Visualizing Waves of Evolutionary Activity of Alleles

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

Shareen Joshi
Santa Fe Institute
1399 Hyde Park Road
Santa Fe NM 87501
shareen@santafe.edu

Benjamin Lillie
Reed College
3203 SE Woodstock Blvd
Portland OR 97202
lillieb@reed.edu

We illustrate a method for visualizing adaptive evolutionary phenomena in evolving systems [3]. The method was originally illustrated very briefly and abstractly at the level of alleles [3], and it has subsequently been applied in great detail at the level of whole genotypes [2]. Here we apply the method in some significant detail to alleles in three different evolving systems: a model of the evolution of sensory-motor strategies, a model of traders buying and selling securities in a financial market using an evolving set of market-forecasting rules, and an analogue of the financial market model in which natural selection is replaced by random selection. The underlying hypothesis behind the visualization method is that "activity wave diagrams" highlight the quality of the main adaptive events and adaptive phenomena in an evolving system. This abstract contains wave diagrams showing a variety of evolutionary phenomena such as competetive exclusion, cooperation, and frozen accidents.

The evolutionary activity statistic we visualize are computed from data obtained by observing an evolving system. (More details are available elsewhere [3, 4, 2].) We view an evolving system as a population of components participating in a cycle of birth, life and death, with each 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. Our purpose for using the evolutionary activity statistic is to identify innovations that persist and continue to be significant. 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 activity count for an entire lineage

To collect activity statistics, one must settle two questions: (1) What should be counted as a component of the system, and when are these added or subtracted from the system? (2) What should be a new component's initial activity and how should this change over time? In all the figures shown here, the components were chosen to be the evolving alleles encoding behavioral rules used by the agents in the systems under investigation, and such alleles are added to the system with a new mutation and removed from the system through the extinction of all genotypes containing the allele. A rule's initial activity is zero, and it is incremented by one each time the rule is actually used by an agent in the system.

References

[1] Arthur, W. B., J. H. Holland, B. LeBaron, R. Palmer, and P. Tayler. 1997. Asset pricing under endogenous expectations in an artificial stock market. In W. B. Arthur, D. Lane, and S. N. Durlauf, eds., *The Economy as an Evolving, Complex System II* (pp. 15–44). Menlo Park: Addison-Wesley.

[2] Bedau, M. A., and C. Titus Brown. 1999. Visualizing evolutionary activity of genotypes. To appear in *Artificial Life*. Santa Fe Institute Working Paper No. 98-03-023.

[3] 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, Calif.: Addison-Wesley.

[4] Bedau, M. A., S. Snyder, N. H. Packard. 1998. A classification of long-term evolutionary dynamics. In C. Adami, R. K. Belew, H. Kitano, and C. E. Taylor, eds., *Artificial Life VI* (pp. 228–237). Cambridge, Mass.: MIT Press.

