
Four Steps Toward Open-Ended Evolution

C. C. Maley

Computer Science Dept.
University of New Mexico
Albuquerque, NM 87131
cmaley@cs.unm.edu
505-277-8912

Abstract

In 1998, Bedau et al. defined a set of metrics for characterizing the long-term evolutionary dynamics of a system. They argued that no known artificial system demonstrates the unbounded evolutionary activity observed in the fossil record. In response we have developed a series of toy models that approach and eventually succeed in demonstrating unbounded evolutionary activity. The underwhelming success of the models suggests that there must be more to open-ended evolution than just unbounded evolutionary activity as the term is currently defined. We derive some potential extensions to the metrics and requirements for developing open-ended evolution in an artificial system.

1 INTRODUCTION

The richness of biological life has never been replicated in our artificial models of evolution. Yet, the challenge to understand the basis of open-ended evolution and to construct an artificial system that demonstrates such a capacity stands as one of the fundamental unsolved problems of Artificial Life. Thanks to Bedau et al. (1998) we now have the tools to begin this project. We have used those tools to construct a model that arguably shares the same evolutionary dynamics as biological life.

The attempt to understand the processes biological evolution encompasses both a scientific and an engineering problem. As biologists we would like to know why life has diversified over the last 3.8 billion years (Maley, 1998; Benton, 1990). Why isn't life just one big homogeneous mass of photosynthetic or thermophilic sludge? As engineers, we would like to know

how to tap into the astonishing capacity for innovation demonstrated by biological evolution (e.g., Koza, 1992). Why can't we replicate the richness of biological life? Why can't we build an open-ended evolutionary system?

1.1 OPEN-ENDED EVOLUTION

One obstacle to the development of open-ended evolutionary systems has been the lack of a clear definition of the problem. What is the goal? What do we mean by "open-ended evolution?" Our intuition is that we would like an evolutionary system to produce "individuals" (organisms?) of a greater complexity and diversity than the initial individuals of the system. What is more, we would like the evolutionary system, like life, to continue to produce individuals of increasing complexity and diversity. In part this last requirement is based on a misperception of the history of life (McShea, 1991). Much of modern life is single celled and its complexity probably hasn't changed much over the last few billion years.

In order to get a handle on the problem of open-ended evolution, let us separate out the issues of diversity and complexity, setting aside the evolution of complexity for the moment. The diversity of life seems to have increased (non-monotonically) over time (Benton, 1995). How might we construct an evolutionary system that also demonstrates an ever-increasing production of diversity?

1.2 METRICS AND CLASSES OF EVOLUTIONARY DYNAMICS

Bedau et al. (1998) have provided a set of metrics for measuring the evolutionary dynamics of a system, be it artificial or natural.¹ The basis of Bedau et al.'s met-

¹See Adami (1994) for an entirely different approach to measuring evolutionary dynamics.

rics is the “component.” One must decide at what level to analyze the evolutionary dynamics of the system. In the following models, we will identify the genotypes of the individuals as the components of the systems. They then define $a_i(t)$ as the activity of component i at time t . While different measurements of activity may be chosen, a convenient one, used by Bedau et al. is to let $a_i(t)$ be the number of time steps that component i has existed in the system since its origin. In other words, $a_i(t)$ is the age of component i . However, if component i has gone extinct, then $a_i(t) = 0$.

We can then use $a_i(t)$ to define three metrics of the evolutionary dynamics of a system. Let $D(t)$ be the diversity of the system at time t .

$$D(t) = \#\{i : a_i(t) > 0\} \quad (1)$$

where $\#\{\cdot\}$ means the number of members of the set (set cardinality).²

The mean cumulative activity, $\bar{A}_{\text{cum}}(t)$ is the sum of all the activities of the components that exist in the model at time t divided by the diversity. That is,

$$\bar{A}_{\text{cum}}(t) = \frac{1}{D(t)} \sum_i a_i(t) \quad (2)$$

The final metric is slightly more subtle. We would like to measure the amount of novel *adaptive* activity in the model. The problem comes in the distinction between new activity that is essentially just random as opposed to new activity that is the result of natural selection preserving and augmenting that component. Bedau et al. choose to identify components that have survived past some threshold a' as new adaptive activity³. We can then define the *new adaptive activity* per component, $A_{\text{new}}(t)$, as the sum of the activities of the components that reach this threshold in time step t divided by the diversity.

