In Proceedings of the 1999 Genetic and Evolutionary Computation Conference Workshop Program, Orlando, Florida, July 13, 1999, Annie S. Wu, ed., pp. 34-37.

Quantifying the Extent and Intensity of Adaptive Evolution

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

1 Evolvability and Adaptive Evolution

Evolvability is the capacity to create new adaptations, and especially new kinds of adaptations, through the evolutionary process. Evolvability is important both as a theoretical issue in biology and as a practical issue in evolutionary computation. But it is difficult to study evolvability, in part because it is difficult to objectively and feasibly quantify evolvability in a general enough way to compare it across different evolving systems.

This paper is intended as an incremental step toward solving the problem of quantifying evolvability. The progress here is only incremental because I do not address the problem of quantifying evolvability *per se*; rather, I address the related problem of quantifying the degree to which a system exhibits adaptive evolution. This is a step in the right direction, though, for two reasons. First, since evolvability is the capacity to evolve new adaptations, measuring a system's adaptive evolution can tell you something about its evolvability. Second, since the method presented here is objective, feasible, and facilitates the quantitative comparison of adaptive evolution across a wide variety of different evolving systems, it could spread those same virtues to the study of evolvability. This paper explains a method for measuring adaptive evolution and then outlines how the method can be applied in the study of evolvability.

2 The Extent and Intensity of Evolutionary Activity

The method for quantifying adaptive evolution presented here involves using the evolutionary activity statistics originally devised by Bedau and Packard [1]. These statistics have been applied to a variety of evolving systems for a variety of purposes, including visualizing adaptive evolutionary phenomena [1, 2, 3], studying punctuated equilibrium dynamics in evolution [8], identifying long-term evolutionary trends [4, 6], and classifying evolutionary dynamics [5]. Evolutionary activity statistics are computed from data obtained by observing an evolving system, where an evolving system as a population of components participating in a cycle of birth, life and death, with each new 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 fundamental presumption behind evolutionary activity statistics is that adaptive components in an evolving system are innovations that persist and continue to be significant in the system. To identify and quantify such innovations, 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. Components can be identified at any number of levels of analysis; previous work has studied components on the level of individual alleles [1], whole genotypes [4, 5, 8], and taxonomic families [4, 5]. Activity counters are attached 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, $a_i(t) = \sum_{k \le t} \Delta_i(k)$, 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. For example, in contexts in which a component's adaptive value tends to be correlated with its persistence in the system, one could increment a component's activity simply by its age [4, 5]. Or one could increment a component's activity with its concentration in the system [2] or the extent to which it is used or expressed [1, 3] in contexts in which either of those properties tended to be correlated with adaptive value.

The values of the activity counters of each component in the system over all time can be collected in the *component activity distribution*, $C(t, a)$, as follows:

$$C(t,a) = \sum_i \delta(a - a_i(t)) \ , \quad (1)$$

where $\delta(a - a_i(t))$ is the Dirac delta function, equal to one if $a = a_i(t)$ and zero otherwise. Thus, $C(t, a)$ indicates the number of components with activity a at time t. Normalizing the component activity distribution by the diversity, $\frac{C(t,a)}{D(t)}$, gives the *fraction* of components in the population with activity a at time t. Graphing a system's component activity distribution vividly visualizes the various kinds of adaptive evolutionary phenomena occurring in the system [2, 3].

It is possible to define different kinds of evolutionary activity statistics based on component activity distributions. These statistics fall into two broad classes: those reflecting evolutionary activity's *extent* and those reflecting its *intensity*. 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. This corresponds roughly to the mass of activity accumulated in the activity distribution. By contrast, the *intensity* of evolutionary activity reflects the rate at which new adaptive structure are being created. This corresponds to the rate at which nedw activity is flowing into the activity distribution. The extent and intensity of adaptive evolutionary activity are two independently varying aspects of the degree to which a system exhibits adaptive evolution. For example, if a set of adaptive components continue to persist indefinitely without changing and no adaptive innovations invade the system, then the extent of evolutionary activity will be positive and perhaps grow over time, but the intensity of evolutionary activity will fall 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 system's extent of evolutionary activity will be very low if none of those adaptations persist for a significant amount of time or make up a significant amount of the system's evolved structure.

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*, $A_{\mathrm{cum}}(t)$, which simply sums the evolutionary activity of all the components at a given time:

$$A_{\mathrm{cum}}(t) = \sum_i a_i(t) \quad (2)$$

$$\rightarrow \int_0^\infty aC(t,a)\,da \ . \quad (3)$$

(In practice, we compute activity statistics using the sum; the integral indicated is obtained in the limit when activity takes on a continuum of values.) As the integral shows, you can think about $A_{\mathrm{cum}}(t)$ as the mass in the component activity distribution weighted by its level of activity. So, the cumulative activity per component, or *mean cumulative evolutionary activity*, $\bar{A}_{\mathrm{cum}}(t)$, is simply the cumulative evolutionary activity $A_{\mathrm{cum}}(t)$ divided by the diversity, $D(t)$:

$$\bar{A}_{\mathrm{cum}}(t) = \frac{A_{\mathrm{cum}}(t)}{D(t)} \ , \quad (4)$$

where the system's diversity $D(t)$ is the number of components present at time t in the system, $D(t) = \#\{i : a_i(t) > 0\}$. We sometimes refer to mean cumulative evolutionary activity simply as "mean activity." Total and mean cumulative evolutionary activity are measures of the *extent* of a system's adaptive evolution.

