

Exploring the Dynamics of Adaptation with Evolutionary Activity Plots

Seth Bullock
School of Computing
University of Leeds
Leeds, LS2 9JT, UK
seth@comp.leeds.ac.uk

Mark A. Bedau*
Reed College
3203 SE Woodstock Blvd.
Portland OR 97202
mab@reed.edu

Abstract Evolutionary activity statistics and their visualization are introduced, and their motivation is explained. Examples of their use are described, and their strengths and limitations are discussed. References to more extensive or general accounts of these techniques are provided.

Keywords
Evolutionary activity, evolutionary adaptation, visualization

1 Introduction

The evolutionary activity statistics and visualization techniques introduced by Bedau and Packard [8] are perhaps unique within artificial life in that they are indigenous to the field, are long-lived, and are still being developed in interesting new directions. The evolutionary activity approach was developed explicitly to deal with issues central to artificial and natural evolutionary adaptation, and has persisted for over a decade, being extended and generalized along the way in a series of articles [10, 11, 20, 19, 4, 6, 5, 9, 22, 17]. Moreover, they have been taken up by other researchers, who have used them to help explore a wide range of questions on evolvability and adaptation (e.g., [21, 26, 23, 12, 13, 16, 24]).

Here, we provide a simple description of the basic technique and some example applications, before considering the strengths and limitations of the approach. Interested readers are directed to a more extensive account [5] and more general formulations of these ideas [9, 17].

2 Determining and Visualizing Evolutionary Activity

A paradigm of the idea of evolutionary activity is the notion of measuring the continued *persistence* of elements in an adaptive system as a proxy for their adaptive *significance*—if a particular element persists in the system for a long time, this is likely to be because it is being maintained by selection. In order to do this one must identify (i) the class of elements to be tracked (genotypes, alleles, equivalence classes of alleles, etc.) and (ii) a measure of identity—what makes an element existing at time t the same element as one logged at time $t-1$. For example, if we are tracking the evolutionary activity of alleles, we might specify that an allele has persisted over time if two genotypes, one drawn from the population at time t and another drawn from the same population at time $t+1$, share the same *value* at the same *location* on their respective genomes.

* Corresponding author.

Once these decisions have been made, a measure of activity must be defined. Different kinds of elements call for different definitions, but in each case the aim is to capture evidence that an element is an adaptation, that is, that it is participating in the system because of the benefit it provides. The definition operationalizes the notion of being an adaptation, and continually reflects the element's adaptive history as so operationalized. The most appropriate definition in a particular context will depend on the type of element in question, as well as the availability of the relevant data.

Assume, for example, that we are interested in genotypes all of whose members are related by historical descent (so that independent originations of exactly the same set of genes would be considered different genotypes). In many contexts, such a genotype's adaptive success is reflected by its frequency in the population, so an appropriate measure of its activity would be its cumulative frequency over a period of evolutionary time. In this case, a genotype's activity would become defined once its first instance appears in the population. From then on, the genotype's activity is incremented at each moment by its frequency in the population. As long as the genotype remains in the population, its activity increases monotonically, with the rate of increase being proportional to its frequency in the population. Should the genotype go extinct, its activity is no longer defined. To track evolutionary activity in the entire population, one would simply make these measurements for each unique genotype generated during the course of evolution.

If an element's activity is plotted over time, it will appear as a monotonically increasing wave (by convention waves are terminated upon extinction of the corresponding element). If we plot this measure for every unique element in an adaptive system, we will see a forest of such *activity waves*, some large, some small, some rising early, some late, some rising slowly, some rapidly, some terminating immediately, some persisting for long periods of time. It is these differences between waves that can tell us useful things about the character of the adaptation that the population is undergoing. But before we can really make use of this type of visualization, we need some way of determining the significance of differences between activity waves—for example, how long must an activity wave persist before we can be confident that it represents a response to selective pressure, rather than merely the transient effects of evolutionary drift?

In order to deal with this issue, we need some way of *normalising* our measures of activity. This is achieved through constructing a *neutral*, or nonadaptive, analogue of the adaptive system that we are exploring—often, an exact replica save that all selective pressures have been removed—and using the level of evolutionary activity in this neutral model as a baseline against which to compare the statistics generated by the actual system that we are studying. One method of constructing a neutral model N of an adaptive system S is to record the time at which every birth and death event in S occurred during a period of evolution, and then to impose the same events on N at the equivalent points in time, but to choose who reproduces or dies at random rather than according to any selection pressure. The evolutionary activity in N is thus a neutral analogue of that in S , and subtracting the activity in N from that in S normalizes our evolutionary activity statistics.