$$A_{\text{new}}(t) = \frac{1}{D(t)} \sum_{a_i(t)=a'} a_i(t) \quad (3)$$

In some cases, when you don't have fine grained information about the evolution of a system, it may be more convenient to count the number of components whose activity falls into a range $a_0 \leq a' \leq a_1$. We haven't yet described how to determine the adaptive

²See Solow et al. (1993) for a more sophisticated treatment of the issues surrounding the measurement of biological diversity for the purposes of conservation.

³They use the term *new evolutionary activity*. However, it is really the issue of adaptive activity that they are interested in so I have chosen to use the slightly more descriptive term *new adaptive activity*.

activity threshold a' . This requires the implementation of a neutral shadow of the model. The guiding idea is that components with activities $a_i(t) > a'$ are more likely to exist in the system due to their adaptive value than merely by chance. To determine the value of a' we record the sequence of births and deaths for every time step of the model. Then we run a “neutral shadow” of the model in which the exact same sequence of births and deaths occurs but the individuals that reproduce and die are chosen randomly from the current population. Thus there is no natural selection in the neutral model, because fertility and mortality are independent of genotype. However, all other aspects of the models are equivalent. Thus the neutral shadow is an experimental control for studying adaptive activity in the model. We then compare the distributions of the activity counters a_i of the model and its shadow for all the components that went extinct as well as the components that remain at the end of the run of the model. If we can identify an activity level beyond which a component is more likely to have survived due to its adaptive value than by chance. This point is a' .

Finally, Bedau et al. describe the long term behavior of the metrics. A metric is considered to be unbounded if it generally increases over time. More precisely, a function $f(t)$ is unbounded if and only if

$$\lim_{t \rightarrow \infty} \left(\frac{\sup(f(t))}{t} \right) > 0 \quad (4)$$

where $\sup(\cdot)$ is the supremum function (the maximum so far). So the maximum of the metric must grow at least linearly with time.

Similarly, a metric is called positive if it does not descend to 0 over time. That is, $f(t)$ is positive if and only if

$$\lim_{t \rightarrow \infty} \left(\frac{\int_0^t f(t) dt}{t} \right) > 0 \quad (5)$$

Bedau et al. (1998) use these metrics to classify the evolutionary dynamics of both artificial and natural systems. Biological life seems to demonstrate unbounded diversity, positive new adaptive activity and bounded mean cumulative activity (unbounded cumulative activity). Intriguingly, no artificial evolutionary system has been shown to share this class. Thus the gauntlet has been thrown. Can we design an artificial evolutionary system that shares these characteristics with life? If not, why not? And if so, what does it tell us about evolution and our metrics for measuring it?

2 PICKING UP THE GAUNTLET

I love fools' experiments. I am always making them. Darwin (1896)

All evolutionary systems have a diversity ceiling or carrying capacity defined by the physics of the system. In our world, we don't know exactly what this ceiling is. However, we can derive some upper bounds. The diversity of life on Earth must be bounded by the number of molecules that compose the planet and its atmosphere. We can guess at some other bounds on the Earth's diversity as it probably depends on quantities of limiting resources like carbon atoms and the input of free energy, as well as the minimum sustainable population sizes for each species. Life is probably nowhere near these upper limits. Most of the radiant energy that strikes the Earth goes unutilized (Begon et al., 1990). This suggests that there must be many open niches in adaptive space. In other words, I do not believe that we have yet observed the asymptotic behavior of life on Earth. An important difference between life on earth and life in the simulated worlds of artificial life is that typically the artificial life worlds fill up and utilize all the available energy.

