A Generic Neutral Model for Measuring Excess Evolutionary Activity of Genotypes

Andreas Rechtsteiner
Systems Science Ph.D. Program
Portland State University
P.O. Box 751, Portland OR 97207
andreas@sysc.pdx.edu
(503) 725-4992

Mark A. Bedau
Department of Philosophy
Reed College
3203 SE Woodstock Blvd, Portland OR 97202
mab@reed.edu
(503) 771-1112, ext. 7337

Abstract

We introduce and study a simple generic model of neutral evolution of genotypes, designed to provide a feasible and general method for quantifying excess evolutionary activity—the extent to which evolutionary activity is the product of adaptive evolution. We compare the behavior of the generic neutral model against two other models: Packard's agent-based model of the evolution of sensory-motor functionality and a neutral "shadow" of Packard's model. Diversity and evolutionary activity of these three models across the mutation rate spectrum illustrate the feasibility and general applicability of the generic neutral model, confirm the appropriateness of using neutral models to quantify the extent of the continual adaptive success of genotypes, and reveal power-law dependences of evolutionary activity on mutation rate.

1 The Need for Neutral Models

Although it is commonly accepted that adaptive evolution produces much of the structure and functionality in complex systems [6, 4], it is often difficult to distinguish adaptive change from other evolutionary phenomena such as random genetic drift [3]. Some even question whether adaptations can be objectively identified at all [3]. The ultimate goal of this paper is to facilitate the investigation of universal laws of adaptive evolution. Toward this end, this paper aims to develop statistics for objectively identifying and quantifying adaptive evolutionary activity, especially statistics feasible and general enough to apply to a broad enough range of natural and artificial evolutionary systems. This paper illustrates evolutionary activity statistics in the context of a simple artificial model of sensory-motor evolution—Packard's Bugs model—and we apply the method to a broadly applicable level of analysis—whole genotypes. In this setting we do find simple law-like regularities involving adaptive evolution. Part of what makes this especially interesting is that evolutionary activity statistics apply to myriad evolutionary systems at myriad levels of analysis, so we can investigate whether the same regulatities hold in evolving systems in general.

We use the approach of Bedau and Packard [1, 2] to identify the extent to which a system's evolutionary dynamics depend on adaptation rather than other evolutionary forces like chance and necessity. That is, we screen off the effect of non-adaptive evolutionary forces by comparing the evolutionary dynamics observed in target evolutionary systems with those observed in analogous evolutionary systems in which adaptive evolution cannot happen. We term these non-adaptive evolutionary data filters "neutral models" of evolution. Filtering observed data with a neutral model yields a measure of excess evolutionary activity—that activity due to adaptation. In effect, neutral models are null hypotheses against which the action of adaptive evolution stands out in relief.

One method for making neutral models is to craft a system that "shadows" the target evolutionary system in all relevant respects *except* that a shadow genotype's presence or concentration or longevity cannot be due to the genotype's adaptive significance [2]. Since such "neutral shadows" are specifically tailored to the target system of interest, they create sharp no-adaptation null hypotheses. But because they are tailor-made for specific target systems, studying new target systems requires constructing new shadow models, and it is vexing to compare models that shadow different target systems.

An obvious way to address these problems is to create a simple *generic* neutral model—one neutral model that reasonably approximates a host of different shadow models. The immediate goal of the present paper is to define and study such a generic neutral model. There are a series of steps involved in proving the value of this generic model, such as comparing it with many different shadow models, discerning how its behavior depends on crucial model parameters, exploring how adaptation alters its behavior, and connecting it with related theoretical and empirical work. The present paper takes the first step in this process by comparing the generic neutral model with a simple evolutionary system and its neutral shadow.

2 Evolutionary Activity Statistics

Evolutionary activity statistics are computed from data obtained by observing an evolving system. We view an evolving system as a population of components participating in a cycle of birth, life and death, with each component largely determined by inherited traits. Birth and mutation introduce innovations into the population. Adaptive innovations persist in the population because of their beneficial effects for component survival or reproduction, and non-adaptive innovations either disappear or persist passively.

