

General Framework for Evolutionary Activity

Michael J. Raven¹ and Mark A. Bedau^{2*}

¹ NYU, 503 Silver Center, 100 Washington Square East, New York NY 10003, USA

² Reed College, 3203 SE Woodstock Blvd., Portland OR 97202, USA,

Email: mab@reed.edu Web: <http://www.reed.edu/~mab>

* To whom correspondence should be addressed.

Abstract. Evolutionary activity statistics have been used to visualize and quantify the adaptive evolutionary dynamics in a wide variety of artificial and natural evolving systems, but the formalism for the statistics has evolved over the years. Furthermore, the statistics can be applied to many different aspects of an evolving system, and application in any given context requires settling certain choices. In addition, the statistics involve normalization with a special-purpose “neutral” system, which requires making even more choices. So, to help make these statistics easier to use and understand, we situate them in a new and more general formal framework and then show how this framework applies to earlier work with the statistics.

1 The need for a framework for evolutionary activity

It is commonly accepted that the process of adaptation produces much of the order and functionality evident in complex systems [11,10,8], but it is often difficult to distinguish adaptive change in a system from other evolutionary phenomena, such as random genetic drift [9,12]. For natural systems the problem is often the unavailability of the relevant data. Those studying artificial evolving systems have the luxury of being able to collect virtually complete data; aside from storage space, only imagination limits what kinds of data are gathered. But this compounds rather than alleviates the problem, which is the inability to highlight the *relevant* data. The study of evolutionary dynamics in natural and artificial systems dearly needs an effective method for identifying and measuring the creation of adaptations in the course of evolution.

A decade ago Bedau and Packard devised a method for visualizing adaptive phenomena in evolving systems [2]. The method rests on the calculation of *evolutionary activity* statistics which can be applied to various kinds of components of evolving systems, including individual alleles [2], various classes of alleles [5], and whole genotypes [3,4]. A significant part of the appeal of evolutionary activity statistics is their applicability to a wide variety of evolving systems. Evolutionary activity has been measured in many artificial life systems, including Packard’s Bugs [3,14],

Ray’s Tierra and its derivatives [1,3,4], Lindgren’s evolving iterated prisoner’s dilemma [15], and Holland’s Echo [4]. It has also been measured in some

natural evolving systems, such as the biosphere as reflected in the fossil record [3, 4] and the evolution of technology as reflected in the patent record [16]. Comparing evolutionary activity phenomenology within or between systems shows how evolutionary phenomena vary as a function of such factors as mutation rate and mode of selection [5], and quantifying evolutionary statistics enables adaptive evolutionary dynamics in various artificial and natural systems to be directly compared [2–4, 14, 16]. Evolutionary activity statistics have also been used to study evolutionary contingency [19], punctuated equilibria [15], mutualism [13], diversity [18], and classifications of evolutionary dynamics [4, 17, 6, 7].

Nevertheless, there are barriers to realizing the full potential of evolutionary activity measurements. One is their very flexibility. Their applicability to an open-ended variety of different kinds of entities obscures the constraints on their proper use. Furthermore, the statistics have been defined in different ways in different publications, and this hides their underlying equivalence. It can be similarly unclear what unifies the different kinds of “neutral” systems that are used to normalize the statistics. So providing a consistent and general formal framework would help make evolutionary activity statistics more useful and more used. That is the purpose of the present paper.

2 A formal framework for evolutionary activity

The purpose of evolutionary activity is to measure the extent to which components of an evolving system are and have been resisting selection pressure. Persisting in the face of exposure to selection is the sign of an adaptation (though there are exceptions to this rule—see below). So the essential idea is to measure a component’s exposure to selection, because continual exposure to selection is evidence of resistance to selection. To get started on this project, one must answer three questions about the system’s components:

Question 1. What should be the components of the system?

Question 2. What should be a component’s initial activity?

Question 3. What should be a component’s current activity?

Choosing different answers to these questions allows one to examine evolutionary dynamics at different levels of analysis. We will consider each question in turn.