An example plot of evolutionary activity in an adaptive system and some of the phenomena it reveals is given in Figure 1. However, it should be noted that there are many variants of the approach described above that have been tailored to tracking and visualizing different adaptive systems.

3 Applications

Evolutionary activity plots of the kind described above have been used to track elements at a variety of different levels: alleles [8], allelic equivalence classes [9], genotypes [5], and even taxonomic families in the fossil record [10, 11] and the patent record [22].

Across these different scales, the evolutionary activity approach is able to shed light on the dynamics of adaptation, extinction, succession, competitive exclusion, speciation, neutral drift, and so on. Key questions addressed include classifying long-term evolutionary trends [10, 15, 12, 13], studying punctuated equilibrium dynamics [21], measuring evolvability [4], and exploring the




Figure 1. An example evolutionary activity plot (adapted from [5, Figure 5]). The upper panel displays evolutionary activity waves for a Tierra-like [18] and an Avida-like [2] system, comprising a population of simple self-reproducing programs. The evolutionary activity in the corresponding neutral analogue (not shown here; see [5]) is orders of magnitude smaller than even the smallest labeled wave. The lower panel depicts the average population fitness. The termination of a significant wave and arrival of subsequent waves marks a succession event. For instance, at or around iteration 600 the ancestral genotype is driven extinct by a small group of more efficient replicators, one of which eventually dominates, until it is succeeded by a group of neutral variants of an even more efficient phenotype. Around iteration 2500 a mutant genotype invades the population but cannot outcompete the incumbent genotype, due to its tiny fitness advantage.

evolution of evolvability [17]. These concerns cluster around a central and fundamental topic of enquiry: how can an adaptive system support continued adaptation indefinitely, rather than eventually equilibrating at some stable configuration that prohibits further adaptive change? This question has been recognized as one of the key open problems faced by artificial life—“determine what is inevitable in the open-ended evolution of life” [7].

4 Strengths and Limitations

The continued and increasing use of evolutionary activity statistics and visualization evidences its strength in efficiently depicting adaptive trends and evolutionary phenomena that are otherwise difficult to track or describe. In particular, the technique benefits from its simplicity: both the relative ease with which the relevant values may be calculated and the relative ease with which they can be depicted and interpreted, piggybacking on our familiarity with the 2D line graph of variation over time. A further key strength lies in the generality of the evolutionary activity approach, spanning many levels of description and many different kinds of ALife system.

However, there are issues that currently limit its application. For instance, where it is hard to obtain complete data describing the population makeup as it changes over time (as in most natural adaptive systems), the use of evolutionary activity statistics can be restricted as a consequence.

Moreover, even in cases where we have complete access to a system's state as it evolves, difficulties in effectively implementing evolutionary activity statistics may stem from problems associated with uniquely identifying the elements to be tracked.

Where a population contains simple, fixed-length genotypes comprising genes that map directly, one to one, onto phenotypic traits, there is little difficulty in tracking entire genotypes as they increase or decrease in frequency, or the concentration of individual alleles in a population as they evolve, or even tracking the behavioral episodes that correspond to a particular allele (e.g., when a particular production-system rule is fired). However, where the genotype-phenotype relationship is less straightforward, it can be difficult to attribute a behavioral response or phenotypic trait to a particular genetic element. Epistasis or morphogenesis may ensure that multiple alleles are responsible for a particular phenotypic trait, or that multiple phenotypic traits are influenced by a single allele. Here, tracking single alleles may fail to clearly reflect adaptive change in phenotypic traits [17]. Furthermore, where genes are represented by continuous rather than discrete values, even tracking the evolutionary activity of genotypes can prove problematic, since equivalence between genotypes becomes a matter of degree and hence difficult to determine. More sophisticated approaches to tracking the activity of elements in these type of scenarios are required.

5 Conclusions

Evolutionary activity statistics and visualization have proved to be durable and effective tools in the study of adaptive dynamics. They have been generalized to deal with a very wide range of systems and questions. With further work on epistatic, developmental, and real-valued systems, we hope to see their successful application to adaptive systems in general.

