

# A Generic Neutral Model for Quantitative Comparison of Genotypic Evolutionary Activity

Andreas Rechtsteiner<sup>1</sup> and Mark A. Bedau<sup>2</sup>

<sup>1</sup> Systems Science Ph.D. Program, Portland State University  
P.O. Box 751, Portland OR 97202, USA

andreas@sysc.pdx.edu, <http://www.sysc.pdx.edu/Alife/aendy.html>

<sup>2</sup> Reed College, 3203 SE Woodstock Blvd, Portland OR 97202, USA  
mab@reed.edu, <http://www.reed.edu/~mab>

**Abstract.** We use a new general-purpose model of neutral evolution of genotypes to make quantitative comparisons of diversity and adaptive evolutionary activity as a function of mutation rate among two versions of Packard's Bugs model and their neutral shadows. Comparing diversity and evolutionary activity of all these models across the mutation rate spectrum shows that the generic neutral model may have broad applicability in discovering quantitative laws involving adaptive evolutionary activity in different evolving systems.

## 1 The Need for a Generic Neutral Model

Adaptive evolution is thought to produce much of the order and functionality evident in complex systems [9, 7, 5], but it is often difficult to distinguish adaptive change from other evolutionary phenomena such as random genetic drift and architectural necessity [6, 10], and some even question whether adaptations can be objectively identified at all [6]. Recent progress on identifying adaptive evolutionary phenomena includes Bedau and Packard's statistical methods for measuring adaptive evolutionary activity. Here, we apply these methods to the problem of determining how adaptive evolutionary activity depends on mutation rate. Our ultimate aim is to develop methods for objectively identifying and measuring adaptive evolutionary activity in all evolutionary systems, both natural and artificial, so that we can seek universal laws of adaptive evolutionary activity. Here, we test such a method, applied at the level of whole genotypes, in the context of two simple models of sensory-motor evolution. But the same method can be applied at other levels of analysis in other evolutionary systems. The ultimate significance of this work comes from the possibility of quantitatively comparing evolutionary adaptations across all evolving systems.

The centerpiece of our method is Bedau and Packard's evolutionary activity statistics. (We also measure system diversity,  $D$ , which is simply the number of different genotypes present in a system at a given time.) Detailed definitions and motivations for evolutionary activity statistics are readily available elsewhere [2, 1, 3, 4, 13]. Evolutionary activity statistics aim to identify evolutionary innovations (here, new genotypes) that persist and continue to play a significant role

in a system because of their adaptive value. These statistics fall into two broad classes: those reflecting evolutionary activity’s *extent* and those reflecting its *intensity*. Here, we attend only to the *extent* of evolutionary activity, measuring it with *mean cumulative evolutionary activity* (sometimes simply called “activity”),  $\bar{A}$ , which in the present context is operationalized as the mean age of the genotypes present in a system at a given time. So, in this context, the higher a system’s mean activity, the higher the mean age of the system’s genotypes, which means the greater the continual adaptive success of those genotypes. Intuitively, the *extent* of evolutionary activity concerns how much adaptive structure is present in a system; one might refer to this as the continual adaptive success of the system’s components. By contrast, the *intensity* of evolutionary activity reflects the rate at which new adaptive structure is being created. The extent and intensity of adaptive evolutionary activity are independent. For example, if a population of highly adaptive genotypes persist indefinitely without changing and no new genotypes invade the system, then the extent of evolutionary activity is positive and perhaps grow over time, but the intensity of evolutionary activity falls to nil.

To ensure that evolutionary activity statistics reflect the adaptive success of the genotypes and not non-adaptive evolutionary forces like chance and necessity, one must use non-adaptive evolutionary systems called “neutral models” as null hypotheses. That is, one must screen off the effects of non-adaptive evolutionary forces like chance by comparing the evolutionary dynamics observed in target evolutionary systems with those observed in analogous neutral models. Such neutral analogues have heretofore been constructed by crafting systems that “shadow” the target 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 [3, 4]. Since neutral shadows are tailored to target systems, they sharply show the target systems’ deviation from the no-adaptation null hypothesis. But neutral shadows have significant drawbacks, too, for studying a new target system involves constructing and studying a new neutral shadow, and it is vexing to make meaningful quantitative comparisons among different tailor-made neutral shadows.

The obvious way to solve these problems is to create a *generic* neutral model—one neutral model that can approximate many different neutral shadows. The immediate goal of this paper is to define such a generic neutral model and test its usefulness for quantifying evolutionary activity across different systems. We pursue this goal by comparing the generic neutral model with two simple evolutionary systems and their neutral shadows.