The idea behind evolutionary activity is to identify innovations that persist and continue to be significant. Counters are attached to components for bookkeeping purposes, to update each component's current activity as the component persists. If the components are passed along during reproduction, the corresponding counters are inherited with the components, maintaining an increasing count for an entire lineage. Previous work has studied components on the level of individual alleles [1] as well as genotypes [2] and taxonomic families [2]. For simplicity, here we restrict our attention to entire genotypes.

To measure activity contributions we attach a counter to each component of the system, $a_i(t)$, where i labels the component and t labels time. A component's activity increases over time as follows, $a_i(t) = \sum_{k \leq t} \Delta_i(k)$, where $\Delta_i(k)$ is the activity increment for component i at time k. Various activity incrementation functions $\Delta_i(t)$ can be used, depending on the nature of the components and the purposes at hand. Since in the present context more adaptive genotypes tend to persist longer, it's natural to measure a component's contribution to the system's evolutionary activity simply by its age. So we choose an activity incrementation function that increases a component's activity counter by one unit for each time step that it exists:

$$\Delta_i(t) = \begin{cases} 1 & \text{if component } i \text{ exists at } t \\ 0 & \text{otherwise} \end{cases} \quad . \tag{1}$$

Though there are ways to refine this simple counting method [1, 2], this version facilitates direct comparison with many other systems.

Now, we can define various statistics based on the components in a system and their activity counters. There are various ways to quantify diversity and evolutionary activity (e.g., [1, 2]). Here we choose statistics that make it easy to compare diversity and evolutionary activity across a wide variety of evolving systems.

Perhaps the simplest statistic is the system's *diversity*, $D(t)$, which is the number of components present at time t,

$$D(t) = \#\{i : a_i(t) > 0\} \ , \tag{2}$$

where $\#\{\cdot\}$ denotes set cardinality. A measure of the continual adaptive success of the components in the system at a given time is provided by the *total cumulative evolutionary activity*, $A(t)$, which simply sums the evolutionary activity of all the components at a given time:

$$A(t) = \sum_i a_i(t) \ . \tag{3}$$

Then, the cumulative activity per component, or *mean cumulative evolutionary activity*, $\bar{A}(t)$, is simply the cumulative evolutionary activity $A(t)$ divided by the diversity $D(t)$:

$$\bar{A}(t) = \frac{A(t)}{D(t)} \ . \tag{4}$$

In this paper we usually refer to mean cumulative evolutionary activity simply as "mean activity" or, even more simply, as "activity."

Activity statistics most clearly reflect a system's *adaptive* evolutionary dynamics after they have been normalized by a "neutral" model [2], to screen off the contribution of non-adaptive or maladaptive genotypes. This normalization can be accomplished in various ways. Here, we measure typical evolutionary activity in a neutral model and then divide the difference between activity observed in the evolving system and activity in the neutral model by neutral model activity. We call this divergence between observed and neutral activity, expressed as a fraction of neutral activity, a system's *excess* activity:

$$\bar{A}_{\text{excess}} = \frac{\bar{A}_{\text{observed}} - \bar{A}_{\text{neutral}}}{\bar{A}_{\text{neutral}}} \ . \tag{5}$$

So, for example, if observed activity is ten times higher than neutral activity, then $\bar{A}_{\text{excess}} = 9$; also, excess activity of any neutral model is nil. Excess mean cumulative activity is our measure of the amount of the persistence of extant genotypes that can be attributed to the adaptive success of those genotypes. It should be noted that excess activity does *not* measure how well components are adapted to the environment. There is no particular correlation between excess activity and fitness. If excess activity becomes significantly higher, that does not mean that the components in the system have become significantly better adapted; rather, it means that there has been a significant increase in the history or activity (here, persistence) of components that is due to their adaptive value.

It should be emphasized that mean evolutionary activity is an *extensive* rather than *intensive* measure. That is, it measures the *extent* of the continual adaptive success of the components in a system, rather than the *intensity* with which new adaptive components are being created. Thus, if a set of adaptive components continue to persist, the extent of their continual adaptive success increases over time, even if the intensity with which new adaptations are being created falls to nil. On the other hand, if evolution is continually creating new adaptations and destroying older adaptive components, the intensity of adaptive evolution will be positive; but the extent of the continual adaptive success in the system will be very low if none of those adaptations persist for very long. (Component activity counters can also yield measures of the intensity of adaptive evolution [1, 2].)

