Visualizing Adaptive Evolutionary Activity of Allele Tokens and of Phenotypic Equivalence Classes of Alleles

Mark A. Bedau* and **Michael J. Raven**
Reed College, 3203 SE Woodstock Blvd., Portland OR 97202, U.S.A.

*To whom correspondence should be addressed. Email: mab@reed.edu. Web: www.reed.edu/˜mab.

Abstract

We present a method for visualizing the important adaptive dynamics in evolving systems. The method consists of measuring evolutionary activity, (as conceived by Bedau and Packard) of components of evolving systems and then plotting this as a function of time. This method has already been applied to individual alleles and whole genotypes. We extend the method to phenotypically equivalent types of alleles and apply it to data generated by Packard's Bugs model, a simple artificial system that consists of agents with sensorimotor genes competing for resources in a two-dimensional world. One novelty of the Packard model studied here is that it allows the evolution of sensory thresholds—a simple form of the evolution of evolvability. Plotting the evolutionary activity of phenotypically equivalent types of alleles provides an especially vivid empirical quantitative picture of the significant adaptive phenomena in an evolving system.

Visualizing adaptive dynamics

It is commonly accepted that the process of adaptation produces much of the order and functionality evident in complex systems (Maynard Smith 1975; Holland 1992; Dawkins 1979; 1987; Holland 1995), but it is often difficult to distinguish adaptive change in a system from other evolutionary phenomena, such as random genetic drift (Gould & Lewontin 1979; Mayr 1988; Burian 1992; West-Eberhard 1992). For natural systems the problem is often the unavailability of the relevant data. Those studying models of artificial 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 of filtering this data to find a picture that reveals a system's significant adaptive phenomena. For the computational systems studied in artificial life, the problem is not a shortage of data but the inability to highlight the *relevant* data. The study of evolutionary dynamics in natural and artificial systems dearly needs an effective method for visualizing 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 (Bedau & Packard 1992). The method rests on the calculation of *evolutionary activity* statistics. The underlying hypothesis behind the method is that plotting evolutionary activity highlights a variety of adaptive evolutionary phenomena such as competitive exclusion, cooperation, frozen accidents, and neutral genetic drift. This method can be applied to various kinds of components of evolving systems, including individual alleles (Bedau & Packard 1992; Bedau *et al.* 1999) and whole genotypes (Bedau & Brown 1999). It can also be applied to data generated by a variety of artificial and natural evolving systems, including Tom Ray's Tierra and its derivatives (Bedau 1996; Bedau *et al.* 1997; 1998), John Holland's Echo (Bedau *et al.* 1998), as well as the fossil record (Bedau *et al.* 1997; 1998). Comparing evolutionary activity waves within or between systems can show how these evolutionary phenomena vary as a function of time, space, mutation rate, mode of selection, or other factors. In addition, the data displayed in evolutionary activity graphs can be quantified with various statistics (Bedau & Packard 1992; Bedau 1995; Bedau *et al.* 1997; Bedau & Brown 1999; Rechtsteiner & Bedau 1999), thus enabling evolutionary activity in various artificial and natural systems to be directly compared.

Here we apply the method to data generated by Packard's Bugs model—a simple artificial two-dimensional system for the evolution of sensorimotor strategies. The Packard model studied here has the novel feature of allowing for the evolution of sensory thresholds—a simple form the evolution of evolvability or, as we term it, "meta-evolution"—and we use evolutionary activity to visualize evolution and meta-evolution simultaneously. The evolution of evolvability is likely to be an especially interesting and important mechanism in the unbounded creativity of evolutionary systems (Bedau 2002). But the primary novelty of the present paper is to extend the evolutionary activity method to a new and especially significant component of evolving systems: phenotypically equivalent types of alleles. The central result of this exercise is to underscore how evolutionary activity highlights the existence and quality of the main adaptive events and adaptive phenomena in evolving systems.

Packard's Bugs model

All of our observations are from computer simulations of a certain model—originated by Norman Packard (Packard 1989; Bedau & Packard 1992)—that is designed to be a very simple model of the evolution of sensorimotor strategies. Packard's model consists of agents sensing the resources in their local environment, moving as a function of what they sense, ingesting the resources at their current location, and reproducing or dying as a function of their internal resource levels. The model's spatial structure is a grid of sites with periodic boundary conditions, i.e., a toroidal lattice. In the versions of Packard's model that we study here, resources are immediately replenished at a site whenever they are consumed. The agents constantly extract resources and expend them by living and reproducing. Agents ingest all of the resources (if any) at their current location and store them internally.[1] Agents expend resources at each time step by paying (constant) "existence taxes" and (constant) "movement taxes".[2] If an agent's internal resource supply drops to zero it dies.

Each agent moves each time step as dictated by its genetically encoded sensorimotor map: a table of behavior rules of the form IF (environment j sensed) THEN (do behavior k). An agent receives sensory information about the existence of resources (but not the other agents) in the von Neumann neighborhood of five sites centered on its present location in the lattice. Thus, each sensory state j corresponds to one of $2^5 = 32$ different detectable local environments. Each behavior k is a jump vector between one and sixteen sites in any one of the eight compass directions, so an agent's genome continually instructs it to move somewhere. Thus, an agent's genotype represents a sensorimotor map, which is just a lookup table of sensorimotor rules.

