In R. Brooks and P. Maes, eds., Artificial Life IV:
Proceedings of the Fourth International Workshop on the Synthesis and Simulation of Living Systems
(Cambridge, MA: MIT Press, 1994), pp. 258-268.

Bifurcation Structure in Diversity Dynamics

Mark A. Bedau
Reed College, 3213 SE Woodstock Blvd., Portland, OR 97202, USA
Email: mab@reed.edu

Alan Bahm
Photon Kinetics, 9405 SW Gemini Drive, Beaverton, OR 97005, USA
Email: bahm@reed.edu

Abstract

We propose a measure of total population diversity D of an evolving population of genetically specified individuals. Total diversity D is the sum of two components, within-gene diversity W_g and between-gene diversity B_g. We observe the dynamics of diversity in the context of a particular model, a two-dimensional world with organisms competing for resources and evolving by natural selection acting implicitly on genetic changes in their movement strategies. We examine how diversity dynamics and population performance—measured as the efficiency with which the population extracts energetic resources from its environment—depend on mutation rate and the presence or absence of selection.

Systematic exploration of mutation rates reveals a bifurcation into qualitatively different classes of diversity dynamics, whether or not selection is present. Class I: At low mutation rates, diversity dynamics exhibit "punctuated equilibria"—periods of static diversity values broken by rapid changes. Class II: At intermediate mutation rates, diversity undergoes large random fluctuations without always approaching any evident equilibrium value. Class III: At high mutation rates, diversity is stable, with small fluctuations around an equilibrium value. Optimal population performance occurs within a range of mutation rates that straddles the border between class I and class II. The relationships among diversity D and its components W_g and B_g reflects the typical features of these different classes of diversity dynamics as well as corresponding differences in the gene pool, which ranges from genetically similar individuals in class I to genetically dissimilar individuals in class III. The fact that class I dynamics occur whether or not selection is present suggests that stochastically branching trait transmission processes have an intrinsic tendency to exhibit punctuated equilibria in population diversity over a critical range of branching (mutation) rates.

1 The Evolution of Diversity

Complex adaptive systems are embodied in many settings, ranging from ecological populations of organisms, through immune systems of antigens and antibodies, even to networks of neurons in the brain. By abstracting away the diverse details, one can model complex adaptive systems at a level of generality that might reveal fundamental principles governing broad classes of such systems—this we take to be the working hypothesis of artificial life [?].

One reason for the impressive effects in many artificial life models is their "emergent" architecture: The system's global adaptive behavior emerges unpredictably from an explicitly modeled population of low-level individuals. We have been studying a class of models consisting of a population of computation agents (basically, individual computer programs) that interact with each other and with their environment in a way that allows natural selection implicitly to shape their strategies for achieving various global computational goals [?, ?, ?, ?, ?]. We define statistical "macrovariables"—loosely akin to thermodynamic macrovariables like pressure or temperature—that reflect fundamental aspects of a system's adaptive behavior. Then we try to identify simple laws relating these macrovariables to other fundamental system parameters and we try to use these macrovariables to identify and explain basic classifications of systems.

An obvious but striking feature of complex adaptive systems is the evolutionary dynamics of population diversity. How can population diversity be defined and measured? How does diversity change as a population evolves? How do diversity dynamics vary as a function of other fundamental system parameters, such as mutation rate and selection pressure? Does population diversity define qualitatively different kinds of evolving systems? The present study addresses these questions (see also [?, ?]).

2 A Simple Model of Evolution

The model used here is designed to be simple yet able to capture the essential features of an evolutionary process [?, ?, ?, ?, ?]. The model consists of organisms (sometimes called "bugs") moving about in a two dimensional world. The only thing that exists in the world besides the organisms is food. Food is put into the world in heaps that are concentrated at particular locations, with levels decreasing with distance from a central location.

Food is refreshed periodically in time and randomly in space. The frequency and size of the heaps are variable parameters in the simulation.

The food represents energy for the organisms. Organisms interact with the food field by eating it at their current site at each time step, decrementing the food value in the environment and incrementing their internal food supply. Organisms must continually replenish their internal food supply to survive. Surviving and moving expend energy. Organisms pay a tax just for living and a movement tax proportional to the distance traveled. If a organism's internal food supply drops to zero, it dies and disappears from the world. On the other hand, an organism can remain alive indefinitely if it can continue to find enough food. Any evolutionary learning that occurs in the model is the effect of the one stress of continually finding enough food to remain alive. A good strategy for flourishing in this model would be to efficiently acquire and manage vital energetic resources.

It is important to note that selection and adaptation in the model are "intrinsic" or "indirect" in the sense that survival and reproduction is determined solely by the contingencies involved in each organisms finding and expending food. No externally-specified fitness function governs the evolutionary dynamics [?, ?].

The organisms in this model follow individually different strategies for finding food (and hence are sometimes called "strategic bugs" [?]). The behavioral disposition of bugs is genetically hardwired. A behavioral strategy is simply a map taking an organism's current sensory state—information about its present local environment (the five site von Neumann neighborhood)—to a vector indicating the magnitude and direction of its subsequent movement:

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

A bug's sensory state has two bits of resolution for each site in its its local environment, allowing the bug to recognize four food levels at each site (least food, somewhat more food, much more food, most food). 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 $\vec{v}$ thus produces a displacement (x, y) in a square space of possible spatial destinations from a bug'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 organism's overall behavioral strategy. The entries are input-output pairs that link a specific behavior (output) with each sensory state (input) that an organism could possibly encounter. The input entries in the look-up table represent genes or loci, and the movement vectors assigned to them represent alleles. Since bugs have 1024 genes or loci, each containing one out of a possible 128 alleles or behaviors, the total number of different genomes is 128^{1024}. Although finite, this space of genomes allows for evolution in a huge space of genetic possibilities, which simulates the much larger number of possibilities in the biological world.