Adaptive innovations correspond to new components flowing into the system and proving their adaptive value through their persistent activity. Let a_0 and a_1 define a strip through the component activity distribution function, $C(t, a)$, such that activity values a in the range $a_0 \le a \le a_1$ are among the lowest activity values that can be interpreted as evidence that a component has positive adaptive significance. Then, one reflection of the rate of the evolution of adaptive innovations is the *new evolutionary activity*, $A_{\mathrm{new}}(t)$, which sums the evolutionary activity per component with values between a_0 and a_1:

$$A_{\mathrm{new}}(t) = \frac{1}{D(t)} \sum_{i, a_0 \le a_i(t) \le a_1} a_i(t) \quad (5)$$

$$\rightarrow \frac{1}{D(t)} \int_{a_0}^{a_1} C(t,a)\,da \ . \quad (6)$$

A_{new} reflects the rate at which new activity signaling the positive adaptive value of system components is flowing into the activity distribution. We sometimes refer to new evolutionary activity per component just as "new activity." New activity is a measure of the *intensity* of a system's adaptive evolution.

To ensure that evolutionary activity statistics clearly reflect the degree to which a system's evolutionary activity depends on adaptation rather than other evolutionary forces like chance and necessity, we must screen off the effect of these non-adaptive evolutionary forces. We can accomplish this by comparing the evolutionary dynamics observed in target evolutionary systems with those observed in analogous evolutionary systems in which adaptive evolution cannot happen. I 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 specifically to adaptation. In effect, neutral models are null hypotheses against which the action of adaptive evolution stands out in relief. This neutral-model normalization can be accomplished in various ways, detailed elsewhere [4, 5, 7], but the gist is easy enough to suggest. For example, to measure the amount of the *extent* of evolution that can be attributed to the adaptive success of the components involved, we can measure the difference between the mean cumulative activity observed in the evolving system and the mean cumulative activity in a corresponding neutral model. Likewise, to measure the amount of the *intensity* of evolution that can be attributed to components' adaptive success, we can use a neutral model of the target system to determine that activity level, a', at which we can begin to have confidence that a component's activity reflects its positive adaptive value, and we use a_0 and a_1 (in Eqn. 5 above) to define a small window surrounding a'.

3 Studying Evolvability with Evolutionary Activity Statistics

It is straightforward to study the evolvability of evolutionary algorithms using evolutionary activity statistics. If an evolutionary algorithm supports a high degree of evolvability, then it should create significant levels of adaptive evolution in a wide variety of contexts. But evolutionary activity statistics provide an objective and feasible method for quantitatively comparing the extent and intensity of adaptive evolution in a wide variety of different systems. So, one can use these statistics to measure the extent and intensity of the adaptive evolution generated by different evolutionary algorithms, and thus infer their level of evolvability.

The general methodology for this study of evolvability is to observe evolutionary activity statistics generated by a variety of different evolutionary algorithms applied to a variety of different problems. An algorithm with a high degree of evolvability would be able to generate lots of evolutionary activity in a wide variety of contexts. For example, one could start by applying the algorithm to small-scale problems and then observe how the extent and intensity of evolutionary activity change as larger-scale problems are tackled. Pooling the evolutionary activity data from different scale problems would produce a general activity profile for the algorithm. This activity profile would reveal which sizes of problems an evolutionary algorithm can handle and at which scale the algorithm breaks down. Then one could compare the activity profiles of different algorithms and explore how changing the algorithm changes its activity profile. This would yield quantitative evidence about how different genetic encoding schemes or different developmental processes affect an algorithm's evolvability.

An algorithm supporting a high degree of evolvability would show significant levels of both the extent and intensity of adaptive evolution in a wide variety of contexts. This allows you to diagnose the evolutionary shortcomings of an algorithm if either kind of evolutionary activity is missing. For example, if the extent of adaptive evolution were significant but its intensity were not, then the system is missing a sufficiently creative source of adaptive innovations.

The use of evolutionary activity statistics does not depend on any prior analysis or understanding of the problem to which the evolutionary algorithms are applied. Thus, statistics like these make it feasible to automate large-scale sweeps of an algorithm's parameter space and problem space, thus enabling us to determine an algorithm's generic evolutionary capacities. Although by no means solving the whole problem of quantifying a system's evolvabilty, judicious use of evolutionary activity statistics can enable us to take a concrete step in the right direction.

Acknowledgements. Thanks to Norman Packard for long-term collaboration on evolutionary activity statistics. Thanks to Paul Marrow for comments on the manuscript. Thanks to 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, CA: Addison-Wesley.

[2] Bedau, M. A. and C. T. Brown. 1999. Visualizing evolutionary activity of genotypes. *Artificial Life* 5: 17–35.

[3] Bedau, M. A., S. Joshi, B. Lillie. 1999. Visualizing waves of evolutionary activity of alleles. *Proceed-*

ings of the 1999 GECCO Workshop on Evolutionary Computation Visualization.

[4] Bedau, M. A., E. Snyder, C. T. Brown, N. Packard. 1997. A comparison of evolutionary activity in artificial evolving systems and in the biosphere. In P. Husbands and I. Harvey, eds., *Proceedings of the Fourth European Conference on Artificial Life* (pp. 125–134). Cambridge, MA: MIT Press/Bradford Books.

[5] 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). Cambridge, MA: MIT Press/Bradford Books.

[6] Maley, C. C. 1999. Four steps toward open-ended evolution. In Banzhaf, W., 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.

[7] Rechtsteiner, A. and M. A. Bedau. 1999. A generic neutral model for measuring excess evolutionary activity of genotypes. In Banzhaf, W., 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.

[8] Shannon, T. 1998. Generic behavior in the Lindgren non-spatial model of iterated two player games. In C. Adami, R. Belew, H. Kitano, and C. Taylor, eds., *Artificial Life VI* (pp. 316–325). Cambridge, MA: MIT Press/Bradford Books.