## 2 The Models

**Packard’s Line and Block Models.** The Bugs simulation is a series of models originated by Norman Packard [11, 2] and subsequently modified in various ways. 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 resource distributions studied here take two forms:

- Line:** 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;
- Block:** a square block of resources, 15 cells on a side, with all other sites in the world entirely devoid of resources.

In each case, 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. In the Line world, 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. In the Block world, there are exactly 14 detectable local environments: those detected by agents either just on one of the four edges, or just off one of the four edges, or in one of the four corners, or in the middle of the block, 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  $120^4 \approx 10^8$  and  $120^{14} \approx 10^{29}$  distinct possible genotypes in the Line and Block worlds, respectively. 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 given simulation starts with randomly distributed agents containing randomly chosen sensory-motor strategies. The model contains no *a priori* fitness function, as Packard [11] 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. In the Line world, the main adaptations that occur 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., 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. In the Block world, as in the Line world, one adaptive pressure is to “flock” along with the other bugs, so as to minimize the changes of getting bumped into the resource desert. But the basic adaptive strategy needed to survive on a resource block is to move in a given direction and speed when in the middle of the block until you detect the edge of the block, and then to jump back in the opposite direction into the middle of the block. A subpopulation of bugs following the same strategy will form a flow that rolls across the block and reflects off its edge. Since all bugs in the Line world must flock in step not to bump each other off the resource strip, the fitness landscape in the Line world has relatively narrow peaks. By contrast, bugs that reflect different distances off the edge can co-exist, and different subpopulations can form along different edges in a resource block, in effect filling different niches, so the Block’s fitness landscape allows for more diversity and thus contains relatively broad peaks.

**Neutral Shadows for Packard’s Line and Block Models.** The crucial property of a “neutral shadow” of a model with emergent genotype dynamics is that the shadow system’s evolutionary dynamics are like its target model *except* that a shadow genotype’s activity cannot be due to its adaptive significance—for it has no adaptive significance. The neutral shadow a Packard Bugs model consists of a population of nominal “bugs” with 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 often), and go out of existence; its only properties are its genotype and the times of its birth, reproductions (if any), and death.

Each neutral shadow run corresponds to a specific Line or Block model run. The neutral shadow’s birth and death events and mutation rate are directly copied from those in the target run. When some creature is born in the target run, 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 or Block run 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 or Block model. When adaptive genotypes are evolving in a Bugs run, one would expect their genotype activity levels to be significantly higher than those in the corresponding neutral shadows. For, although individuals in the Bugs model and its neutral shadow have the same birth, reproduction, and death rates, and their mutation rates are the same, in the Bugs model natural selection can 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 Bugs and its neutral shadow shows how much natural selection affects Bugs activity.

**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 is defined by 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,  $m_i$ , the probability that the allele at a given locus will be mutated when an individual is born. (The probability that an offspring will have mutation somewhere in its genome, i.e., the mutation rate per individual is  $m_i = 1 - (1 - m_l)^l$ .) The parameters together determine the model's generic 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. The population wanders through the space stochastically, spreading and clustering at random.

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 does not closely correspond to those systems in which some of the generic model parameters are variable. E.g., in Tierra [12] 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 [7] 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 Bugs model. One of the goals of this study is to assess the usefulness of the generic neutral model under such an approximation.