When a bug's internal food supply crosses a threshold, it produces some number of offspring by asexual budding. After reproduction, the parental food supply is divided equally among the new children and the parent(s). Parental genes are inherited with some probability of mutation. Point mutations of the genes change the output values of entries in a child's look-up table. The mutation rate μ determines the probability with which individual loci mutate during reproduction. At the limit of $\mu = 1$, every allele has probability one of mutating and thus each child's alleles are chosen completely randomly.

While mutation rate is an explicit parameter of the model, selection pressure is controlled indirectly by adjusting other explicit parameters. The parameter *output noise*, P_0, is defined as the probability that the behavior *actually* performed by a bug on a given occasion in a given local environment will be chosen *randomly* from the 2^7 possible behaviors, rather than determined by the bug's genes. If $P_0 = 1$, then natural selection has no opportunity to "test" the usefulness of the behavioral traits encoded in a bug's genome. The bugs are still subject to differential survival and reproduction, and so there is a sort of "selection," but the alleles or traits transmitted in reproduction reflect only random genetic drift. There is heritable genetic variation but no heritable *phenotypic* variation, so natural selection plays no role in shaping the evolution of either genotypes or phenotypes. In simulations reported in this paper, output noise P_0 was set to either zero or one, thus creating pairs of simulations in which all model parameters were identical except for the presence or absence of selection's effects on the course of evolution.

This model is a very abstract and idealized representation of a population of evolving organisms, and has many biologically unrealistic respects. Nevertheless, our working hypothesis is that this model captures many fundamental aspects of evolving systems, and is thus a useful way to investigate the essential aspects of more complex evolving systems.

3 Measures of Diversity

How might population diversity be measured? (To simplify terminology, in what follows "diversity" always means *population* diversity.) Our proposal, very roughly, is to represent the population as a cloud of points in an abstract genetic space, and then define its diversity as the spread of that cloud. In the present model, an allele is a movement vector, a spatial displacement triggered by the sensory state corresponding to a given local environment. An individual's genotype is a complete set of spatial displacements, one for each possible sensory state. To capture the total population diversity, D, then, collect all the displacements of all bugs in all environments into a cloud, and measure the spread or variance of that cloud. (One can define related measures of diversity based on information-theoretic uncertainty rather than variance [?].) More explicitly, we define total diversity as the mean squared deviation between the average movement of the whole population, averaged over all individuals and over all sensory states, and the individual movements of particular individuals subject to particular conditions, i.e.,

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

where I is the number of individuals i, J is the number of sensory states for local environments (or, in the present model, genes) j, (x_{ij}, y_{ij}) is the movement vector of individual 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 individuals i and sensory states (genes)j.

We can divide this total diversity D into two components. (Additional diversity components can also be defined and studied [?].) Collect the spatial displacements of each bug in the population corresponding to a given sensory state, i.e., the traits encoded across the population at a given gene locus, and calculate the spread or variance of this gene cloud. The average spread of *all* such gene clouds is a population's within-gene diversity W_g. More explicitly,

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

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 sensory state (gene) j averaged over all individuals i.

Now, form another, second-order collection of the centers of gravity of each gene cloud, i.e., a cloud of each "average" displacement across the population in a given gene. The spread or variance of this second cloud is the population's between-gene diversity B_g, which measures the diversity of the average population response across all sensory states (genes), thus:

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

An easy calculation shows that the total diversity is the sum of the within- and between-gene components, $D = W_g + B_g$.

Absolute diversity values presented here reflect the size of the model's output space. (To compare diversity measurements across different size output spaces, measurements could be normalized by the size of output space; since all our simulations have the same size output space, we have not done this.) The maximum possible diversity value corresponds to the distribution in output space that is peaked at the four corners; in this case, each point is maximally distant from the mean (in this case, the center of output space). In the present context in which the maximum displacement is fifteen squares, the diversity value of this "corner post" distribution is the sum of the x and y displacements from the mean, i.e., $15^2 + 15^2 = 450$. In a flat random distribution in our modified polar coordinate system of 128 possible movements, since movements in the NE, SE, SW, or NW directions are $\sqrt{2}$ times larger than movements in the N, E, S, or W directions, the diversity value of the flat distribution is $\frac{4[(1^2+\sqrt{2}^2)+(2^2+(2\sqrt{2})^2)+\cdots+(15^2+(15\sqrt{2})^2)]}{128} = 116.25$.

Intuitively, the flat distribution, not the "corner post" distribution, is maximally diverse, but the diversity measures defined here are higher for the "corner post" than for the flat distribution. (This situation is reversed when using analogous diversity measures in which variance is replaced by information-theoretic uncertainty [?].) However, recall that food is placed in the simulated world in heaps that slope away from the center and that the bugs pay a movement tax proportional to distance traveled. So, is is not surprising that observed diversity values exceed the value for the flat distribution only in special circumstances when selection is absent.

The relative proportions of the two diversity components reflects a population's genetic structure. Consider a population consisting of "random individuals," in the sense that each bug's alleles are chosen randomly from the set of possible alleles, different bug's genes being chosen independently. In this case, the distribution across the population at any given environment-gene will be a huge cloud covering the whole set of possible spatial displacements, so the population's within-gene diversity W_g will be quite large. Since the center of gravity 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-gene diversity will be nearly zero, $B_g \approx 0$. The population's total diversity will approximately equal the within-gene diversity, $D \approx W_g$.