What should be the components of the system? An evolving system typically contains many different kinds of evolving components, e.g., individual genes, combinations of genes or schemata, individual phenotypic traits and clusters of them, whole genotypes, species and higher taxonomic groups. One might want to measure the adaptive evolution of any of these components. In general, there is no unique right choice of component to study, but some choices are wrong because some kinds of components in some systems are ill defined. After choosing what type of component to study, one must focus on which instances of

those components are present in the system at a given time. If S is an evolving system and F is a property that identifies the components of interest, then the set of components present in S at t is:

Definition 1. $C_t = \{c : c \text{ exists in } S \text{ at } t \wedge c \text{ has property } F\}$,

while the set of all such components across T , the set of all time steps, is:

Definition 2. $C = C_{t_0} \cup C_{t_1} \cup \dots$,

where $t_1, t_2, \dots \in T$.³ For example, if F is the property *being a genotype*, then C_t is the set of all genotypes extant at t and C is the set of all genotypes extant at some time or other. In most evolving systems the extant components change over time, so we have a notion of the addition (“birth” or “origination”) and subtraction (“death” or “extinction”) of a component, as follows:

Definition 3. c is *added to* C_t if and only if $c \notin C_{t-1} \wedge c \in C_t$.

Definition 4. c is *subtracted from* C_t if and only if $c \in C_{t-1} \wedge c \notin C_t$.

For example, instances of a particular allele at particular loci exist in a system at t if some agent has an instance of that allele at that locus at t , and a genotype exists in a system at t if some agent has that genotype at t .

What should be a component’s initial activity? Computing evolutionary activity involves tracking each component’s evolutionary activity over time. As a bookkeeping matter, a component’s activity is stored in its *activity counter*. When a new component is added to an evolving system, its activity counter must be initialized. In some cases the proper initial value will be the same for all components at all times, but in other cases it will depend on the context. For example, if one wished to record the activity in a component’s lineage, one might initialize a new component’s activity with the activity of its immediate ancestor (more on this later). We can represent this formally with $\Delta_{init}^C : C \rightarrow \mathfrak{R}$, where the initial value of a component c ’s activity counter is $\Delta_{init}^C(c)$.

What should be a component’s current activity? A component’s activity counter is a historical record (sum) of its activity over its entire lifetime. At each moment that the component exists, its activity counter is incremented by its current activity, i.e., its current exposure to selection. Exposure to selection can be measured in various ways; some methods are easier than others and some methods reveal selection exposure more clearly. The simplest measurement of selection exposure is always a component’s existence, but this method is crude. A more sensitive measure of a genotype’s exposure to selection is its concentration in the population, and a more sensitive measure of an allele’s exposure to selection is its expression or use. How best to measure a component’s current activity

³ We assume that time is discrete since artificial evolving systems usually assume discrete time. The formalism can be extended to continuous time.

will depend upon what one wants to learn about the target system and whether one hopes to compare this information *across* systems. A component’s current activity can be represented with the partial function $\Delta^C : C \times T \rightarrow \mathfrak{R}$, where $\Delta^C(c)$ is c ’s activity at t if c is present in S at t (and is otherwise undefined).

3 Definitions of evolutionary activity statistics

Evolutionary activity statistics aim to reflect how the evolutionary process is creating adaptations by observing resistance to selection pressures. A component of an evolving system can resist selection pressure only when it is “active” or exposed to selection. So, we assign each component an activity counter to record its exposure to selection pressure over its entire history.

Evolutionary activity and excess activity of a component. More precisely, we define the evolutionary activity (or activity counter) of component c at time t as c ’s initial activity plus the sum of c ’s activity increments up to t . Formally, the value of c ’s activity counter at t is given by $\alpha^C : C \times T \rightarrow \mathfrak{R}$:

Definition 5.
$$\alpha^C(c, t) = \Delta_{init}^C(c) + \sum_{i=Birth(c)}^t \Delta^C(c, i).$$

where $Birth : C \rightarrow T$ gives the time step at which c is added to the system.

Every time a component is exposed to natural selection, selection can provide feedback about its adaptive value. Obviously, it will not continue to be tested by natural selection unless it has passed previous tests. So, the amount that a component has been tested by selection reflects how *successfully* it has passed those tests. If a sufficiently well-tested component persists and spreads through the population, we have positive evidence that it is persisting because of its adaptive value, i.e., that it is an adaptation. But natural selection is not instantaneous. Repeated trials might be needed to drive out maladaptive components. So persistence in the face of *some* selection is no proof of being an adaptation. Thus nonadaptive items can generate “noise” in evolutionary activity data, and to gauge resistance to selection we must filter out this noise.