References

1. Adami, C., Belew, R., Kitano, H., & Taylor, C., (Eds.) (1998). *Artificial life VI*. Cambridge, MA: MIT Press.
2. Adami, C., & Brown, C. T. (1994). Evolutionary learning in the 2D artificial life system "Avida." In R. Brooks & P. Maes (Eds.), *Artificial life IV* (pp. 377–381). Cambridge, MA: MIT Press.
3. Banzhaf, W., Daida, J., Eiben, A. E., Garzon, M. H., Honavar, V., Jakiela, M., & Smith, R. E. (Eds.) (1999). *Genetic and Evolutionary Computation Conference*. San Francisco: Morgan Kaufmann.
4. Bedau, M. A. (1999). Quantifying the extent and intensity of adaptive evolution. In [27] (pp. 34–27).
5. Bedau M. A., & Brown C. T. (1999). Visualizing evolutionary activity of genotypes. *Artificial Life*, 5(1), 17–35.
6. Bedau, M. A., Joshi, S., & Lillie, B. (1999). Visualizing waves of evolutionary activity of alleles. In [27] (pp. 96–98).
7. Bedau M. A., McCaskill J. S., Packard N. H., Rasmussen S., Adami C., Green D. G., Harvey I., Ikegami T., Kaneko K., & Ray T. S. (2000). Open problems in artificial life. *Artificial Life*, 6, 363–373.
8. Bedau, M. A., & Packard, N. H. (1992). Measurement of evolutionary activity, teleology, and life. In [14] (pp. 431–461).
9. Bedau M. A., & Raven M. J. (2002). Visualizing adaptive evolutionary activity of allele tokens and of phenotypic equivalence classes of alleles. In T. M. C. Smith, S. Bullock, & J. Bird (Eds.), *Beyond fitness: Visualizing evolution—workshop proceedings of Eighth International Conference on Artificial Life* (pp. 119–130). UNSW.
10. Bedau, M. A., Snyder, E., Brown, C. T., & Packard, N. H. (1997). A comparison of evolutionary activity in artificial evolving systems and the biosphere. In P. Husbands & I. Harvey (Eds.), *Fourth European Conference on Artificial Life* (pp. 125–134). Cambridge, MA: MIT Press.
11. Bedau, M. A., Snyder, E., & Packard, N. H. (1998). A classification of long-term evolutionary dynamics. In [1] (pp. 228–237).
12. Channon, A. D. (2001). Passing the ALife test: Activity statistics classify evolution in Geb as unbounded. In J. Keleman & P. Sosik (Eds.), *Sixth European Conference on Artificial Life* (pp. 417–426). Heidelberg: Springer-Verlag.

13. Channon, A. D. (2002). Improving and still passing the ALife test: Component-normalised activity statistics classify evolution in Geb as unbounded. In [25] (pp. 173–181).
14. Langton, C. G., Taylor, C., Farmer, J. D., & Rasmussen, S. (Eds.) (1992). *Artificial life II*. Redwood City, CA: Addison-Wesley.
15. Maley, C. C. (1999). Four steps toward open-ended evolution. In [3] (pp. 1336–1343).
16. Pachepsky, E., Taylor, T., & Jones, S. (2002). Mutualism promotes diversity and stability in a simple artificial ecosystem. *Artificial Life*, 8, 5–24.
17. Raven, M. J., & Bedau, M. A. (2003). A general framework for evolutionary activity statistics. In W. Banzhaf, T. Christaller, P. Dittrich, J. T. Kim, & J. Ziegler (Eds.), *Seventh European Conference on Artificial Life* (pp. 676–685). Heidelberg: Springer-Verlag.
18. Ray, T. S. (1992). An approach to the synthesis of life. In [14] (pp. 371–408).
19. Rechtsteiner, A., & Bedau, M. A. (1999). A generic model for measuring excess evolutionary activity. In [3] (pp. 1366–1373).
20. Rechtsteiner, A., & Bedau, M. A. (1999). A generic model for quantitative comparison of genotypic evolutionary activity. In D. Floreano, J.-D. Nicoud, & F. Mondada (Eds.), *Fifth European Conference on Artificial Life* (pp. 109–118). Heidelberg: Springer-Verlag.
21. Shannon, T. (1998). Generic behavior in the Lindgren non-spatial model of iterated two player games. In [1] (pp. 316–325).
22. Skusa, A., & Bedau, M. A. (2002). Towards a comparison of evolutionary creativity in biological and cultural evolution. In [25] (pp. 233–242).
23. Standish, R. K. (2000). An Ecolab perspective on the Bedau evolutionary statistics. In M. A. Bedau, J. McCaskill, N. Packard, & S. Rasmussen (Eds.), *Artificial life VII* (pp. 238–242). Cambridge, MA: MIT Press.
24. Standish, R. K. (2002). Diversity evolution. In [25] (pp. 131–137).
25. Standish, R. K., Bedau, M. A., & Abbass, H. A. (Eds.) (2002). *Artificial Life VIII*. Cambridge, MA: MIT Press.
26. Taylor, T., & Hallam, J. (1998). Replaying the tape: An investigation into the role of contingency in evolution. In [1] (pp. 256–265).
27. Wu, A. (Ed.) (1999). *Genetic and Evolutionary Computation Conference: Workshop program*. Unpublished.