Despite the fact that most of the energy that strikes the Earth goes unutilized, there is evidence that some clades reach a diversity plateau (Jablonski and Bottjer, 1990). Indeed, the phenomenon of punctuated equilibria (Gould and Eldridge, 1977) implies that diversification is constrained for long periods of time, with occasional bursts of innovation when a species colonizes some new adaptive space. Thus, a model that mimicked the evolutionary dynamics of life should have at least two time constants. A fast constant that describes the expansion of "life" given the capacity to make use of currently unutilized resources. Artificial life models often lack a slower time constant under which innovations arise that open up new adaptive space and so increase the current carrying capacity of the model. In reality occasional innovations arise that do not destroy the previously filled niches, but rather add to the complexity of the ecosystem. The fossil record suggests that diversity

Tierra (Ray, 1992) comes close to this. It has one time constant of expansion and another, slower one, of ecological innovation. However, with the exception of the first parasite, the innovations tend to drive the previous species extinct. Most of the innovations do not open up previously unutilized resources. Thus, with the exception of the first parasites, you get turnover but no real expansion of diversity.

In deference to the fundamental limitations of energy

and space, we will only require that an evolutionary system demonstrate unbounded diversity up until those limits are approached. This is a fundamental shift in the approach laid down by Bedau et al. (1998). We expect that in most systems there will be short term transient dynamics while the initial conditions adjust to the constraints of the model. This is followed by a growth period. And finally, that growth will approach an asymptote determined by the resource limitations. Since we believe that life on earth is still in the growth phase, we will be interested in models that mimic the dynamics of biology during the growth phase. Note that from this perspective, the problem with Tierra is that the model hits its asymptote during its transient phase and all we are left with are the asymptotic behaviors of the model. And so we have our first requirement:

Requirement 1 *An open-ended evolutionary system must demonstrate unbounded diversity during its growth phase.*⁴

2.1 URMODEL 1

Let us start with an extremely simple toy model to help us illuminate the issues. Urmodel 1 is an attempt to create a model that demonstrates unbounded diversity. This can be realized by a simple diffusive process. We might specify such a model as:

- A completely neutral (flat) adaptive landscape.
- Mutation (1.0) which should diversify the organisms.
- Stop before adaptive landscape fills.

Consider each position in a two-dimensional array as a viable niche. Then, as long as there is still empty niche space, and some probability of an occupied niche expanding into an empty neighbor, diversification will happen. Let each species have a 32 bit genotype (0 indicates an empty position in the environment). These are the components of the model. At reproduction (speciation) the new component may flip one bit with probability *mutation rate*. In this case the mutation rate is 1. We start the model with a single species in the central niche of a 256 by 256 grid. The niches

⁴One might also be interested in finding models with bounded diversity but unbounded cumulative activity in their asymptotic phase. Perhaps the Earth will demonstrate those dynamics once it reaches its asymptote. However, I will only focus on the growth phase of evolutionary systems.

are updated sequentially. A non-empty niche reproduces into the empty niches amongst its four adjacent neighbors. However, a component that has originated in time step t may not reproduce until time step $t + 1$. We stop the model when half of the niches (32K) have been filled. This prevents edge effects.

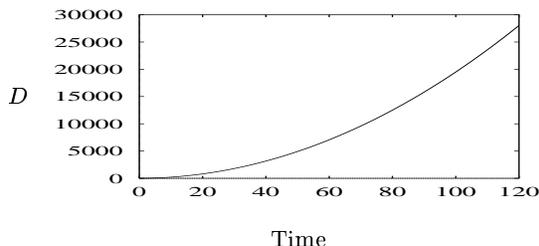


Figure 1: The diversity D of components in Urmodel 1 over time. The data is the average of 50 runs of the model.

Diversity is indeed unbounded, as shown in Figure 1. Yet, there is an important flaw in this model. Urmodel 1 demonstrates evolution in the strict population geneticist definition of the term. The gene frequencies change. Yet, it fails to meet one of the basic criteria of natural selection: the heritable variation in the population has no effect on fertility. The distribution of activities of Urmodel 1 and its neutral shadow are identical. We may now introduce a requirement on our models for open-ended evolution.

Requirement 2 *An open-ended evolutionary system must embody selection.*

This effectively rules out purely neutral evolution. It is only through the termination of the maladaptive and adaptively neutral components that A_{new} can be used to distinguish new adaptive activity by its longevity. This is subsumed by Bedau et al.’s requirement:

Requirement 3 *An open-ended evolutionary system must exhibit continuing (“positive”) new adaptive activity.*

All the following models have the same structure as Urmodel 1 except where noted.