3 Packard's Line Model

Packard's Bugs simulation is a series of models originated by Norman Packard [7, 1]. Packard's simulation is designed to be a very simple model of the evolution of sensory-motor strategies. It consists of agents sensing the resources in their local environment, moving as a function of what they sense, ingesting the resources they find, and reproducing or dying as a function of their internal resource levels. The model's spatial structure is a grid of sites with periodic boundary conditions, i.e., a toroidal lattice. The version of Packard's model that we study here has an especially simple resource distribution: a thin continuous strip, one cell in width, that wraps entirely around the world, with all the other sites in the world entirely devoid of resources. We call this the "Line" model. Resources are immediately replenished at a site whenever they are consumed. The agents constantly extract resources and expend them by living and reproducing. Agents ingest all of the resources (if any) found at their current location and store them internally. Agents expend resources at each time step by "paying" (constant) "existence taxes" and "movement taxes" (variable, proportional to distance moved). If an agent's internal resource supply drops to zero, it dies and disappears from the world.

Each agent moves each time step as dictated by its genetically encoded sensory-motor map: a table of behavior rules of the form IF (environment j sensed) THEN (do behavior k). An agent receives sensory information about the resources (but not the other agents) in the von Neumann neighborhood of five sites centered on its present location in the lattice. There are exactly 4 detectable local environments: those detected by agents either on the resource strip, immediately to the strip's left or right, or anywhere else. Each behavior k is a jump vector between one and fifteen sites in any one of the eight compass directions. Thus, an agent's genotype, i.e., its sensory-motor map, is just a lookup table of sensory-motor rules. But the space in which adaptation occurs is fairly large, consisting of $120^4 \approx 10^8$ distinct possible genotypes.

An agent reproduces (asexually, without recombination) if its resource reservoir exceeds a certain threshold. The parent produces one child, which starts life with half of its parent's resource supply. The child inherits its parent's sensory-motor map, except that mutations may replace the behaviors linked to some sensory states with randomly chosen behaviors. A time step in the simulation cycles through the entire population and has each agent, in turn, complete the following sequence of events: sense its present von Neumann neighborhood, move to the new location dictated by its sensory-motor map unless that site is already occupied, in which case randomly walk to the first unoccupied site, consume any resources found at its new location, expend resources to cover existence and movement taxes, and then, if its resource reservoir is high enough or empty, either reproduce or die.

A given simulation starts with randomly distributed agents containing randomly chosen sensory-motor strategies. The model contains no *a priori* fitness function, as Packard [7] has emphasized. Agents with maladaptive strategies tend to find few resources and thus to die, taking their sensory-motor genes with them; by contrast, agents with adaptive strategies tend to find sufficient resources to reproduce, spreading their sensory-motor strategies (with some mutations) through the population. The main adaptations that occur in the Line model are learning how to stay on the resource strip and learning to do so in step with the other bugs on the strip (i.e., precisely meshing with the "flock" of other bugs on the line). Another, secondary adaptation is optimizing the jump size on the strip (smaller jumps are better). Furthermore, there is a slight adaptive advantage to learning how to get back on the strip when immediately adjacent to it.

4 A Neutral Shadow

The crucial property of a "neutral shadow" of a model with implicit genotype dynamics like Packard's models is that shadow system's evolutionary dynamics are like the normal model *except* that a shadow genotype's presence or concentration or longevity cannot be due to its adaptive significance—for it has no adaptive significance. The shadow model of a Packard Line model consists of a population of only nominal "bugs" with only nominal "genotypes." A shadow "bug" has no spatial location and it cannot ingest resources or interact with other "bugs." All it ever does is come into existence, perhaps reproduce (perhaps many times), and go out of existence; its only properties are its genotype and the times of its birth, reproductions (if any), and death.