One way to filter the nonadaptive noise is to determine how activity would accrue if components were persisting due solely to nonadaptive factors like random drift or architectural necessity. A general way to measure the expected evolutionary activity of such nonadaptive items is to construct a *neutral variant* of the target system, that is, a system that is similar to the target in all relevant respects except that none of its components has any adaptive significance. For example, if natural selection affects only births and deaths in a target system, then a neutral system could be just like the target system except that births and deaths are the result of random rather than natural selection. (More details about neutral systems are available elsewhere [1, 3, 4, 14, 16, 5].) The accumulated activity in neutral systems provides a no-adaptation null hypothesis for the target system, which is used to screen off nonadaptive activity. If we

observe significantly more evolutionary activity in the target system than in its neutral variant, we have good evidence that this “excess” activity cannot be attributed to nonadaptive factors. That is, we have good evidence that the components with excess activity are adaptations. So, we normalize target systems by subtracting the evolutionary activity accrued in the corresponding neutral system, and call the result *excess* activity. Specifically, we define the *excess activity* $\alpha_{\text{excess}}^C : C \times T \rightarrow \mathfrak{R}$ of a component at a time as:

Definition 6.
$$\alpha_{\text{excess}}^C(c, t) = \begin{cases} \alpha^C(c, t) - \nu(c, t) & \text{if } \alpha^C(c, t) > \nu(c, t) \\ 0 & \text{otherwise} \end{cases}$$

where the function ν is determined by the specific neutral system used. As this definition indicates, the excess activity of a component is positive only if the component’s raw observed activity exceeds the value of the activity observed in the neutral system.

The choice of appropriate neutral system depends on details of the target system. The ν functions implied by three recent measurements of excess evolutionary activity illustrate some possible forms of ν . Sometimes ν is a constant function which ignores c and t , in which case the significance of a component’s activity depends only on its level and not on the identity of the particular component or time. But this need not be the case (see the third example below). The first example is from Bedau, Snyder, & Packard [4], who used a neutral system to determine a threshold a' above which activity could be regarded as more likely than not to be the result of an adaptation. In this case, ν is a simple constant function:

Example 1. $\nu(c, t) = a'$.

The second example is from Skusa & Bedau [16]. They determined ν by noting the maximum activity value that the neutral system produced for any component. If we let N^α be the set of values of activity counters of the neutral system at the end of the simulation, time t_{end} , then one can define ν as:

Example 2. $\nu(c, t) = \text{Max}(N^\alpha)$.

Third, Channon [7] uses a more fine grained neutral model, with a unique neutral component corresponding to each component in the target system, and he normalizes “on the fly” by taking the activity of each component target system and subtracting the activity of its corresponding component in the neutral system. If N_t is the set of neutral system components at t and $\eta : C \times T \rightarrow N_t$ is a function giving the neutral component $n \in N_t$ that corresponds to the $c \in C_t$, then one can define ν as:

Example 3. $\nu(c, t) = \alpha^N(\eta(c, t), t)$.

Extent and intensity of activity of an evolving system. Activity and excess activity are “micro” statistics defined for each component. It is possible to define various “macro” statistics that summarize the evolutionary activity in

a whole system. In particular, one can measure a system's *intensity* of adaptive evolution, that is, the rate at which new adaptations are being produced by natural selection. In addition, one can measure a system's *extent* of adaptive evolution. The extent and intensity of evolutionary activity are two independently varying aspects of a systems adaptive evolution. While these statistics could be applied to either raw activity or excess activity, our concern here is with the latter.⁴

The *excess extent* of evolutionary activity is given by the function $E^C : T \rightarrow \mathfrak{R}$, defined simply as the sum of the excess activity of all extant components:

Definition 7. $E^C(t) = \sum_{c \in C_t} \alpha_{\text{excess}}^C(c, t)$.

This statistic measures the total continual adaptive success of all the components in the system. The *excess intensity* is given by the function $I^C : T \rightarrow \mathfrak{R}$, defined as the number of components that have newly given evidence of being adaptations by having positive excess activity:

Definition 8. $I^C(t) = \#\{c \in C_\tau : \alpha_{\text{excess}}^C(c, \tau - 1) = 0 \wedge \alpha_{\text{excess}}^C(c, \tau) > 0\}$.

Sometimes one is interested in the mean or median of extent and intensity statistics. It is easy to define these notions.