2.2 URMODEL 2

A simple unbounded diversity model with natural selection:

- Mutation (0.1) diversifies species on a 2D grid of niches
- Neighboring species compete for both empty and previously (but not newly) filled niches.

- “Dissimilarity” provides a competitive advantage.
- Stop the model before the adaptive landscape fills.

This time a species’ genotype determines its spread, in the context of its neighbors. The only question is how we measure the most “dissimilar” from a group of two or more. We can tally the hamming distance between each species and its competitors. Then break ties randomly. The winner, the species most different from its competitors, gets to reproduce into the central position of the neighborhood. As in Urmodel 1, the model is updated sequentially and a newborn organism cannot be displaced by competition in the same time step in which it is born, nor can it reproduce in that time step. This time the *mutation rate* is 0.1. That is, there is a 1 in 10 chance that a single bit is flipped in a new component at the time of reproduction.

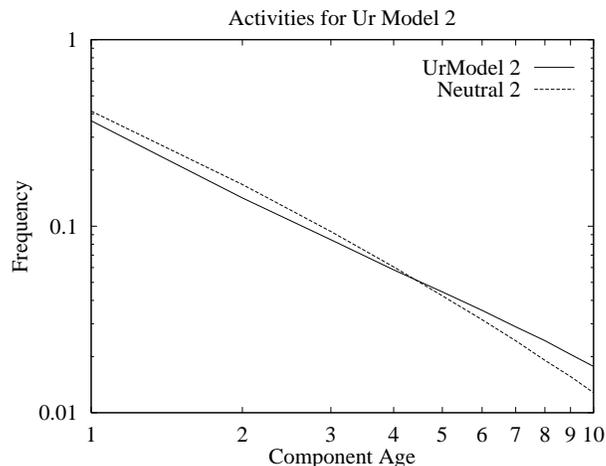


Figure 2: A log-log plot of the age of components versus their frequency in the population over the entire run of Urmodel 2 and its neutral shadow. The point at which they cross, a' is the age threshold after which the survival of a component is more likely due to its adaptive value than to pure chance. We set $a' = 5$. This data is the average of 50 runs of both Urmodel 2 and the neutral shadows for each of those runs.

Selection for differences can be justified by niche overlap theory (Levins, 1968). The more that two species’ niches overlap (share common resources) the more likely that one will drive the other to extinction through competitive exclusion. Thus, being different from your competitors makes it easier to survive.

The distributions of activities between Urmodel 2 and its neutral shadow are similar, though the fact that the curves in Figure 2 cross just before $a' = 5$ suggests that some components in Urmodel 2 have adaptive value. If we examine the evolutionary dynamics of Urmodel

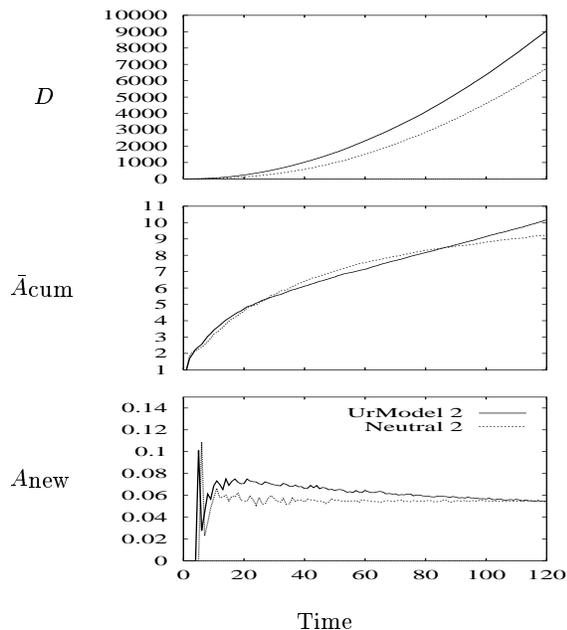


Figure 3: A plot of the diversity D , the mean cumulative activity \bar{A}_{cum} , and the new adaptive activity A_{new} over time in Urmodel 2. This data is the average of 50 runs. A_{new} , defined with $a' = 5$, seems to be slowly descending over time, and so its long-term dynamics are not positive.