Each model run has its own corresponding neutral shadow run. The neutral shadow's birth and death events and mutation rate are directly copied from those in the Line simulation. When some creature is born in the Line simulation a shadow parent is chosen at random (with equal probability) from the shadow population to reproduce. The new shadow child inherits its parent's genotype unless a mutation gives the child a new genotype. When some creature dies in the Line simulation a "creature" is chosen at random from the shadow population and killed. Thus, all selection in the neutral shadow is random.

The evolutionary dynamics in a neutral shadow is a neutral diffusion in genotype space. Genotypes arise and go extinct, and their concentrations change over time, but the genotype dynamics are at best weakly linked to adaptation through the birth and death rates determined by adaptation in the Line model. When adaptive genotypes are evolving in a Line simulation, one would expect their genotype activity levels to be significantly higher than those in the corresponding neutral shadows. For, although individuals in the Line model and its neutral shadow have the same birth, reproduction, and death rates (indeed, all model parameters are identical), in the Line model natural selection can preferentially cull poorly adapted genotypes and preserve well adapted genotypes while the selective force in the neutral shadow is entirely random. The difference between the activity levels in the Line model and its neutral shadow shows how much natural selection affects the activity counts in the Line model.

5 A Generic Neutral Model

The generic model of neutral genotype evolution consists of a population of individuals that reproduce and die in a fixed genotype space. The genotype space consists of some number of loci at each of which some number of alleles are segregating. Parameters that need to be specified in the generic neutral model are N, the size of the population of individuals, r, the reproduction rate (the number of individuals that die and reproduce per time step), l, the number of loci, a, the number of possible alleles per locus, m_l, the probability that the allele at a given locus will be mutated when an individual is born. (The probability that an offspring will have a mutation somewhere in its genome, i.e., the mutation rate per individual, is $m_i = 1 - (1 - m_l)^l$.) The parameters remain fixed during any given instance of the model, and together they determine the model's behavior. The genotype space is a hypercube of dimension l and size a^l (number of possible genotypes), with each location in this space corresponding to a given genotype. The current state of the model is given by the distribution of N individuals in genotype space. Over time, the population wanders through the space stochastically, spreading and clustering at random locations.

The individuals in the initial population are assigned genotypes at random. Time is discrete, and moves forward each time step by iterating the following two-step algorithm: (1) r individuals (selected at random, with replacement) each produce a child that is genetically identical to itself except for mutations. Mutant alleles are chosen at random from the set of possible alleles. (2) r individuals (selected at random, without replacement) die and are removed from the population and are replaced by the r children produced at step (1).

This neutral model is somewhat disanalogous to those systems in which some of the generic model parameters are variable. E.g., in Tierra [8] the number of loci is variable; indeed, it is not clear exactly what to count as a locus in Tierra. In addition, population size and reproduction rate vary over time in many artificial models of evolution, such as Echo [4] and Packard's Bugs models. Still, the neutral model might apply reasonably well to these systems if the relevant neutral model parameters are set to plausible corresponding values. For the comparisons here we set N and r to the mean observed value of the corresponding parameter in the Line model. A main goal of this study is to assess the usefulness of the generic neutral model under such an approximation.

6 Experimental Methods

We observed the behavior of the Line model, the neutral shadow model, and the generic neutral model across the mutation rate spectrum (varied on a log scale). All Line model simulations were started with a randomly initialized populations of 500 individuals. We did at least 10 runs at every mutation rate we studied, varying the simulation time between 5×10^5 and 5×10^7 depending on the mutation rate. The transient time is longer at lower mutation rates, and we aimed to have long-enough simulations to minimize the variance in the time series data. The parameters for the generic neutral model were set to correspond to the Line model (four loci and 120 alleles per locus). We observed average population size N and reproduction rate r from each Line simulation and set corresponding parameter values in generic model runs.