### 3 Experimental Methods

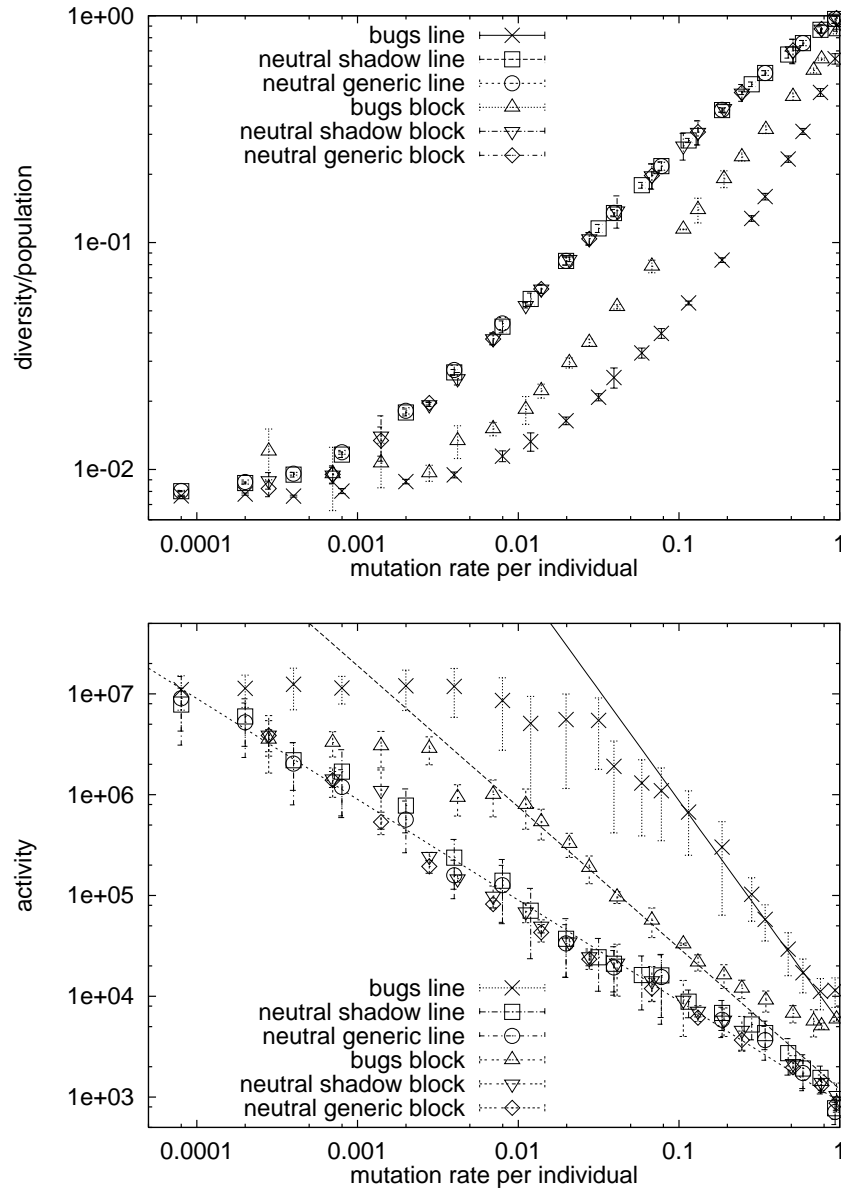
We observed the behavior of the Line and Block models, the Line and Block neutral shadows, and the Line and Block generic neutral models across the mutation rate spectrum (varied on a log scale). All Packard model simulations were started with a randomly initialized populations of 500 individuals. We did at least 10 runs at every mutation rate in each model. We varied the simulation time between  $5 \times 10^5$  and  $5 \times 10^7$  depending on the mutation rate. The transient time is longer at lower mutation rates, and we aimed to have simulations that

were long enough to minimize variance due simply to simulation time. The parameters for the generic neutral model were set to correspond to the Line model (four loci and 120 alleles per locus) and the Block model (fourteen loci and 120 alleles per locus). We determined average population size  $N$  and reproduction rate  $r$  from each Line or Block model run 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 made sure data dumping frequencies did not influence our results. 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 Bugs 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. This bias was minimized by using shorter simulation time and data dumping intervals for high mutation rates.

## 4 Results and Discussion

The top of Figure 1 shows the time average of the diversity normalized by the time average of the population size,  $\frac{D}{N}$ , for the Line and Block models, and for shadow and generic neutral models of each, as a function of mutation rate,  $m_i$ . At very low mutation rates diversity levels off. This diversity floor is an artifact of the finite population size. Indefinitely larger populations would indefinitely lower this floor. These diversity data at mutation rates above the diversity floor show three salient results: First, the dependence of normalized diversity on mutation rate for all four neutral models is strikingly similar. This provides evidence that the generic neutral model is a good approximation of the neutral shadows. Second, the normalized diversity of both Line and Block models is strikingly lower than that for the neutral models—the expected result of natural selection in the Line and Block models versus the random selection in the neutral models. Third, the normalized diversity of the Line model is strikingly lower than that for the Block model. This can be explained by the different fitness landscapes in the Block and Line world. Each fitness landscape has several local peaks, but the Block peaks are broader than the Line peaks and broader peaks support a more diverse population.

The bottom of Figure 1 shows the time average of activity,  $\bar{A}$ , as a function of mutation rate,  $m_i$ , for the Line and Block models and for shadow and generic neutral models of each. The activity ceiling in the Line and Block models at lower mutation rates is partly an artifact of simulation time; a genotype’s observed age cannot exceed simulation time and the longest runs we did lasted  $5 \times 10^7$  time steps. Longer simulations would raise the observed activity values at very low mutation rates. Where unaffected by the activity ceiling artifact, the time-averaged activity data show three significant results.