2, shown in Figure 3, we discover that D and \bar{A}_{cum} are unbounded, but A_{new} is not positive. In the limit, it appears that A_{new} will tend toward 0.

Perhaps selection for local dissimilarity is too weak to produce a clear effect in the activity levels. Or perhaps the metric A_{new} is not well suited to the detection of this sort of adaptation. In either case, we still have not satisfied Requirement 3.

2.3 URMODEL 3

Let us consider a more traditional description of evolution. Organisms evolve in heterogeneous environments. When a mutation turns out to produce an adaptive change, that mutation tends to spread rapidly in the population in what is called a “selective sweep.” Would such a model produce open-ended evolution?

- We will model the evolution of parasite species by matching the bit pattern of their genotypes against the heterogeneous but static environment of host genotypes. The parasite’s fitness is determined by the degree of match (number of bits with the same value in the same positions) be-

tween the bit patterns. We may think of each bit in the pattern representing the locus of some character of the parasite relevant to the exploitation of its host.

- Mutation (0.1) occurs in the parasite bit patterns at speciation (reproduction).
- A new parasite species may displace a neighboring species if it would have an equal or higher fitness than the neighbor.
- The probability of invading a new host species that has never been parasitized (0.1) is lower than the probability of invading a host that has been infected before (1 if it is a superior competitor relative to the resident parasite).
- Stop the model when half the adaptive landscape fills.

We have chosen to call the evolving species “parasites” and their environments “hosts” in anticipation of Urmodel 4. For now, note that Urmodel 3 has abandoned frequency based selection and dynamic fitness functions. The host genotypes were set using Urmodel 1 with a mutation rate of 0.1. However the host genotypes and thus the selective pressures on the parasites remain constant. The model is seeded with a single parasite species that fails to match its host in any of its 32 bits.

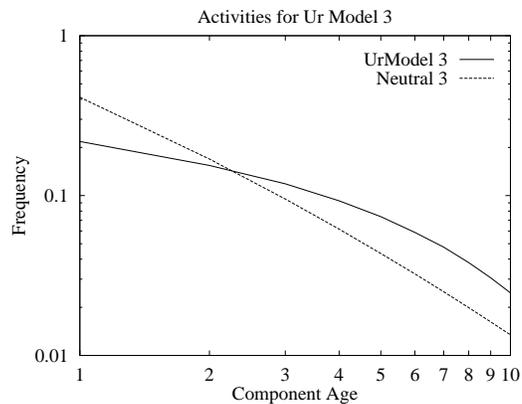


Figure 4: A log-log plot of the age of components versus their frequency in the population over the entire run of Urmodel 3 and its neutral shadow. We conservatively set $a' = 3$. This data is the average of 50 runs of both Urmodel 3 and the neutral shadows for each of those runs.

Figure 5 shows that Urmodel 3 is our first example of unbounded evolutionary activity. It shares the class of evolutionary dynamics with biological life, except