We dumped 5000 data points in each simulation, so the time interval between data dumps varied with run length. We tested shorter and shorter data-dumping frequencies until the statistics we were observing converged. In the generic neutral model evolutionary activity was calculated continuously, so the exact activity value could be recorded in each data dump. But in the Line and neutral shadow models genotype data was only *sampled* at each data dump. So, for simplicity, we assumed that a genotype that first appeared at a certain time arose immediately after the previous data dump. This procedure loses all information about short-lived genotypes that arose and went extinct between data dumps, and it significantly overestimates the age of short-lived genotypes that appear in only a few data dumps. To minimize this bias, we shortened the data dumping time at higher mutation rates.

7 Results

Figure 1 shows diversity and activity time series data for a typical run with a mutation rate of $m_i = 0.19$ and a simulation time of 5×10^6 time steps. Note that the diversity in the Line model is on average about one quarter of that in the neutral models. The diversity for the two neutral models are indistinguishable. The activity time series for the Line simulation is very different from the neutral models. First, it is much higher, indicating much longer lived genotypes. It also has larger-magnitude events like the steep drops of activity indicating the extinction of one or more long-lived genotype. Note that up to time 5×10^5 the shadow and generic activity differ. This reflects the coupling of shadow and Line model, and is mainly due to Line model reproduction rate fluctuations and population fluctuations which are transmitted to the shadow neutral.

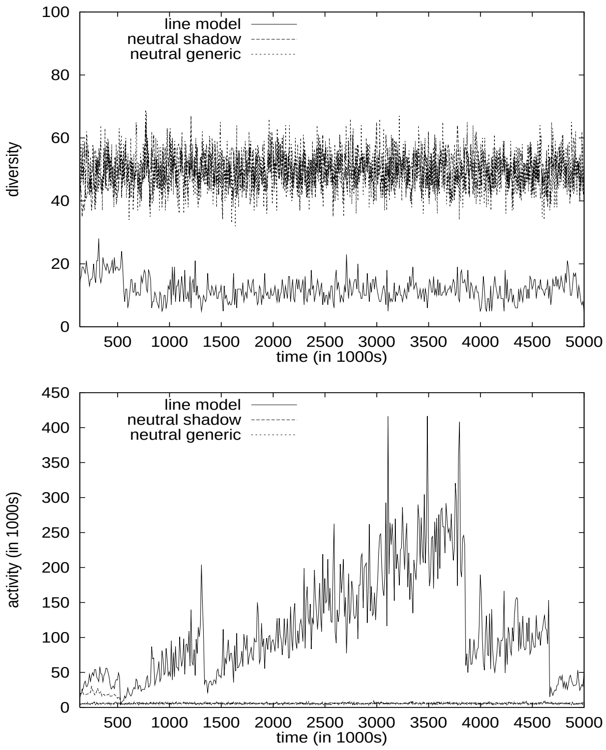

Figure 1: Time series of diversity, D, and mean evolutionary activity, A, in a typical 5,000,000 time-step run with mutation rate per individual $m_i = 0.19$. Note that the diversity and evolutionary activity data for the generic and shadow neutral models are quite similar, so similar that they are indistinguishable in this figure. On the other hand, the Line model shows significantly lower diversity and significantly higher evolutionary activity than the neutral models.

Figure 2 shows the time average of diversity for the Line model and for both neutral models plotted as a function of the mutation rate per individual, m_i. As illustrated in Fig. 1, diversity is everywhere lower in the Line model than the neutral models except for very low mutation rates (when diversity reaches its floor of unity, $D = 1$) and very high mutation rates (when diversity reaches its ceiling of the population size, $D = N$). Note how well the diversities of the two neutral models match over the whole mutation rate spectrum. Note also that the significantly higher diversity in the neutral models compared with the Line model (at those mutation rates away from the diversity floor and ceiling), which is the expected effect of the