Another extreme case is a population consisting of genetically identical bugs that are "sensitive" to their environment in that they behave differently when they sense different environments. In this case, the within-gene diversity is zero, $W_g = 0$, since the average spread of the cloud of behavioral displacements at each environment-gene is nil. On the other hand, since the average behaviors in different environments are quite different, the between-gene diversity is large and equal to the total diversity, $D = B_g$.

Typical results from the first 10,000 time steps of simulations with mutation rate $\mu = 0$ is shown in Figure 1. Notice that the within-gene diversity W_g drops over time. In fact, when selection is present W_g reaches zero within 3000 time steps, at which point the entire population has become genetically identical. Furthermore, the between-gene diversity B_g increases over time until $B_g = D$ when $W_g = 0$. These effects also happen when selection is absent, but they typically take substantially more time.

When selection is absent, it is to be expected that pure genetic drift will produce this crossing of the W_g and B_g components: W_g will drop as stochastic sampling fixes more loci, and B_g will rise as different traits become

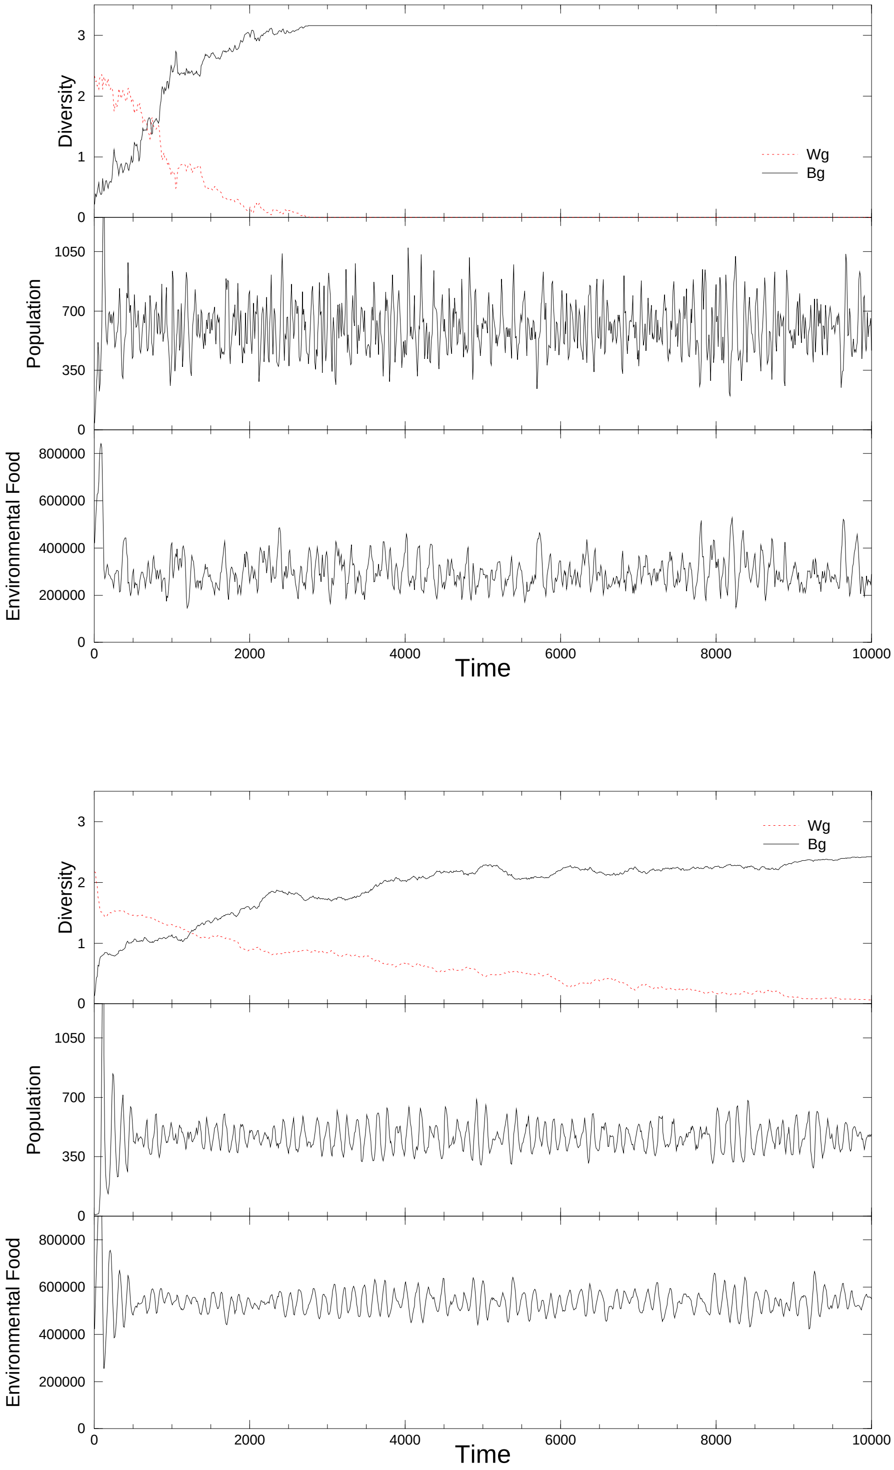

Figure 1: Within-gene W_g and between-gene B_g diversity, population level, and residual food in the environment in the first 10,000 time steps of a typical simulation with mutation rate $\mu = 0$. Top: selection ($P_0 = 0$). Bottom: no selection ($P_0 = 1$).

fixed. When selection is present, one expects the same component crossing but at a faster rate, since the effects of genetic drift are augmented by selection: W_g will drop as selection progressively weeds out traits at given loci, and B_g will rise as traits at different loci become specialized in different directions. Population level and residual environmental food data (Figure 1) also reveal the effects of selection, with selection supporting larger populations that extract more food from the environment.