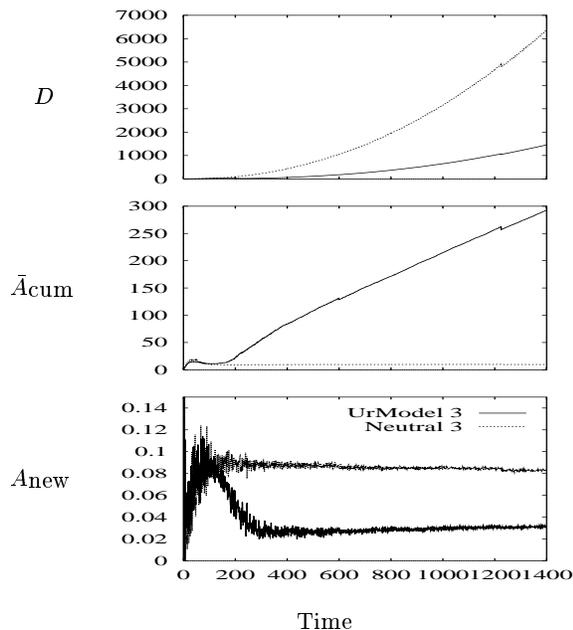


Figure 5: A plot of the diversity D , the mean cumulative activity \bar{A}_{cum} , and the new adaptive activity A_{new} over time in Urmodel 3. This data is the average of 50 runs. A_{new} , defined with $a' = 3$, is positive and might even be unbounded. Urmodel 3 shows less new activity than its neutral shadow (the upper curve). However the new activity of Urmodel 3 rises over time while the neutral shadow's new activity decreases. There is no reason to think they would not cross as the system continued to grow. Urmodel 3 is the first known artificial evolutionary system demonstrating unbounded evolutionary activity.

that it does life one better by exhibiting unbounded, rather than bounded, mean cumulative activity. However, if we were to let death intrude on the model in forms other than competition, we might well see $\bar{A}_{cum}(t)$ reach a natural bound. The previous statements should only be accepted with the caveat that we must extrapolate the dynamics of the model past the point where A_{new} surpasses the new activity in its neutral shadow.

2.4 URMODEL 4

The most important aspect of an organism's environment are the other organisms with which it interacts. Urmodel 3 lacks an ecology, however primitive. It is informative to examine how the dynamics change when we add coevolution to the model.

Call the organisms of Urmodel 3 "parasites" and the environment they are exploiting "hosts." Then, if we

let the hosts evolve defenses against the parasites and the parasites evolve responses to those defense, we have a new form of non-static fitness function for both the parasites and the hosts. Urmodel 4 is thus similar to Urmodel 2 with their dynamic fitness functions. However, we have abandoned an abstract frequency based selection in favor of a more traditional dynamic fitness function: coevolution.

- We will model this coevolution by the matching of two bit patterns. The host's fitness is determined by the degree of mismatch and the parasite's fitness is determined by the degree of match between the bit patterns. We may think of each bit in the pattern representing the locus of some character of the species relevant to the coevolutionary interaction.
- Mutation (0.1) occurs in both the host and parasite bit patterns at speciation.
- A new species may displace a neighboring species if it would have a higher fitness than the neighbor.
- Only hosts may expand into novel adaptive space and their probability of doing so (0.1) is lower than the probability of invading a niche that has been filled before (which depends only on competition).
- Stop the model before the adaptive landscape fills.

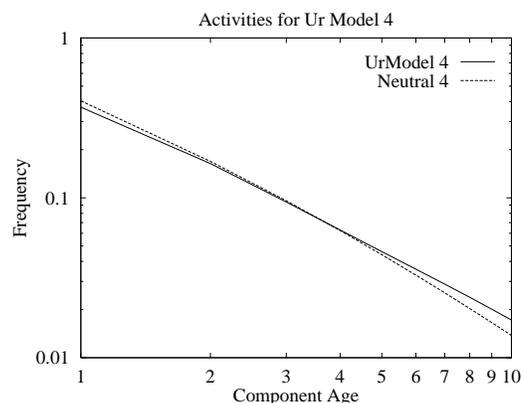


Figure 6: A log-log plot of the age of components versus their frequency in the population over the entire run of Urmodel 4. We set $a' = 4$ based on this data. The data is the average of 50 runs of both Urmodel 4 and the neutral shadows for each of those runs.

The activity distributions of Urmodel 4 and its neutral shadow do not show as strong a difference as in Urmodel 3. Natural selection is generating upheaval

in the components, not stability. A mutation that is initially adaptive will spread through the population of hosts or parasites. But the very success of this mutation spells its doom. Consequently, there is strong pressure on the coevolving species to respond to that mutation, and when it does, the initial mutation becomes maladaptive. We can now see that the action of natural selection on the activity distribution of Urmodel 2 (Figure 2) prevented a clear distinction between the model and its neutral shadow because selection was biased toward novel genotypes.