Note that two different kinds of time indicators appear in the definition of the excess intensity of activity. This complication arises because the intensity statistic is supposed to reflect the rate at which significant adaptations are arising, but it takes some time to determine whether a component is a significant adaptation. We let τ be the time when a component first provides evidence that it is an adaptation, i.e., the time when its excess activity first become positive. And we let t represent the time when one considers a significant component to have arisen. Now, consider some component c which eventually has positive excess activity. The choice of how to define t is the choice of when to consider c to have arisen. One option is to let t be the time when c first arises in the system, i.e., $t = \text{Birth}(c)$, in which case the intensity statistic would measure the rate of origination of new components that will *eventually* provide evidence of being an adaptation. At the other extreme, one could let t be the time when c first *does* provide significant evidence that it is an adaptation, i.e., $t = \tau$, in which case the intensity statistic would measure the rate at which components (which might have arisen some time in the past) are showing they are adaptations. The best approach will depend on the details of the system under investigation. For example, Skusa and Bedau [16] let $t = \text{Birth}(c)$ because in their system excess activity first becomes positive long after the origination of a component, but Bedau, Snyder, and Packard [4] let $t = \tau$ because this time lag in their systems was minimal. However one chooses t , it is important of course to be consistent.

⁴ Here, we normalize activity before rather than after defining the macro statistics, but essentially the same result can be achieved by reversing the procedure. E.g., Rechtsteiner & Bedau [14] define excess extent after summing component activity.

4 Application to different kinds of components

Different levels of activity can be measured in one and the same system at the same time. Measurements of evolutionary activity in previous work can illustrate how the present framework applies to various kinds of components. It might also help suggest how to extend the framework to new kinds of components. (In the examples below, we indicate the kind of component in question with an index on C_t and c ; e.g. C_t^κ and c^κ for allele tokens, C_t^v and c^v for allele types, etc.

Allele tokens. An allele token is an individual allele at a locus in some particular agent. When the activity of allele tokens in asexual populations has been measured in previous work [2, 5], the concern has been to see lineages of highly used alleles. Accordingly, if an allele token was inherited from an ancestor, its activity counter was initialized to the value of the ancestral allele; otherwise, if the allele token was the product of a mutation, it was initialized to zero. So, if we let C^κ be the set of all allele tokens present at some time in the system, the activity initialization function can be defined as follows:

$$\textit{Example 4. } \Delta_{init}^{C^\kappa}(c^\kappa, t) = \begin{cases} \alpha^{C^\kappa}(A(c^\kappa), t) & \text{if } A(c^\kappa) \text{ is defined} \\ 0 & \text{otherwise} \end{cases}$$

where $A : C^\kappa \rightarrow C^\kappa$ gives c^κ 's immediate ancestor and is otherwise undefined.⁵ Then, if we let C_t^κ be the set of allele tokens present in the system at t , the activity counter of an allele token is to be incremented just in case that allele token was used or expressed, thus:

$$\textit{Example 5. } \Delta^{C^\kappa}(c^\kappa, t) = \begin{cases} 1 & \text{if } c^\kappa \in C_t^\kappa \wedge c^\kappa \text{ is used at } t \\ 0 & \text{if } c^\kappa \in C_t^\kappa \wedge c^\kappa \text{ is not used at } t \end{cases}$$

Allele types. Two agents with exactly the same kind of allele token at the same locus have the same *type* of allele at that locus. Let C_t^v and C^v be sets of allele types, defined analogously to the previous examples. An allele type, $c^v \in C^v$, is present in S at t just in case there is at least one allele token of type c^v present in some agent in S at t . The lineage of a given allele is in effect an allele type,⁶ so it makes most sense to initialize the activity of all allele types to zero: $\Delta_{init}^{C^v}(c^v) = 0$. When activity of allele types has been measured in previous work [5], the activity of an allele type c^v was defined as the sum of the activity of allele tokens c^κ of that type:

$$\textit{Example 6. } \Delta^{C^v}(c^v, t) = \sum_{c^\kappa \in c^v} \Delta^{C^\kappa}(c^\kappa, t)$$

⁵ Example 4 can be easily modified for contexts in which reproduction is sexual.

⁶ So-called ‘‘back’’ mutations create an exception to this rule, because our approach lumps together and ignores independent originations of the same allele. One could define a component that exactly corresponded to an individual allele lineage, if one wanted.