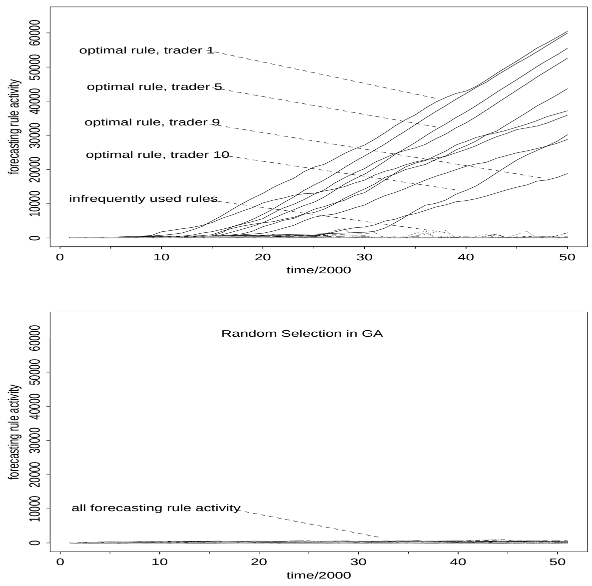

Figure 1: Activity waves of market-forecasting rule alleles in the Santa Fe Artificial Stock Market, an agent-based model of a financial market in which risk-averse traders buy and sell stocks and bonds on the basis of price forecasts they make using market-forecasting rules [1]. The individual alleles in this model genetically encode the market-forecasting rules which comprise a trader's forecasting methods. A forecasting rule's activity is incremented whenever a trader uses it to forecast the market and thereby decide to buy or sell in the market. Each trader's set of market-forecasting rules evolves through the operation of a genetic algorithm (GA). In the simulation shown here, the GA's rate of evolution is moderate (the GA runs a thousand times in the course of the simulation). There are ten traders in this market, and each has one hundred forecasting rules at any given moment. Above: The GA operates normally, i.e., natural selection determines which rules "reproduce" and which rules "die" over time. Below: A "neutral shadow" [4, 2] of the Artificial Stock Market, in which the GA's birth and death processes occur throuth *random* selection. Comparing these two activity wave diagrams shows the dramatic effect of natural selection on the rules the traders use. When natural selection operates normally, each trader independently learns the optimal market forecasting rule somewhere in the middle of the run. Once this optimal rule is learned, it is almost always subsequently used. Note that random selection yields no genotypes with anything like the persistence of the optimal rule under natural selection.

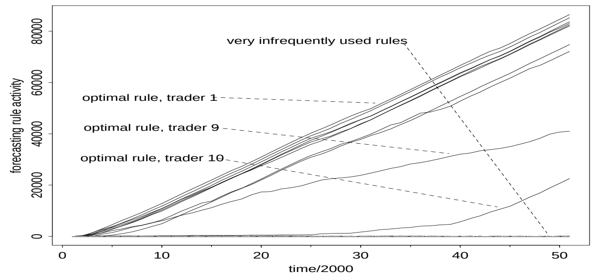

Figure 2: Activity waves of market forecasting rule alleles (as in Figure 1 except that the activity scale is approximately doubled) when the GA operates very frequently (the GA runs ten thousand times in the course of this simulation). Here, the GA operates normally, using natural selection. The activity wave diagram shows that each trader almost immediately learns the optimal forecasting rule and subsequently uses virtually just that rule. One trader takes significantly longer to learn the optimal rule.

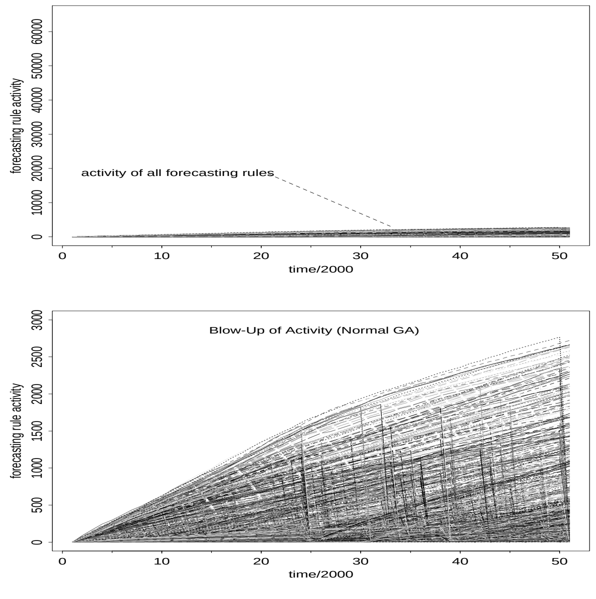

Figure 3: Activity waves of market forecasting rule alleles (as in Figure 1) with a very infrequent GA (the GA runs ten times in the course of the whole simulation). Above: Plotted on the same scale as Fig. 1. Below: A blow-up of activity, with activity plotted on a scale that is an order of magnitude smaller. Note that, when the GA is very slow, no rule is used vastly more than all other rules; i.e., the GA is too slow for the traders to latch onto the optimal rule.