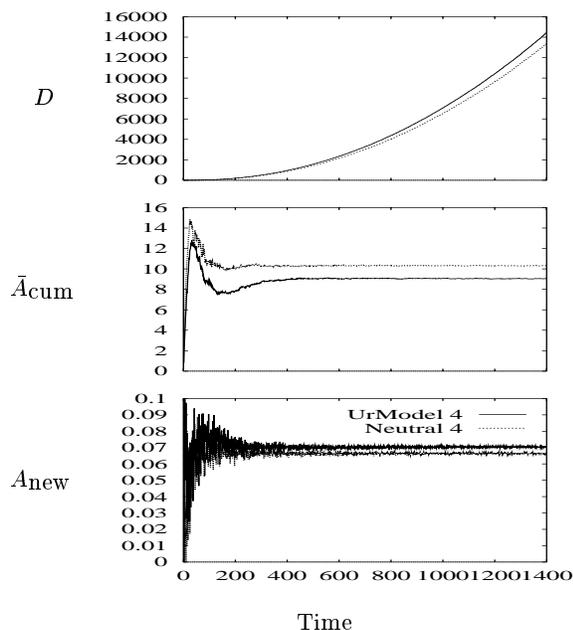


Figure 7: A plot of the diversity D , the mean cumulative activity \bar{A}_{cum} , and new adaptive activity A_{new} over time in Urmodel 4. This data is the average of 50 runs. The lower curve in the plot of A_{new} shows the neutral shadow indicating that selection is maintaining new adaptive components more than the neutral shadow. The fact that \bar{A}_{cum} is bounded shows that there is a lot of turnover amongst the components in this coevolutionary model.

Figure 7 shows the evolutionary dynamics of Urmodel 4. Now \bar{A}_{cum} is bounded but less than its neutral shadow due to the turnover of coevolution. A_{new} is positive and greater than its neutral shadow, but it is nonetheless bounded. Urmodel 4 is another example of unbounded evolutionary activity.

3 DISCUSSION

The Urmodels illustrate two important lessons. First of all, the activity metric chosen by Bedau et al. (1998)

will sometimes fail to capture adaptive activity if natural selection does not tend to preserve components for long periods of time. Bedau et al. point out that there are many reasonable definitions that might be chosen for D , \bar{A}_{cum} and A_{new} . If the biotic or abiotic environment of a species changes over time, then some components that are adaptive may well prove maladaptive in the future. Furthermore, if the selective pressure is particularly strong, then an adaptation may sweep through a population like wild-fire, only to be replaced shortly thereafter by the next improvement. In this case, even a static fitness function might result in a shorter than random life span for the components of the model. We need a different measure of adaptive activity for dynamic fitness functions. For example, we would predict that the rate of replication of an adaptive component should be greater than rates of growth of components in a neutral shadow. If we let $p_i(t)$ be the number of instances of component i at time t in a model, and p' and k' be thresholds set by comparison to a neutral shadow, then we might redefine

$$A_{new}(t) = \#\{i : p_i(t) > p' \text{ and } a_i(t) < k'\} \quad (6)$$

This metric should identify selective sweeps by picking out the components that are expanding at a faster than random rate. Components might be replaced quickly through selective pressures in a fluctuating environment but adaptive activity could be detected in the rapid growth of each newly selected component.

To be fair, Urmodels 3 and 4 are probably not the first artificial systems created that demonstrate unbounded evolutionary activity. If we restrict our analysis to the dynamics of the systems during their growth phase, many artificial life models probably would share this class with Urmodel 3. The only trick is defer the point when the model hits its true asymptotic behavior for long enough that the growth dynamics of the model are themselves asymptotic in some sense. That is, the early transient effects should die away and the growth dynamics of the model should become stable over time. These are the conditions for artificial dynamics that may be best compared to the dynamics of real life.

The second lesson is that the metrics only get us so far. While we have technically succeeded in creating artificially evolving systems that share the class of unbounded evolutionary systems with biological life, no one would be satisfied with Urmodel 3 as an “open-ended” evolutionary system. Thus, something must be missing from our specification of the problem. There are at least two distasteful aspects of Urmodel 3. First, the niches are imposed from the outside, they are not endogenous to the diversity of the biota.

Requirement 4 *An open-ended evolutionary system must have an endogenous implementation of niches.*

The capacity to survive and flourish should only be a (potentially stochastic) function of the individual's phenotype and its match to its environment. Urmodels 1, 2 3, and 4 (in the case of the hosts) do not fulfill this requirement because they allow any phenotype to survive as long as they land in an open grid location. We should note that requiring endogenous niches leads us further astray from the field of evolutionary algorithms and its application of evolution to engineering.