4 The Bifurcation Structure

A bifurcation structure is a change in qualitative behavior caused as a parameter is smoothly varied [?]. Our observations suggest that there is a bifurcation structure in diversity dynamics as a function of the mutation rate μ, regardless of whether selection is present. As μ is varied, long-term diversity dynamics exhibit three qualitatively different patterns (summarized in Table 1). The transitions between these dynamics occurs at roughly the same mutation rates whether or not selection is present.

In class I, which is exhibited at low mutation rates, diversity dynamics exhibit "punctuated equilibria"—periods of static diversity values broken by rapid changes. In class II, which is exhibited at intermediate mutation rates, diversity undergoes large random fluctuations without approaching any evident equilibrium value. In class III, which is exhibited at high mutation rates, diversity is stable, with small fluctuations around an equilibrium value.

Some features of the bifurcation structure remain obscure. It is still uncertain how sharp the transitions are, and even how many transitions there are. For example, the transition between classes II and III might be quite smooth, and class II might be merely an extreme case of class III. Nevertheless, the qualitative character of class II dynamics is distinctive. Moreover, it is quite clear that classes I and III consist of two fundamentally different kinds of evolving systems. So, without prejudging the detailed nature of the bifurcation structure, we will use diversity D and its components W_g and B_g to describe the characteristic features of these three classes of behaviors.

The relationships among total diversity and its components reflects each kind of diversity dynamic. In class I, total diversity is dominated by its between-gene component, $D \approx B_g$ and $W_g \approx 0$. In class III, the situation is reversed; total diversity is dominated by the within-gene component, $D \approx W_g$ and $B_g \approx 0$. In class II, neither component dominates total diversity, $D \gg W_g \sim B_g \gg 0$.

These effects can be made vivid by plotting the diversity of the diversity components, i.e., the extent to which the total diversity D contains a large contribution from *both* components W_g and B_g (Figure 2). The diversity of the components, $C = \frac{4W_gB_g}{D^2}$, reflects what fraction of the area of a square of side D is contained in a rectangle with sides $2W_g$ and $2B_g$. (The factors of 2 scale C so that $0 \leq C \leq 1$.) Note that $C = 0$ if $W_g = 0$ or $B_g = 0$, and $C = 1$ if $2W_g = 2B_g = D$.

These relationships among D, W_g, and B_g indicate corresponding differences among the genetic structure of the population in the three classes. Class I populations remain highly similar while class III populations remain highly dissimilar. Class II populations are intermediate between being similar and dissimilar; the extent of similarity or dissimilarity is continually shifting rapidly within an intermediate range.

It is striking that the full bifurcation structure occurs whether or not selection is present. When selection is absent because $P_0 = 1$, there is no connection between the behavioral strategy encoded in a bug's genes and the bug's chances of survival and reproduction. This creates something very much like a stochastically branching trait transmitting process. So, the fact that the qualitative dynamics does not depend on whether selection is in effect strongly suggests that there is a general tendency for evolving systems generally to exhibit all three kinds of dynamics.

It is also striking that maximal population performance, measured by the efficiency with which the population extracts energetic resources from its environment, occurs in a range of mutation values spanning the border between classes I and II. This suggests that the mutation rates around the transition between classes I and II optimally balance evolutionary learning's competing demands for "memory" (reflected in low mutation rates) and "novelty" (reflected in high mutation rates). Experiments in which the mutation rate itself is allowed to evolve corroborate this suggestion, for there is a robust tendency for mutation rates to evolve toward this same transition zone and the evolved mutation distribution can be pushed higher (or lower) by engineering situations that call for greater novelty (or memory) [?].

Further discussion of diversity dynamics occurs elsewhere [?, ?].

4.1 The Simulations

We measured total diversity D and its within-gene W_g and between-gene B_g components in a series of pairs of selection/no-selection simulations, smoothly varying the mutation rate μ (on a log scale). All other parameters of the model, including the size of the world and the food environment, were held constant. Many parameter settings was multiply sampled. We simultaneously measured two crude aspects of the "performance" of the population—the population level and the amount of residual food in the environment—on the assumption that higher population level and lower residual food reflects better evolutionary learning on the part of the population.

Each simulation consisted of 10^6 time steps (every bug moves, and perhaps eats and reproduced, each time step). Diversity was sampled one thousand times in each simulation. To make the diversity more clearly reflect the action of selection, we collected statistics over only that portion of the genome that received the vast bulk of use (the thirty six most frequently used genes).

Founder populations in all simulations consisted of 100 bugs that were assigned traits randomly, with displace-

Table 1: Three qualitatively different kinds of diversity dynamics.