Phenotypic equivalence classes. Traits that are phenotypically the same can be grouped into *phenotypic equivalence classes*. One can measure the evolutionary activity of such phenotypic equivalence classes by attaching activity counters to them. A given phenotypic equivalence class is present at a moment if some member of the class is present at that moment. Let C_t^ϕ be the set of all phenotypic equivalence classes present at t . Phenotypic equivalence classes are like allele types in that lineages of phenotypically equivalent traits make up such classes. So it is natural to initialize the activity counters of phenotypic equivalence classes to zero: $\Delta_{init}^{C^\phi}(c^\phi) = 0$. Provided the notion of the use of a trait is well defined, it would also be natural to increment the activity of a phenotypic equivalence class c^ϕ by the use of its member traits:

$$\textit{Example 7. } \Delta^{C^\phi}(c^\phi, t) = \begin{cases} 1 & \text{if } c^\phi \in C_t^\phi \wedge c^\phi \text{ is used at } t \\ 0 & \text{if } c^\phi \in C_t^\phi \wedge c^\phi \text{ is not used at } t \end{cases}$$

See [5] for one way to measure evolutionary activity of phenotypic equivalence classes.⁷

Genotypes. It is usually easy to measure the evolutionary activity of entire genotypes. Let C_t^γ and C^γ be sets of genotypes, as in previous examples. A genotype c^γ is present in a system at a time just in case some agent in the system has c^γ as its genotype. A genotype covers all the instances of the genotype in a genotype lineage, so it is natural to initialize the activity of genotypes at zero: $\Delta_{init}^{C^\gamma}(c^\gamma) = 0$. It is also natural to increment a genotype c^γ 's activity by its concentration in the population at t , as defined by $con : C^\gamma \times T \rightarrow \mathfrak{R}$, thus:

$$\textit{Example 8. } \Delta^{C^\gamma}(c^\gamma, t) = con(c^\gamma, t)$$

where this function is defined only if $c^\gamma \in C_t^\gamma$. See [3, 4, 6, 7, 13–15, 17–19] for measurements of evolutionary activity statistics of genotypes.

5 Challenges and the future

It is straightforward to measure and interpret evolutionary activity of genotypes or phenotypic equivalence classes in almost any system. This is the most common level at which the statistics have been applied to date [3, 4, 6, 7, 13–15, 17–19]. It can be more difficult to apply the statistics at the level of genes. The evolutionary activity of individual genes is most easy to interpret when each gene has a clearly identifiable phenotypic function, because the method depends on a correlation between an component's activity and its utility for coping with selection pressure. Complications arise when genotype-phenotype mapping if many genes together affect single characters (e.g., epistasis) or if many characters are affected by single genes (e.g., pleiotropy). For instance, a gene coding for a given connection

⁷ Bedau and Raven were able to identify phenotypic equivalence classes with sets of allele tokens, because there is a one-to-one genotype-phenotype mapping in the model they studied.

strength in a neural network might influence the network's behavior in different ways depending on how activation is flowing through the other nodes. In this case, it is very difficult to decompose the network's traits and behavior and assign responsibility for different pieces of it to different genes.

Nevertheless, one can usually still find a useful way to measure the evolutionary activity of genes in systems with genetic context sensitivity. For example, although a genome encoding weights a neural net typically is highly epistatic, one could still increment the activity of individual genes by their age. Long-lived weights would tend to be those that are highly adaptive in many contexts or in a few critical contexts, and these would accrue high activity. Alternatively, one could increment the activity of a network weight whenever its activation in the network exceeds some threshold. Those highly adaptive weights through which a lot of network activation consistently flows would accrue high activity.

The most significant obstacle to applying activity statistics to natural systems is collecting enough of the right kind of data, specifically, a time series of a some sort of census of components in the system. The relative ease of collecting the necessary data is the main reason why activity statistics have been applied mostly to artificial systems to date. High throughput automated methods are currently generating massive biological data bases of various kinds (genomic, proteomic, metabolomic, etc.), and electronic media are increasingly accumulating an ever growing variety of data about the evolution of various aspects of culture (patents, financial markets, internet sites, newspapers, news groups, etc.). So our ability to measure evolutionary activity in natural systems will grow significantly in the near future.

Acknowledgements Thanks to Norman Packard and our other colleagues in the exploration of evolutionary activity statistics over the years. Thanks to Seth Bullock for suggestions about how to measure activity in evolving neural nets. And thanks to the anonymous ECAL'03 reviewers for helpful comments.