An agent reproduces (asexually, without recombination) if its resource reservoir exceeds a certain threshold. The parent produces one child. The child starts life at the unoccupied site nearest to its parent with one half of its parent's resource supply. The child also inherits its parent's sensorimotor map (and sensory threshold), except that point mutations can replace the behaviors linked to some sensory states with randomly chosen behaviors (the sensory threshold can be mutated analogously). The sensorimotor mutation rate is controlled by a model parameter and is set here so that the probability of a mutation at each sensorimotor locus is $p_{\mathrm{sm}} = 0.05$. There is thus roughly one sensorimotor mutation per reproduction event, on average.

The agents' genotypes also include a special *sensory threshold gene* that controls an agent's sensory discriminability. The resource level at any given site is an integer between 0 and some maximum value, and an agent's *sensory threshold* determines the resource level that it can discriminate at the sites. If an agent's sensory threshold is set to 100, for example, then it can distinguish between those sites with resource level less than 100 and those sites with resource level 100 or greater. The sensory threshold gene can mutate during reproduction. Usually we fixed this gene at 1 by setting the sensory threshold mutation rate to $p_{\mathrm{th}} = 0.0$. In this case the agents can distinguish between no resources and some resources. In one specific environment (the Ridge, discussed below) we made the sensory threshold mutation rate positive ($p_{\mathrm{th}} = 0.2$), so the agents' sensory thresholds could adapt.

An agent's sensory threshold determines the *semantics* of the agent's sensory experience. For example, a given sensory experience could mean that the resource level at a site either is simply positive or is above 100, depending on the sensory semantics. So the sensory semantics sets the meaning of all the sensorimotor genes in the agents' behavioral strategies. One sensory semantics might afford much more useful information than another about a given environment. Thus, the ability of the agents to evolve useful sensorimotor strategies depends on their sensory thresholds. If their sensory thresholds obscure useful information about the environment, natural selection cannot create very adaptive behavioral strategies because no such strategies exist. But if the thresholds highlight key environmental distinctions, then highly adaptive behavioral strategies exist for evolution to find. Thus, changes in the sensory semantics affects the ability of evolution to craft adaptations. In other words, mutation of the sensory-threshold gene opens the door to the evolution of evolability, or "meta-evolution." Meta-evolution is the interweaving of two simultaneous and connected evolutionary processes: the evolution of sensorimotor genes and of the semantics that determines their meaning.

A time step, or update, in the simulation cycles through the entire population and has each agent, in turn, complete the following sequence of events: sense its present von Neumann neighborhood, move to the new location dictated by its sensorimotor map (if that site is already occupied, it randomly walks to the first unoccupied site), consume any resources found at its new location, expend resources to cover existence and movement taxes, and then, if its resource reservoir is high enough

---

[1] Note that in some environments it is possible for bugs to ingest the resources at their current location even if they are unable to sense the existence of any resources at that location. The Ridge environment, which shall be described later, is just such an environment.

[2] Usually the "movement taxes" have been variable and proportional to the distance moved. We set the "movement tax" to a constant value so as not to bias selection toward bugs that do not move very far.

or empty, either reproduce or die, respectively. A given simulation starts with randomly distributed agents containing randomly chosen sensorimotor strategies. The model contains no *a priori* fitness function (Packard 1989). Agents with maladaptive strategies find few resources and thus die, taking their sensorimotor genes with them; by contrast, agents with adaptive strategies tend to find sufficient resources to reproduce, spreading their strategies (with mutations) through the population.

Tokens, types, and phenotypic equivalence classes

We distinguish particular instances or "tokens" of an allele at a locus and "types" of alleles at that locus. This distinction applies to any gene, including both sensorimotor genes and sensory threshold genes. An allele *token* is the allele at a particular locus in some particular agent. If a second agent has exactly the same kind of allele at the same locus, even if that allele arose from an independent mutation in a different lineage, it has a different token of the same *type* of allele.

Sensorimotor alleles code for individual sensorimotor rules. A *type* of sensorimotor allele is a particular size jump in a particular direction given a particular local sensory environment. Sensory threshold alleles code for an integer sensory threshold value between 0 and the maximum resource level. A *type* of sensory threshold allele is simply such an integer.

Natural selection acts on phenotypes in the first instance, so we want to identify classes of allele tokens that are phenotypically equivalent. This is trivial to do for sensorimotor alleles since all tokens of the same sensorimotor allele type are phenotypically equivalent. They all code for the same kind of behavior in the same sensory context. Identifying classes of phenotypically equivalent sensory threshold allele tokens is a little more complicated. One complication is that phenotypic equivalence classes of sensory threshold alleles depend on the environment. Different environments will have different phenotypic equivalence classes. For example, a Heap environment contains sites with just two resource levels, zero and the height of the heap. If h is the height of the heap, all of the positive sensory threshold values θ up to the hight of the heap, $0 < \theta \leq h$, yield exactly the same sensory information in all local Heap environments. So, although $\theta = \frac{h}{2}$ and $\theta = \frac{h}{3}$ are different allele types, they will have exactly the same phenotypic effect in the Heap environment. By contrast, Ridge environments contain several different resource levels, so there are several classes of sensory threshold tokens that are equivalent from the agents' point of view, i.e., that are phenotypically equivalent. So, to identify phenotypic equivalence classes for sensory threshold alleles in a given environment, we explicitly calculate which ranges of threshold values yield equivalent sensory information.