The second flaw is the fact that the products of Urmodel 3 would never surprise us. In section 1.1 we said that we would set aside the issue of complexity. Clearly, our wonder in the natural world has much to do with its complexity. A puddle of inert, multi-colored and diverse algae would not be nearly so inspirational as the rain forest. Some work has begun in the measurement of complexity in biology (McShea, 1991; Yagil, 1995; Arthur, 1994). A significant contribution to the field would be an adaptation of such a measure to an interesting evolutionary model. If we did have an easily computable metric for complexity, we would like to see unbounded growth in the maximum complexity of the organisms over time, as well as unbounded diversity.

We have intentionally kept a distinction between the still undefined term "open-ended evolution" and Bedau et al.'s "unbounded evolutionary activity." The hope is that through refinements of these metrics we may eventually capture what we mean by "open-ended evolution" and then use those metrics to develop models of open-ended evolution. The gauntlet has been thrown.

Acknowledgments

I would like to thank M. Bedau, S. Forrest and D. Smith for their helpful discussions. In particular, I would like to thank the UNM department of computer science and S. Forrest for her generous support under her Presidential Young Investigators award.

References

Adami, C. (1994). On modelling life. In Brooks, R. A. and Maes, P., editors, *Artificial Life IV*, pages 269–274, Cambridge, MA. MIT Press.

Arthur, W. B. (1994). On the evolution of complexity. In Cowan, G. A., Pines, D., and D.Meltzer, editors, *Complexity: Metaphors, Models, and Reality*, pages 65–81. Addison Wesley Longman, Reading, MA.

Bedau, M. A., Snyder, E., and Packard, N. H. (1998). A classification of long-term evolutionary dynamics. In Adami, C., Belew, R. K., Kitano, H., and Taylor, C. E., editors, *Artificial Life VI*, pages 228–237. MIT Press, Cambridge, MA.

Begon, M., Harper, J. L., and Townsend, C. R. (1990). *Ecology*. Blackwell Scientific Publications, London.

Benton, M. J. (1990). The causes of the diversification of life. In Taylor, P. D. and Larwood, G. P., editors, *Major Evolutionary Radiations*, pages 409–430. Clarendon Press, Oxford, UK.

Benton, M. J. (1995). Diversification and extinction in the history of life. *Science*, 268:52–58.

Darwin, C. R. (1896). *The Life and Letters of Charles Darwin, including an autobiographical chapter; edited by his son Francis Darwin*. Appleton and Co., New York, NY.

Gould, S. J. and Eldridge, N. (1977). Punctuated equilibria: The tempo and mode of evolution reconsidered. *Paleobiology*, 3:115–151.

Jablonski, D. and Bottjer, D. J. (1990). The ecology of evolutionary innovation: The fossil record. In Nitecki, M. H., editor, *Evolutionary Innovations*, pages 253–288. The University of Chicago Press, Chicago, IL.

Koza, J. R. (1992). *Genetic Programming*. MIT Press, Cambridge, MA.

Levins, R. (1968). *Evolution in Changing Environments*. Princeton University Press, Princeton, NJ.

Maley, C. C. (1998). *The Evolution of Biodiversity: A Simulation Approach*. PhD thesis, Massachusetts Institute of Technology, Cambridge, MA.

McShea, D. W. (1991). Complexity and evolution: What everybody knows. *Biology and Philosophy*, 6:303–324.

Ray, T. S. (1992). An approach to the synthesis of life. In Langton, C. G., Taylor, C., Farmer, J. D., and Rasmussen, S., editors, *Artificial Life II*, pages 371–408, Reading, MA. Addison-Wesley.

Solow, A., Polasky, S., and Broadus, J. (1993). On the measurement of biological diversity. *Journal of Environmental Economics and Management*, 24:60–68.

Yagil, G. (1995). Complexity analysis of a self-organizing vs. a template-directed system. In Morán, F., Moreno, A., Merelo, J. J., and Chacón, P., editors, *Advances in Artificial Life*, pages 179–187, Berlin. Springer-Verlag.