References

1. Bedau, M. A., and Brown, C. T. 1999. Visualizing evolutionary activity of genotypes. *Artificial Life* 5: 17-35.
2. Bedau, M. A., and Packard, N. H. 1992. Measurement of evolutionary activity, teleology, and life. In Langton, C. G.; Taylor, C.; Farmer, J. D.; and Rasmussen, S., eds., *Artificial Life II*, 431-461. Redwood City, CA: Addison-Wesley.
3. Bedau, M. A., Snyder, E.; Brown, C. T.; and Packard, N. H. 1997. A comparison of evolutionary activity in artificial evolving systems and the biosphere. In Husbands, P., and Harvey, I., eds., *Proceedings of the Fourth European Conference on Artificial Life*, 125-134. Cambridge, MA: MIT Press.
4. Bedau, M. A.; Snyder, E.; and Packard, N. H. 1998. A classification of long-term evolutionary dynamics. In Adami, C.; Belew, R.; Kitano, H.; and Taylor, C., eds., *Artificial Life VI*, 228-237. Cambridge, MA: MIT Press.
5. Bedau, M. A. and Raven, M. J. 2002. Visualizing adaptive evolutionary activity of allele types and of phenotypic equivalence classes of alleles. In Bilotta, E.; Gross, D;

- Smith, T.; Lenaerts, T.; Bullock, S.; Lund, H. H.; Bird, J.; Watson, R.; Pantano, P.; Pagliarini, L.; Abbass, H.; Standish, R. ; and M. Bedau, eds., *Workshops from the Artificial Life VIII Conference*, 119-130. University of New South Wales, Sydney.
6. Channon, A. D. 2001. Passing the ALife test: Activity statistics classify evolution in Geb as unbounded. In Kelemen, J., and Sosik, P., eds., *Advances in Artificial Life*, 417-426. Heidelberg: Springer-Verlag.
 7. Channon, A. D. 2002. Improving and still passing the ALife test: Component-normalised activity statistics classify evolution in Geb as unbounded. In Standish, R.; Bedau, M. A.; Abbass, H. A., eds., *Artificial Life VIII*, 173-181. Cambridge, MA: MIT Press.
 8. Dawkins, R. 1987. *The Blind Watchmaker: Why the evidence of evolution reveals a universe without design*. New York: Norton.
 9. Gould, S. J., and Lewontin, R. C. 1979. The spandrels of San Marco and the Panglossian paradigm: a critique of the adaptationalist programme. *Proceedings of the Royal Society B* 205: 581-598.
 10. Holland, J. H. 1992. *Adaptation in Natural and Artificial Systems: An introductory analysis with applications to biology, control, and artificial intelligence*, 2nd edition. Cambridge, MA: MIT Press/Bradford Books.
 11. Maynard Smith, J. 1975. *The Theory of Evolution*, 3rd edition. New York: Penguin.
 12. Mayr, E. 1988. *Towards a New Philosophy of Biology*. Cambridge, MA: Harvard University Press.
 13. Pachepsky, E.; Taylor, T.; and Jones, S. 2002. Mutualism promotes diversity and stability in a simple artificial ecosystem. *Artificial Life* 8: 5-24.
 14. Rechtsteiner, A. and Bedau, M. A. A generic model for measuring excess evolutionary activity. In Banzhaf, W.; Daida, J.; Eiben, A. E.; Garzon, M. H.; Honavar, V.; Jakiela, M.; and Smith, R. E., eds., *GECCO-99: Proceedings of the Genetic and Evolutionary Computation Conference*, Vol. 2, 1366-1373. San Francisco, CA: Morgan Kaufmann.
 15. Shannon, T. 1998. Generic behavior in the Lindgren non-spatial model of iterated two-player games. In Adami, C.; Belew, R.; Kitano, H.; and Taylor, C., eds., *Artificial Life VI*, 316-325. Cambridge, MA: MIT Press.
 16. Skusa, A. and Bedau, M. A. 2002. Towards a comparison of evolutionary creativity in biological and cultural evolution. In Standish, R.; Bedau, M.; and Hussein, A., eds., *Artificial Life VIII*, 233-242. Cambridge, MA: MIT Press.
 17. Standish, R. K. 2000. An Ecolab perspective on the Bedau evolutionary statistics. In Bedau, M. A.; McCaskill, J. S.; Packard, N. H.; and Rasmussen, S., eds., *Artificial Life VII*, 238-242. Cambridge, MA: MIT Press.
 18. Standish, R. K. 2002. Diversity evolution. In Standish, R.; Bedau, M. A.; and Hussein, A., eds., *Artificial Life VIII*, 131-137. Cambridge, MA: MIT Press.
 19. Taylor, T. and Hallam, J. 1998. Replaying the tape: An investigation into the role of contingency in evolution. In Adami, C.; Belew, R.; Kitano, H.; and Taylor, C., eds., *Artificial Life VI*, 256-265. Cambridge, MA: MIT Press.