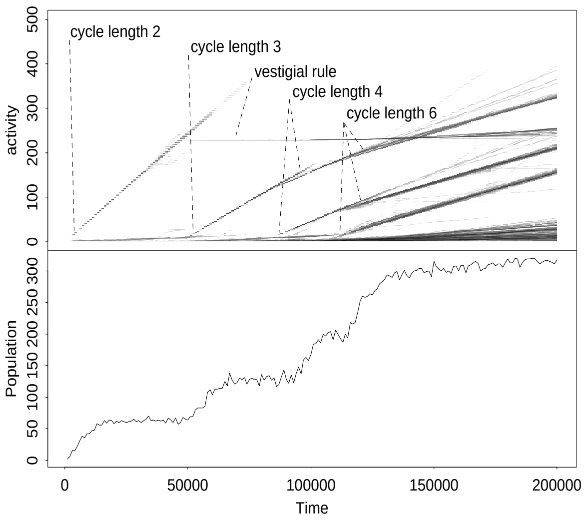

Figure 4: Activity waves of sensory-motor rule alleles and population level in Packard's Bugs—a model of the evolution of sensory-motor strategies in a population of agents competing for energetic resources in a space-limited two-dimensional world [3]. An individual allele genetically encodes a single sensory-motor rule in an agent's sensory-motor strategy. A rule's activity is incremented each time the agent acts on a rule. In this model, successful strategies take the form of short cycles, repeatedly firing a precise sequence of sensory-motor rules. In general, a cycle's *length* is proportional to its fitness. A wave's slope is inversely proportional to the genotype's cycle length. All the evident major adaptive events in the wave diagram are increases in cycle length. The first significant wave in the activity diagram shown here corresponds to a cycle of length two. The next major adaptive innovation (visible in the population graph) is the introduction of a cycle of length three, which corresponds to the second major wave (at slightly lower slope, because the cycle length is longer). The next major adaptive innovation (visible in the population graph) is the introduction of a cycle of length four. This corresponds to the third major wave *and* to the changed slope in the second major wave. This shows that the length-four cycle incorporates some of the same sensory-motor rules involved in the length-three cycle. A similar thing happens with the next major innovation—a length-six cycle that reuses rules from the length-three and length-four strategies. Note also the flat wave left by one of the rules in the original length-two cycle—a vestigial rule, unused (unexpressed) after the advent of the length-three cycle but still in the gene pool.

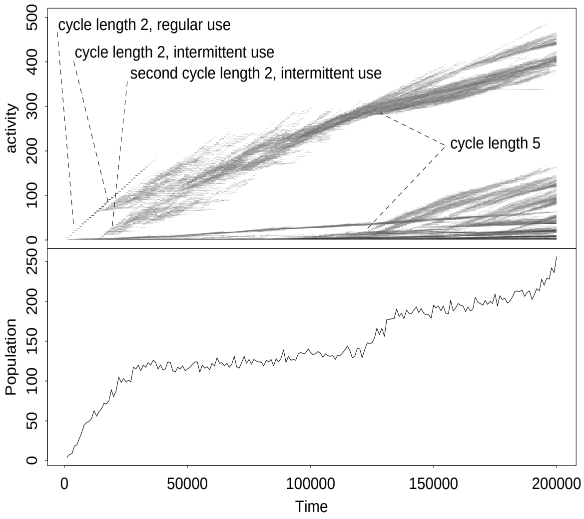

Figure 5: Activity waves of sensory-motor rule alleles and population level in another run of Packard's model (see Fig. 4 caption). This run start with a length-two cycle, and this strategy quickly spreads through all the blocks in the Scatter world. Very early on in the run a second length-two cycle emerges. (This happens so quickly that the population level has no time to level off with the initial length-two cycle.) As it happens, agents can express these two length-two cycles on the same tiny block without the agents colliding. What happens, then, is that agents on a given block express one of the length-two cycles for a while until they are "bumped" into the other length-two cycle. Thus, instead of regularly exercising a given cycle, they exhibit a pattern of switching back and forth between two length-two cycles. Thus, at any given time, the sensory-motor rules in one cycle are being expressed (at the characteristic frequency of length-two cycles) while the rules in the other cycle are not being used at all. All that changes is that at random intervals there is a switch between which cycle is expressed and which is unexpressed. The net effect is that all rules are used *on average* with the characteristic frequency of a length-four cycle; this is evident in the relatively "cloudy" waves left by the co-existing length-two cycles. Eventually (at little after the middle of this simulation), a length-five cycle is introduced and out-competes the pair of co-existing length-two cycles. The fact that one of the length-five cycle waves grows out of the length-two cycle waves shows that some of the rules in the length-five cycle were used in the length-two cycles.