**Fig. 1.** Above: Time average of diversity (normalized by dividing by time average of average population size),  $\frac{D}{N}$ , as a function of mutation rate per individual,  $m_i$ , for the Line and Block models, their neutral shadows, and generic neutral models for them. Below: Time average of evolutionary activity,  $\bar{A}$ , as a function of mutation rate per individual,  $m_i$ , for the Line and Block models, their neutral shadows, and generic neutral model for them, along with lines showing the power laws which approximately fit these data (at certain mutation rates). In both graphs, error bars indicate standard deviations of time averages computed from at least ten runs per mutation rate.

First, activity in all the neutral models is quite similar, especially considering that we did not account for differences in population size and reproduction rate in the Line and Block worlds. Furthermore, activity’s dependence on mutation rate approximately fits a power law of the form  $\bar{A} \propto m_i^{-\alpha}$ , where  $\alpha_{\text{neutral}} = 1.0 \pm 0.1$  (all error bounds are standard deviations). This power law can be explained by adapting an argument of Kimura [8]. Kimura found that the average time it takes for a new mutant gene to reach fixation during neutral evolution, on the assumption that genes get substituted one after another and not at the same time, can be described by two time scales. 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. 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  $m_i^{-1}$ . 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. So, for low mutation rates Kimura’s discussion applies equally well to genotype substitution in our neutral models, with Kimura’s new mutant gene corresponding to our new mutant genotype. For low enough mutation rates only the second time scale is relevant; the other time scale is basically constant (because population is basically constant) and becomes negligible. So, evolutionary activity, which corresponds to the mean lifetime of genotypes, will be proportional to the second time scale, i.e., to  $m_i^{-1}$ .

The second significant result in the activity data is that the magnitude of activity is significantly higher in both Line and Block than in the neutral models, as is the slope of activity’s dependence on mutation rate. Lower neutral model activity can be explained by the lack of adaptation in the neutral models. Random selection in neutral models does not preferentially preserve well-adapted genotypes but natural selection in the Line and Block models does, so Line and Block activity is expected to exceed neutral model activities.

The third significant activity result is that the magnitude of the activity is significantly higher in the Line model than in the Block model, as is the slope of activity’s dependence on mutation rate. Although the data are somewhat ambiguous, at relatively higher mutation rates the Line and Block activity’s dependence on mutation rate might approximately fit power laws with  $\alpha_{\text{Line}} = 2.2 \pm 0.2$  and  $\alpha_{\text{Block}} = 1.4 \pm 0.1$ . Our activity data at lower mutation rates are clearly affected by the activity ceiling artifact; we expect that the slope of the Line and Block activity will fall to  $-1$  at low enough mutation rates, but the activity ceiling prevents us from resolving this here. We expect that the differences in the Line and Block fitness landscapes can explain the differences in activity magnitude and slope in the Line and Block worlds. Evidently, fewer successful genotypes residing on narrow fitness peaks persist significantly longer than more genotypes (therefore lower average population size per genotype) residing on broader peaks. In addition, on narrow fitness peaks more than on broad peaks, the extent of adaptation seems to fall faster as mutation rate rises.



Current work includes resolving the activity's different dependence on mutation rate in the Line and Block models. This project is especially engaging because we have precise quantitative results needing explanation.

## 5 Conclusions

Comparing diversity and evolutionary activity in Packard's Line and Block worlds and in their shadow and generic neutral models, yields a variety of precise and interesting quantitative results. Some are due to the fact that the neutral models are devoid of adaptation, others to the different fitness landscapes in the Line and Block worlds. Both illustrate the power and promise of using neutral models to quantify adaptive evolution in different evolutionary systems.