	CLASS I	CLASS II	CLASS III
μ RANGE	$\mu < 10^{-3}$ (approx.)	$10^{-3} \leq \mu \leq 10^{-2}$	$10^{-2} < \mu$ (approx.)
DYNAMICS	punctuated equilibria	large erratic fluctuations	small, stable fluctuations about equilibrium value
DOMINATING COMPONENT	B_g	none	W_g
GENE POOL	highly similar	somewhat similar, somewhat dissimilar, always rapidly shifting	highly dissimilar

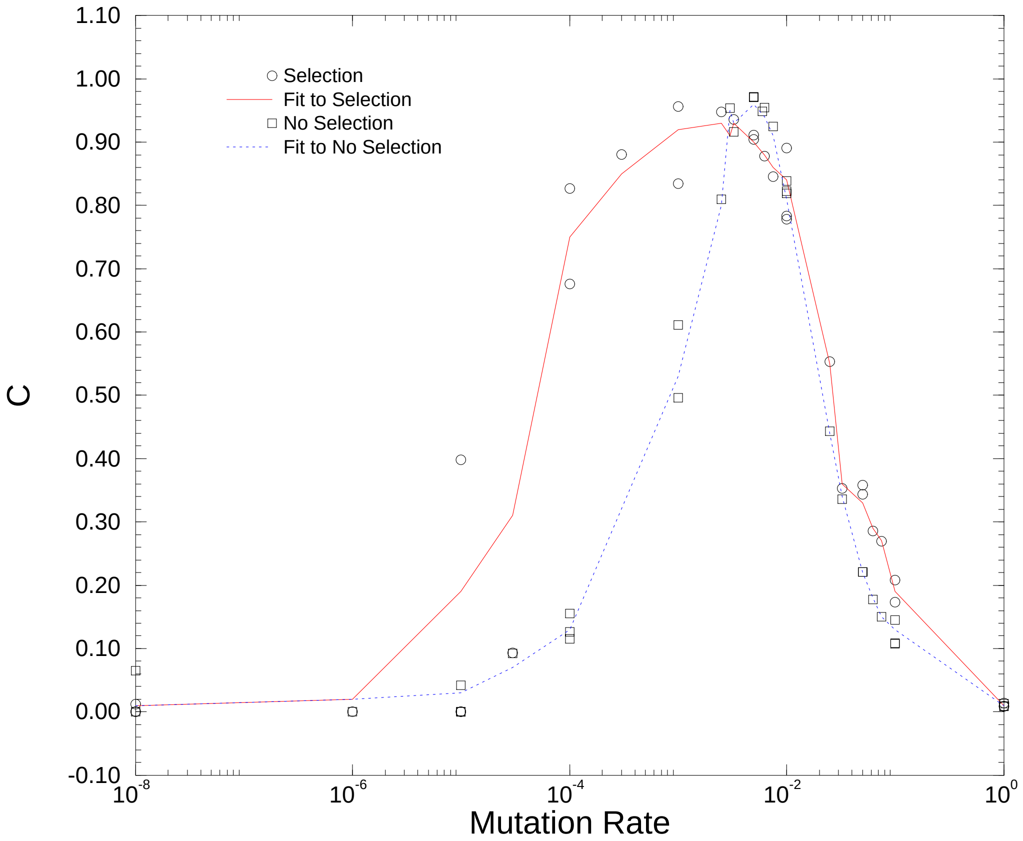

Figure 2: The time average of the diversity of the diversity components, $C = \frac{4W_g B_g}{D^2}$, as a function of mutation rate μ, for selection ($P_0 = 0$) and no selection ($P_0 = 1$). (The extreme left data points are for $\mu = 0$.) Fitted values for selection and no selection are included. The diversity of components evidently shows a sharper transition without selection.

ment 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-gene component, $D \approx W_g$ and $B_g \approx 0$.

4.2 Class I: Punctuated Equilibria

When μ is low, diversity dynamics display vivid punctuated equilibria—this characterizes class I diversity dynamics. Selection and no selection examples of the typical dynamics of diversity for simulations with very low mutation rates are displayed in Figure 3. The total diversity D remains basically static for long periods of time, but this pattern of stasis is punctuated occasionally by very rapid changes in diversity. The resulting picture is characterized by relatively flat plateaus separated by abrupt cliffs.

As one might expect, as μ increases, so does the frequency of the punctuations; in addition, the stability during the periods of stasis declines. If μ gets large enough ($\mu \geq 10^{-3}$, in the present model), punctuated equilibria are no longer evident, and the system is no longer in class I. As μ declines toward 0, the periods of stasis become more stable and longer lasting, until, at $\mu = 0$, the population remains genetically identical forever (after an initial transient) and the diversity dynamics reveal one, indefinitely long, unbroken period of stasis.

Class I systems have populations of genetically identical (or, nearly identical) individuals—in effect, a population of (near) "clones." Different loci encode different traits, and this distribution of traits across loci abruptly shifts from time to time, when B_g is punctuated. Hence, total diversity is dominated by the between-gene diversity, $W_g \approx 0$ so $D \approx B_g$, as Figure 3 reveals. Although B_g always dominates D in class I, as μ approaches the border between class I and class II, W_g tends to occupy an increasingly significant share of D.

4.3 Class II: Erratic Fluctuations

At intermediate μ values, total diversity D continually exhibits large and rapid random fluctuations on different time scales. Illustrations of these characteristic class II dynamics, with and without selection, are shown in Figure 4.

In class II systems, neither W_g nor B_g is close to 0, so neither component dominates D (Figure 2). Although the relative share of D captured by the two components is always shifting—indeed, at some μ values the two components are wildly criss-crossing—each component always claims a significant portion of the total diversity.

The intermediate values of W_g and B_g in class II systems indicates the presence of populations which are neither highly similar nor highly dissimilar, but have a complex and shifting genetic structure that is intermediate between these extremes. Precisely characterizing the variable degree of relationship among individuals in class II populations is a topic of current research.