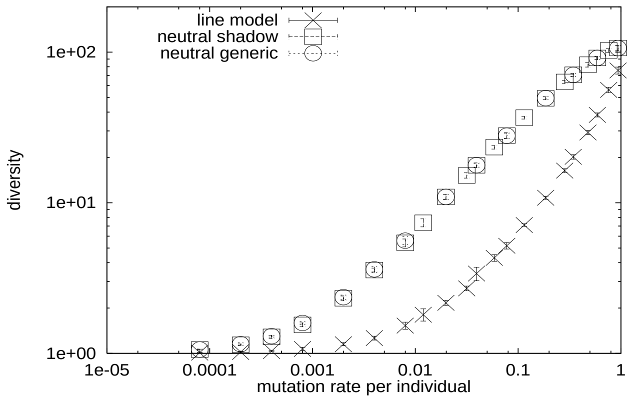

Figure 2: Time average of diversity, D, as a function of mutation rate per individual, m_i, for the Line model, its neutral shadow, and the generic neutral model. At each mutation rate the mean (and standard deviation) of at least ten runs is shown. Note the striking similarity of the two neutral models. Note also the strikingly lower diversity of the Line model, for all but extreme mutation rates (except at very high and very low mutation rates, when the system hits the diversity ceiling and floor).

neutral models' random selection.

Figure 3 shows the time average of mean cumulative evolutionary activity for the Line model and for both neutral models, plotted as a function of the mutation rate per individual, m_i. The longest runs we did were 5×10^7 time steps, so this is the maximum activity we could record. This explains why activity in the Line model levels off for $m_i \leq 0.05$; it is an artifact of too-short simulation times. We are exploring the Line activity for $m_i \leq 0.05$ with longer simulation times in current work.

Note that the average activity of the two neutral models lies well within the standard deviation. We found in other work that neutral model activity is inversely proportional to reproduction rate. We would therefore expect the varying reproduction rate in shadow models would make shadow activity higher than generic neutral model activity, even though the generic model reproduction rate is set to the average reproduction rate in the shadow run. And, in fact, we do observe that generic neutral model activity is generally somewhat lower than shadow model activity. Nonetheless, their correspondence is very close.

Figure 3 also shows that neutral model activity's dependence on mutation rate approximately fits a power

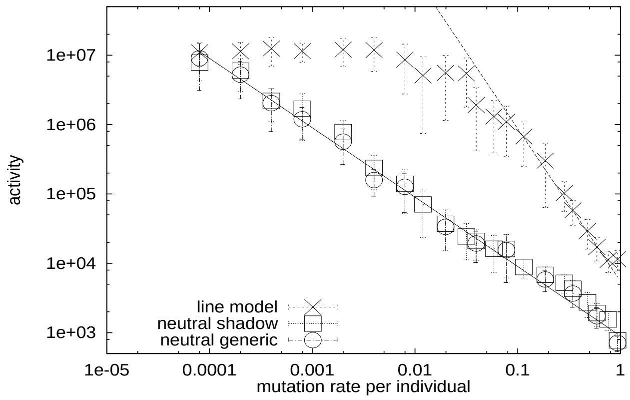

Figure 3: Time average of mean evolutionary activity, $\bar{A}$, as a function of mutation rate per individual, m_i, for the Line model, its neutral shadow, and the generic neutral model. At each mutation rate the mean (and standard deviation) of at least ten runs is shown. Activity in the Line model at lower mutation rates is artificially lowered by the simulation time being too short. Note the striking similarity of the two neutral models and the strikingly higher activity for the Line model (at those mutation rates with long enough simulations). To show how evolutionary activity in the neutral models and in the Line model approximately fit power laws, we have superimposed lines of slope -1.0 and -2.2 over the neutral and normal data, respectively.

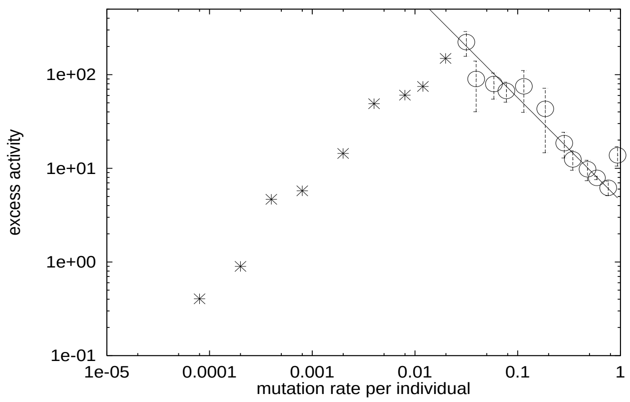

Figure 4: Time average of excess mean evolutionary activity, $\bar{A}_{\text{excess}}$, as a function of mutation rate per individual, m_i, for the Line model. Circles indicate where simulation times were long enough for activity statistics to converge. Data shown are mean (and standard deviation) of at least ten runs. A superimposed a line shows that excess activity approximately fits a power law, with $\alpha_{\text{excess}} = 1.1 \pm 0.2$.