The absence of adaptation in the neutral models explains their relatively high diversity and low activity, compared with the Line and Block worlds, as well as the lower slope with which activity depends on mutation rate in the neutral models. The differences we observed have three related implications: they confirm the appropriateness of using activity statistics to measure the extent of adaptive structure in an evolving system, thereby confirming the appropriateness of using neutral model activity to measure the amount of activity that can be attributed to adaptation as opposed to other evolutionary forces like chance and necessity, and thereby confirming the importance of the generic neutral model. The generic neutral model closely approximates the behavior of different special-purpose neutral shadows; dependence of diversity and activity on mutation rate is remarkably similar in all of them. Having one simple generic neutral model removes the need to make a new neutral shadow for each evolutionary model and allows us to study the general properties of neutral models in one fell swoop. To be sure, the generic neutral model has so far passed only a preliminary test, and its final confirmation can come only if it successfully approximates many more neutral shadows. Conducting these further tests is a subject of current work, as is discerning the generic neutral model's typical behavior. The results discussed above reveal an important difference between the evolutionary dynamics of neutral and adaptive evolutionary systems, and the generic neutral model is an excellent tool for discerning and understanding this difference.

The magnitude and slope of activity's dependence on mutation rate also reveals the fundamental difference between the Line and Block fitness landscapes. This, too, confirms the appropriateness of measuring the extent of adaptive structure in a system with activity statistics, thereby underscoring the utility of neutral models in general and the generic neutral model in particular. For example, we can measure how much the observed activity reflects the force of adaptive evolution by appropriately normalizing observed activity against the corresponding neutral model, e.g., by dividing observed activity by the corresponding neutral activity. If we call this fraction by which observed activity exceeds neutral (non-adaptive) activity a system's *excess* evolutionary activity [13], then a power-law dependence of activity on mutation rate (over part of the mutation rate spectrum) implies simple power-law dependence of *excess* activ-

ity on mutation rate in the Line and Block models. Such excess activity power laws raise intriguing questions: Does excess activity show similar power laws in a broad class of evolutionary models? If so, what exactly explains the magnitude and exponent in the laws? Answering any of these questions would significantly advance our quantitative understanding of adaptive evolution. And a key tool for facilitating these precise quantitative comparisons is the generic neutral model.

**Acknowledgements.** For helpful discussion, thanks to Norman Packard, Tad Shannon, and Emile Snyder. Thanks to Emile Snyder for software consultation. Thanks to the ECAL-99 reviewers for comments on the manuscript. MAB also thanks the Santa Fe Institute, for support and hospitality while some of this work was completed.

## References

1. Bedau, Mark A. 1995. Three illustrations of artificial life's working hypothesis. In W. Banzhaf and F. Eeckman, (Eds.), *Evolution and biocomputation: computational models of evolution*, (pp. 53–68). Berlin: Springer.
2. Bedau, Mark A., and Norman H. Packard. 1992. Measurement of evolutionary activity, teleology, and life. In C. Langton, C. Taylor, D. Farmer, S. Rasmussen, (Eds.), *Artificial life II* (pp. 431–461). Redwood City CA: Addison-Wesley.
3. Bedau, Mark A., Emile Snyder, C. Titus Brown, Norman Packard. 1997. A comparison of evolutionary activity in artificial evolving systems and in the biosphere. In P. Husbands and I. Harvey, (Eds.), *Fourth European conference on artificial life* (pp. 125–134). Cambridge: MIT Press.
4. Bedau, Mark A., Emile Snyder, Norman 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). Cambridge: MIT Press.
5. Dawkins, R. 1976. *The selfish gene*. New York: Oxford University Press.
6. Gould, S. J., Lewontin, R. C. 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.
7. Holland, John H. 1992. *Adaptation in natural and artificial systems*, 2nd edition. Cambridge: MIT Press.
8. Kimura, Motoo. 1983. *Neutral theory of molecular evolution*. New York: Cambridge University Press.
9. Maynard Smith, J. 1975. *The theory of evolution*, 3rd edition. New York: Penguin.
10. Mayr, E. 1988. *Towards a new philosophy of biology*. Cambridge: Harvard University Press.
11. Packard, Norman. 1989. Intrinsic adaptation in a simple model for evolution. In C. Langton, (Ed.), *Artificial life* (pp. 141–155). Redwood City CA: Addison-Wesley.
12. Ray, T. S. 1992. An approach to the synthesis of life. In C. Langton, C. Taylor, D. Farmer, S. Rasmussen, (Eds.), *Artificial life II* (pp. 371–408). Redwood City CA: Addison-Wesley.
13. Rechtsteiner, Andreas and Mark A. Bedau. 1999. A generic neutral model for measuring excess evolutionary activity of genotypes. In W. Banzhaf, J. Daida, A. E. Eiben, M. H. Garzon, V. Honavar, M. Jakiela, and R. E. Smith, eds., *GECCO-99: Proceedings of the genetic and evolutionary computation conference, July 13-17, 1999, Orlando, Florida USA*. San Francisco, CA: Morgan Kaufmann.