4.4 Class III: Stable Dynamics

When μ is toward the high end of the spectrum, total diversity D exhibits small noisy fluctuations around a stable equilibrium value—this is the signature of class III. Figure 5 shows two examples of typical diversity dynamics in class III, one with and one without selection.

As the mutation rate increases, the amplitude of the fluctuations decreases. When $\mu = 1$ the equilibrium diversity value corresponds to the precise numerical value of a "flat" distribution of alleles, regardless of whether selection is present. In the limit of large populations, the amplitude of fluctuations at $\mu = 1$ becomes arbitrarily low. If μ is below 1, selection pulls the equilibrium value down, while the equilibrium value when there is no selection remains at the value of the flat distribution (compare top and bottom, Figure 5).

Class III populations consists of genetically distinct individuals, each of which has a random collection of alleles. The total diversity is well approximated by the within-gene diversity, $D \approx W_g$. When μ falls near the border between class II and class III, B_g starts to comprise a significant share of D.

4.5 Population Performance

The bifurcation structure has an interesting connection with optimal population performance, measured by how much food the population can extract from the environment. Three robust patterns in residual environmental food emerge (Figure 6). First, when selection is absent residual food is flat across the mutation scale and is significantly higher than in simulations with selection. This is to be expected; if $P_0 = 1$ and there is no selection, then a population's ability to find food should not depend on the mutation rate and it should be much worse than it would be if there were evolutionary learning.

Second, when natural selection shapes evolutionary learning, population performance worsens when μ approaches 1, and also when μ is extremely close to 0. This effect no doubt reflects the sacrifice of one of the two competing demands of evolutionary learning. On the one hand, the need to remember what has been learned calls for a sufficiently low mutation rate; on the other hand, the need to explore novel possibilities calls for a sufficiently high mutation rate. Thus, a very high mutation rate sacrifices a population's memory, while a very low mutation rate sacrifices its source of novelty. In either case, suboptimal performance results.

The third effect observed is that the range of maximal food extraction is when μ is broadly in the vicinity of the boundary between class I and class II. Our argument in the preceding paragraph implies that optimal evolutionary learning requires a mutation rate that appropriately balances the competing demands for memory and novelty. These optimal mutation rates evidently are found around the border between regimes I and II. This suggests that the bifurcation structure—which is independent of selection—is exploited to optimize evolutionary learning when selection is present. Experiments with evolving mutation rates tend to corroborate this suggestion [?].

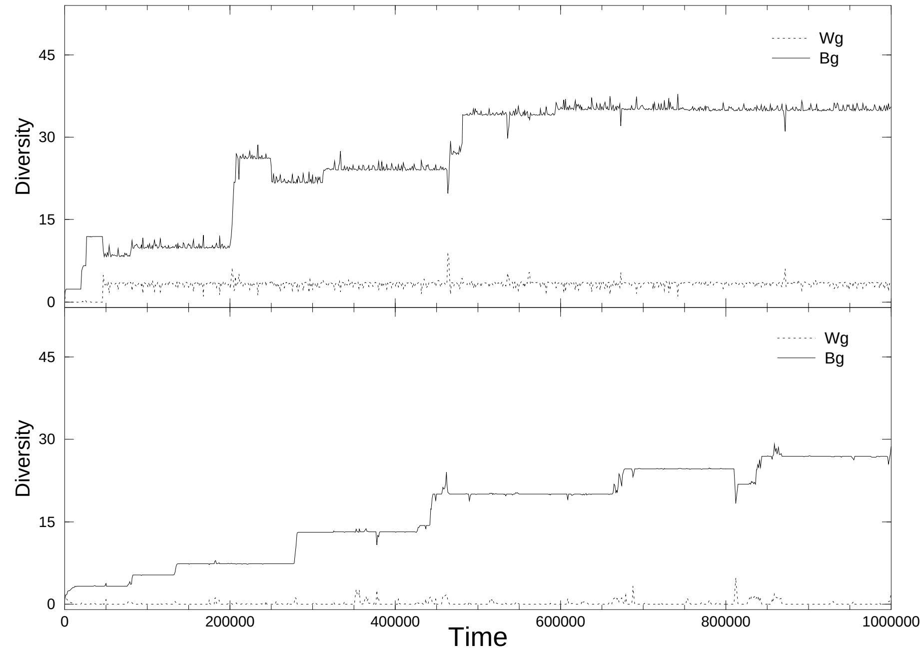

Figure 3: Typical class I punctuated equilibrium dynamics at $\mu = 10^{-5}$. Top: selection ($P_0 = 0$). Bottom: no selection ($P_0 = 1$). Punctuations tend to lead to increased diversity on this time scale in our simulations, since all our founder populations have low initial diversity. In the long run, punctuations continue indefinitely but show no general trend up or down.