law:

$$\bar{A} \propto m_i^{-\alpha} \ , \tag{6}$$

with $\alpha_{\text{neutral}} = 1.0 \pm 0.1$. (Error bounds on all exponents reported here are based on standard deviations of observed activity.) In addition, for $m_i \geq 0.05$, where activity is not affected by simulation time, Line model activity fits the same power law, with $\alpha_{\text{Line}} = 2.2 \pm 0.2$.

Figure 4 graphs excess activity $\bar{A}_{\text{excess}}$ in the Line model as a function of mutation rate. This initial picture of excess activity's dependence on mutation rate is understandably incomplete and noisy. For one thing, we have good data for excess activity measurements only for mutation rates $m_i \geq 0.05$. Furthermore, we typically observed only ten runs at each mutation rate; collecting much more data should reduce noise considerably. Nevertheless, at those mutation rates for which we do have good data, excess activity $\bar{A}_{\text{excess}}$ shows the power-law dependence of Eq. 6, with $\alpha_{\text{excess}} = 1.1 \pm 0.2$. There is one significant deviation from this law at $m_i \approx 1$, discussed below.

8 Discussion

The generic neutral model shows clear evidence of suitably approximating the behavior of special-purpose neutral shadows. In particular, the shadow and generic neutral models show remarkable similarity in how diversity depends on the mutation rate, and the match of their activity response to diversity is also quite significant, especially when the variance in activity is considered. To be sure, the generic neutral model has only passed one preliminary test, and its final confirmation will come only if it shows similar success at approximating a range of other neutral shadows—a subject of current work.

The other results presented here reveal why the generic neutral model is important, because they underscore that excess evolutionary activity is appropriate for quantifying the extent to which evolution creates continual adaptive success. Just as one would expect, the neutral models have higher diversity than the Line model. Comparing neutral diversity with diversity in the Line model (Figure 2) dramatizes how much random selection of genotypes diversifies a system compared with natural selection's preference for conserving well adapted genotypes. Furthermore, again as one would expect, mean genotypic activity (which in the present context is equivalent to persistence) is significantly lower in the neutral models. The neutral activity in Fig. 3 shows precisely how much of the Bug's model activity is not due to the genotype's adaptive success, and this allows us to compute excess activity, $\bar{A}_{\text{excess}}$, the difference between observed and neutral activity, expressed in proportion to neutral activity.

The significant rise in excess activity at $mu_i \approx 1$, clearly evident in Figure 4, bears some discussion because it might seem surprising or even anomalous. At those mutation rates at which our statistics converge, excess activity generally seems to have a power-law dependence on mutation rate, but there is one significant exception to this pattern, at $mu_i \approx 1$. In fact, though, this exception proves the rule, i.e., it *confirms* that excess activity measures continual adaptive success. The most crucial locus for success in the Line model is the Line locus itself. Any agent that can't stay on the Line is doomed to die. So, if there is a mutation at the Line locus, then the child will almost certainly have lost all information about how to stay on the Line. As the mutation rate increases, more and more children loose the information about how to stay on the Line. Therefore, there is less competition for space on the Line and all agents on the Line live longer (the Line model enforces no finite lifetime on the agents). In other words, the potentially indefinite Line agent lifetime give a significant adaptive advantage to founder genotypes, whatever they happen to be, but this adaptive advantage occurs only at $mu_i \approx 1$.

