethology: the ‘Panglossian Paradigm’ Defended,” *Behavioral and Brain Sciences*, 6, 343–390.
13. N. Eldredge and S. J. Gould, 1972, “Punctuated Equilibria: An Alternative to Phyletic Gradualism,” in T. J. M. Schopf, ed., *Models in Paleobiology*, Freeman, Cooper and Company, San Francisco.
14. W. J. Ewens, 1979, *Mathematical Population Genetics*, Springer-Verlag, Berlin.
15. D. S. Falconer, 1981, *Introduction to Quantitative Genetics*, second edition, Wiley, New York.
16. S. J. Gould, 1989, *Wonderful Life*, Norton, New York.
17. S. J. Gould and N. Eldredge, 1977, “Punctuated Equilibria: The Tempo and Mode of Evolution Reconsidered,” *Paleobiology*, 3, 115–151.

18. S. J. Gould and R. C. Lewontin, 1979, "The Spandrels of San Marco and the Panglossian Paradigm: A Critique of the Adaptationist Programme," *Proceedings of the Royal Society B*, 205, 581–598.
19. D. Hillis, 1992, "Simulated Evolution and the Red Queen Hypothesis," Biocomputation Workshop, Monterey, June 22-24.
20. G. R. Iversen and H. Norpoth, 1976, *Analysis of Variance*, Sage Publications, Beverly Hills, CA.
21. K. Lindgren, 1991, "Evolutionary Phenomena in Simple Dynamics," in C. G. Langton, C. E. Taylor, J. D. Farmer, and S. Rasmussen, eds., *Artificial Life II*, SFI Studies in the Sciences of Complexity, Vol. X., Addison-Wesley, Redwood City, CA.
22. J. Maynard Smith, 1978, "Optimisation Theory in Evolution," *Annual Review of Ecology and Systematics*, 9, 31–56.
23. J. Maynard Smith, 1989, *Did Darwin Get It Right?*, Chapman and Hall, New York.
24. E. Mayr, 1983, "How To Carry Out the Adaptationist Program," *American Naturalist*, 121, 324–33.
25. E. Mayr, 1988, "Is Biology an Autonomous Science?" In his *Towards a New Philosophy of Biology*, Harvard University Press, Cambridge, MA.
26. E. Mayr, 1988, "Speciational Evolution through Punctuated Equilibria," in his *Towards a New Philosophy of Biology*, Harvard University Press, Cambridge, MA.
27. N. H. Packard, 1989, "Intrinsic Adaptation in a Simple Model for Evolution," in C. G. Langton, ed., *Artificial Life*, SFI Studies in the Sciences of Complexity, Vol. VI., Addison-Wesley, Redwood City, CA.
28. T. Ray, 1991, "An Approach to the Synthesis of Life," in C. G. Langton, C. E. Taylor, J. D. Farmer, and S. Rasmussen, eds., *Artificial Life II*, SFI Studies in the Sciences of Complexity, Vol. X., Addison-Wesley, Redwood City, CA.
29. A. Rosenberg, 1985, "Adaptationalist Imperatives and Panglossian Paradigms," in J. H. Fetzer, ed., *Sociobiology and Epistemology*, Reidel, Dordrecht.