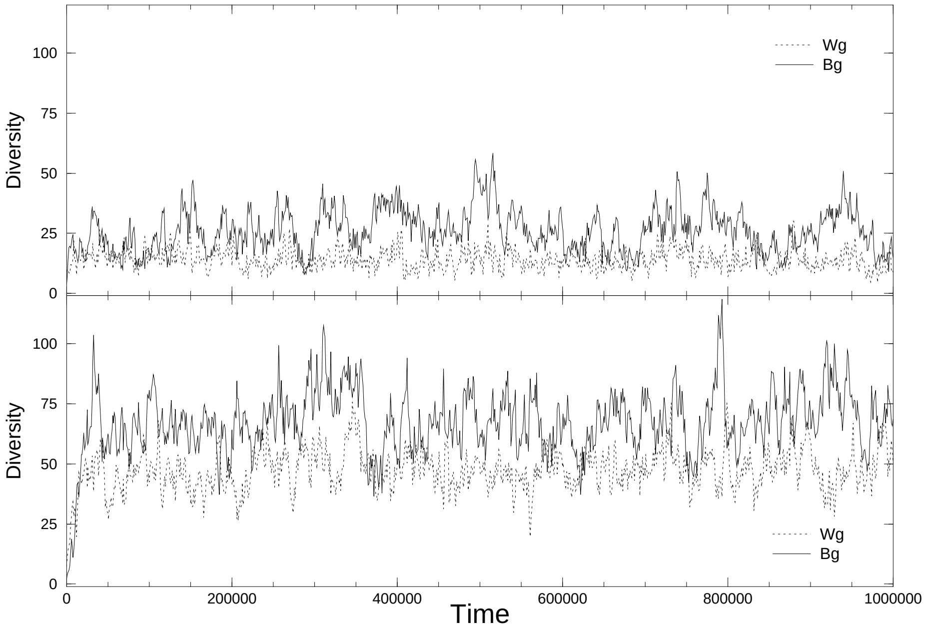

Figure 4: Typical class II dynamics with large random fluctuations at $\mu = 3 \times 10^{-3}$. Top: selection ($P_0 = 0$). Bottom: no selection ($P_0 = 1$).

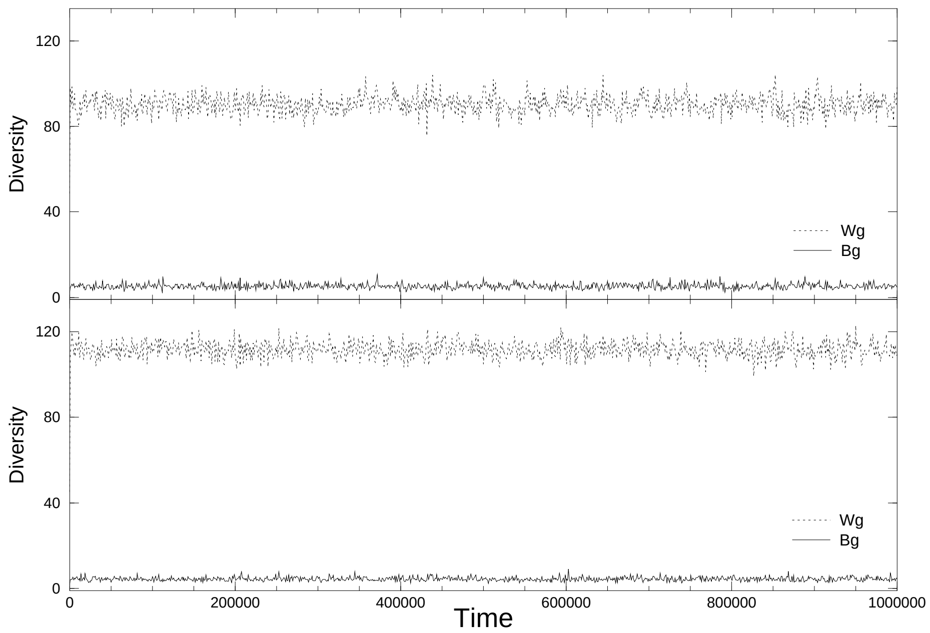

Figure 5: Typical class III stable dynamics at $\mu = 10^{-1}$. Top: selection ($P_0 = 0$). Bottom: no selection ($P_0 = 1$).

4.6 Generic Punctuated Equilibria

Artificial life systems commonly display punctuated equilibria in quantities like species concentration [?] and average fitness [?]. 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 selection plays an essential role.

Our observations question these presumptions, since we routinely observe punctuated equilibria in population diversity when our model is in class I. None of the ecological or genetical complications usually thought to play a role are explicitly present in the model. For example, the population has no explicit division into anything like host and parasite and the genetic structure has explicit no epistasis. However, there might be some *implicit* genetic complications, such as multigene traits, in which environmental regularities trigger regular sequences of genes which caused a characteristic sequence of movements. If such implicit genetic complications are present, their magnitude is up in the air. In addition, it is true that the model could support the emergence of implicit sub-populations following competing or cooperating food-finding strategies. Such sub-populations would be revealed by a substantial within-gene diversity W_g, since the average trait at given loci must differ between the sub-populations. Since punctuated equilibrium dynamics frequently occur when $W_g \approx 0$, (e.g., Figure 3, bottom), implicit sub-populations clearly play no necessary role in punctuated equilibria generally.

What is most striking about these punctuations is their occurrence even when natural selection is absent. Although punctuated equilibria in the absence of selection occur only when the mutation rate μ falls within an appropriate range near 0, the effect is quite robust. The presumption that punctuated equilibria in population diversity require the operation of natural selection is simply wrong. Therefore, even when punctuated equilibria occur *with* selection, without further evidence we cannot assume that selection plays any important role in its genesis. Evidently, there is an intrinsic tendency for evolving systems absent selection—that is, stochastically branching, trait-transmitting processes—to produce punctuated diversity dynamics, provided the branching rate is suitably poised. How to explain this effect remains a topic of current research.

5 A Science of Artificial Life

We are aiming to achieve two goals simultaneously: first, to develop plausible and useful measures of population diversity and, second, to use those measures to discern basic features of the evolution of diversity, ideally features that illuminate the fundamental nature of complex adaptive systems in general. Progress towards these goals is inextricably intertwined. One reason for finding our observations plausible and interesting is that our measures seem appropriate, and one sign that our measures are appropriate is that they reveal seemingly plausible and interesting effects.