How can we explain evolutionary activity's power-law dependence on mutation rate (Eq. 6). For the neutral models we can explain this law with an argument of Kimura [5]. Kimura has calculated the time it takes on average for a new neutral mutant gene to reach fixation during neutral evolution, on the assumption that genes get substituted one after another and not at the same time. Kimura finds that the average time to reach fixation can be described by two different time scales. If we neglect all mutants that go extinct before they reach fixation, the first time scale—the time it takes on average for a neutral mutant to spread throughout the population—is proportional basically to the population size, N. The second time scale—the time it takes on average for such a mutant gene to occur in the population—is proportional to the number of mutations that occur, which is proportional to m_i^{-1}. For low mutation rates Kimura's discussion also applies to genotype substitution in our neutral models. A new mutant gene corresponds in our neutral models to a new mutant genotype. Kimura's assumption that genes are substituted one after another corresponds in our neutral models to the assumption that genotypes are substituted one after another, and this assumption holds when the mutation rate is not too high. For low enough mutation rates only the second time scale—the time it takes for new mutants to occur—is relevant; the other time scale is basically constant (because population is basically constant) and becomes negligible. Activity (here, mean lifetime of genotypes in a population) will therefore show the same relationship with mutation rate as the second time scale, i.e., activity will be proportional to m_i^{-1}, thus explaining why $\alpha_{\text{neutral}} = -1.0$.

It is also striking that the Line model excess activity shows evidence (at least where we have good data) of a simple power-law dependence on mutation rate (except at $mu_i \approx 1$, of course).[1] One would expect Bug model excess activity to be inversely proportional to mutation rate—at least, if the mutation rate is not *too* low—for mutations cause the system to "forget" adaptations. But none of this explains why $\alpha_{\text{Line}} = 2.2$ and $\alpha_{\text{excess}} = 1.1$. These precise quantitative results cry out for explanation and raises a number of interesting questions, all of which are straightforward to answer: Does the same law hold in many, or even any, other models? Is the same exponent found? Exactly what explains the values of the exponents? Answering such questions would significantly advance our quantitative understanding of the foundations of adaptive evolution—all made possible by a feasible generic neutral model of the evolution of genotypes.

Acknowledgements. For helpful discussion thanks to Norman Packard, Tad Shannon, and Emile Snyder. For helpful comments on the manuscript, thanks to the GECCO reviewers. Thanks to Emile Snyder for help with the software for the Line model, the generic neutral model, and the evolutionary activity statistics. MAB also thanks the Santa Fe Institute, for support and hospitality while some of this work was completed.

References

[1] Bedau, M. A. and N. H. Packard. 1992. Measurement of evolutionary activity, teleology, and life. In C. Langton, C. Taylor, J. D. Farmer, S. Rasmussen, eds., *Artificial Life II* (pp. 431–461). Redwood City, Calif.: Addison-Wesley.

[2] Bedau, M. A., E. Snyder, N. Packard, 1998. A classification of long-term evolutionary dynamics. In C. Adami, R. Belew, H. Kitano, and C. Taylor, eds., *Artificial Life VI* (pp. 228–237). MIT Press/Bradford Books.

[3] Gould, S. J., Lewontin, R. C. 1979. The spandrals of San Marco and the Panglossian paradigm: a critique of the adaptationist programme. *Proceedings of the Royal Society B* 205: 581–598.

[4] Holland, J. H. 1992. *Adaptation in natural and artificial systems*, 2nd edition, MIT Press/Bradford Books.

[5] Kimura, M. 1983. *Neutral theory of molecular evolution*. Cambridge University Press.

[6] Maynard Smith, J. 1975. *The theory of evolution*, 3rd edition. New York: Penguin.

[7] Packard, N. 1989. Intrinsic adaptation in a simple model for evolution. In C. Langton, ed., *Artificial Life* (pp. 141–155). Redwood City, Calif.: Addison-Wesley. Addison-Wesley.

[8] Ray, T. S. 1992. An approach to the synthesis of life. In C. Langton, C. Taylor, J. D. Farmer, S. Rasmussen, eds., *Artificial Life II*, Addison-Wesley.

[1] In a related vein, notice that the *intensity* of adaptive evolution should *increase* with the mutation rate (up to an error threshold). One might become confused about why $\bar{A}_{\text{excess}}$ is *inversely* proportional to mutation rate if one forgets that $\bar{A}_{\text{excess}}$ measures the *extent* (not the intensity) of adaptive evolution.