To confirm the extent of our progress toward these goals requires further work. Analytical details remain to be settled, of course, and the need for a more precise statistical analysis of our effects calls for extensive further simulations. But the most important task is to determine the full generality of our results by replicating

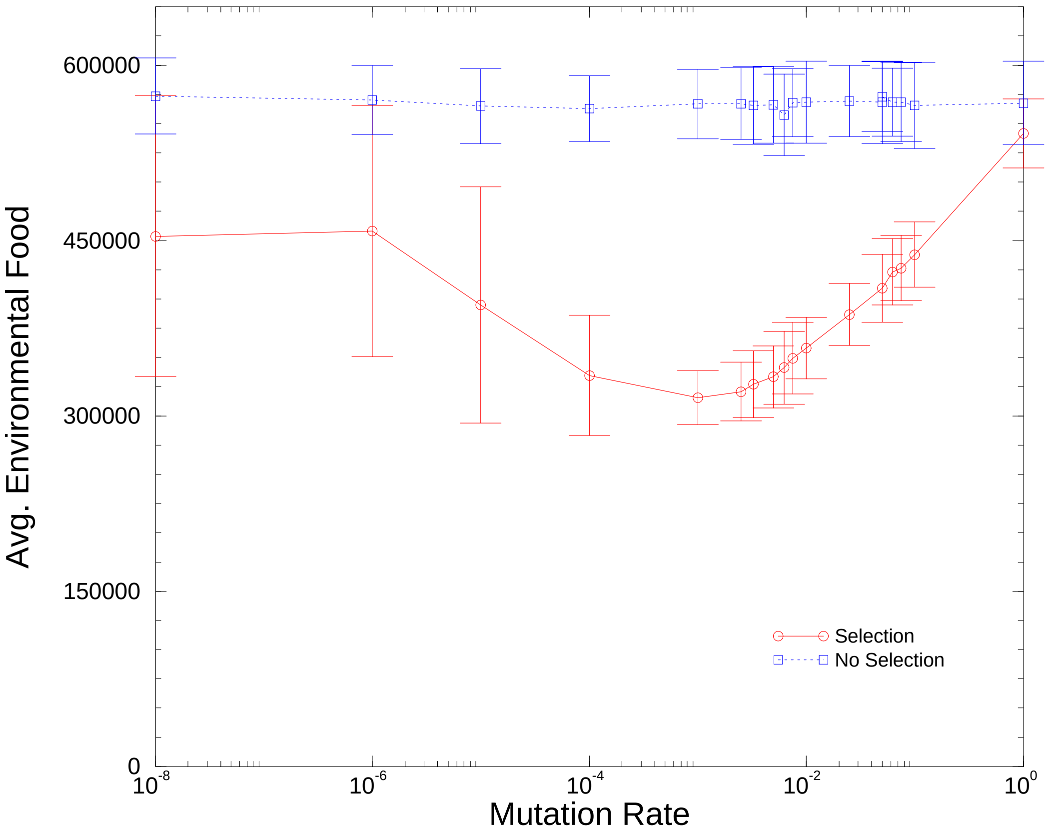

caption Time averages with standard errors of typical residual environmental food as mutation rate μ is varied, for selection ($P_0 = 0$) and no selection ($P_0 = 1$). (The data points on the extreme left are for $\mu = 0$.)

them in other models. Our observations were all made in one simple model of evolution, but the same measures can be implemented in other systems (or portions of systems)—both artificial and natural—in which all alleles at all genes share a metric. (And the uncertainty-based diversity measures can be implemented across a much broader range of systems [?].) The fact that the bifurcation structure exists both with and without selection is itself strong evidence that this structure is a fundamental feature of evolving systems generally. Final confirmation of the importance of our measures and the universality of our effects—and vindication of artificial life's fundamental working hypothesis—can come only from comparing quantitative results across a host of complex adaptive systems. Our effects provide one target for such comparisons.

Acknowledgements

For extensive helpful discussion, many thanks to Norman Packard and Marty Zwick. For expert programming assistance, thanks to Robert Seymour. Thanks also to audiences at Harvard University and the Rowland Institute for Science.

References

[1] M. A. Bedau, "Philosophical Aspects of Artificial Life." In F. J. Varela and P. Bourgine, *Towards a Practice of Autonomous Systems*. (Bradford/MIT Press: Cambridge, MA, 1992).

[2] M. A. Bedau and N. H. Packard, "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, 1991).

[3] M. A. Bedau, F. Ronneburg, and M. Zwick, "Dynamics of Diversity in an Evolving Population." In R. Männer and B. Manderick, eds., *Parallel Problem Solving from Nature, 2*. (New York: Elsevier, 1992).

[4] M. A. Bedau and R. Seymour, "Adaptation of Mutation Rates in a Simple Model of Evolution." Submitted to Complex'94.

[5] M. A. Bedau, M. Zwick, and A. Bahm, "Measures of Diversity in a Simple Model of Evolution." Preprint.

[6] S. Eubank and D. Farmer, "An Introduction to Chaos and Randomness." In D. L. Stein, ed., *Lectures in the Sciences of Complexity*, SFI Studies in the Sciences of Complexity, Lec. Vol. I. (Addison-Wesley: Redwood City, CA, 1989).

[7] D. Hillis, "Simulated Evolution and the Red Queen Hypothesis." Biocomputation Workshop, Monterey, June 22-24, 1992.

[8] K. Lindgren, "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, 1991).

[9] N. H. Packard, "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, 1989).