Three kinds of environments

Here we restrict resources to three kinds of distributions:

- **Heap** environments consist of a single large (31×31) block of resources surrounded by a resource-free desert in a landscape with 127×127 sites.
- **Archipelago** environments consist of 100 small (3×3) resource blocks randomly sprinkled across a resource-free desert in a 127×127 landscape.
- **Ridge** environments consist of a band (15×50) of resources with a triangular cross-section that circles a resource-free desert in a 100×50 landscape.

All of these environments consist of certain sites with resources surrounded by a resource-free desert.

In Heap environments, all evolution concerns strategies for exploiting the resources on the one resource block. Since the width of a block is above 3, the agents' sensory information is ambiguous (Fletcher *et al.* 1996). Agents cannot always tell exactly where they are, for the sites inside the block all look alike, as do the sites along the same edge. Agents in a Heap tend to exhibit a fluidly changing ecosystem of coexisting strategies. These strategies are virtually all edge strategies, i.e., a strategy consisting of moving in a straight line when inside the Heap and jumping back into the middle of the Heap when detecting the edge. A proliferating edge strategy will fill a region of the Heap with a perpetually rolling population. Heaps have two main kinds of adaptive events: lengthening an edge strategy's jump back and discovering a new edge strategy compatible with the ecology of existing strategies.

If an agent is on a 3×3 block in an Archipelago environment, the agent's sensory information always unambiguously indicates its exact location on the block (NW corner, middle, N edge, etc.).[3] Each 3×3 block in an Archipelago environment is an evolutionary island supporting its own evolutionary development and subject to migration from nearby islands. So a behavioral innovation that originates on one island can hop from island to island and eventually colonize the entire archipelago. Agents on a given 3×3 block usually all follow the same strategy, for agents following heterogeneous strategies tend to collide and be bumped into the resource desert. The strategies observed are cycles jumping through a sequence of sites on the block. The simplest cycles (period 2) consist of jumping back and forth between two sites. Since a 3×3 block contains 9 distinct sites, it can support at most a period 9 strategy. A period n strategy has room for at most $n-1$ agents (one agent in the

---

[3] Actually, there is some sensory ambiguity in the Archipelago. The little blocks are scattered randomly in the environment and they sometimes overlap, so overlapping blocks can create little local regions of sensory ambiguity.

cycle must move first and the space to which it is jumping must be unoccupied). Thus, longer period strategies can support larger populations because they can exploit more of the energetic resources on a block. All agents reproduce at the same rate, so a block with a larger population will produce offspring at a higher rate. Thus, blocks with larger period strategies will exert greater migration pressure and thus will have a selective advantage. So evolution in an Archipelago of tiny 3×3 resource islands will exhibit one main kind of adaptive event: lengthening the period of an existing strategy. More information on the strategies that evolve in Heap and Archipelago environments is available elsewhere (Bedau 2001).

In the Ridge environment, there is a smooth gradient of resource levels rising from one edge of the Ridge at a constant slope up to the Ridge top and then falling at the same rate to the other edge of the Ridge. The width of the Ridge is 15 sites, so there are 8 different resource levels along each side of the Ridge. Resource levels change as agents cross the Ridge, but they remain constant as agents move along the Ridge at a constant distance from the edge. Since agents can detect only whether resources at neighboring sites are above or below their sensory threshold, there is rampant sensory ambiguity both across and along the Ridge. In Ridge worlds we let the agents' sensory threshold genes evolve, to see how their sensory resolution of the world would adapt. As their sensory thresholds evolve, their perception of the edge of the Ridge changes. We set taxes in Ridge worlds high enough that agents can survive only if they remain at the spine of the Ridge. Successful adaptation then involves two things: evolving a sensory threshold that can discriminate the spine, and evolving a sensorimotor strategy that runs along the spine and jumps back onto the spine if an edge is detected.

Neutral models of Packard's Bugs

In order to assure ourselves whether the evolutionary activity of some component reflects the component's adaptive value, we also measure activity in *neutral models* paired with the target. The crucial property of a neutral model is that its evolutionary dynamics are like its target model *except* that neutral components' activity cannot be due to their adaptive significance—for they have no adaptive significance.

Our neutral models here are a sequence of instances of a generic neutral model (Rechtsteiner & Bedau 1999). The generic model consists of a population of individuals that reproduce and die in a fixed genotype space. The genotype space is defined by some number of loci at each of which some number of alleles are segregating. Parameters that need to be specified in the generic neutral model are N, the size of the population of individuals, r, the reproduction rate (the number of individuals that die and reproduce per time step), l, the number of loci, a, the number of possible alleles per locus, and m_l, the probability that the allele at a given locus will be mutated when an individual is born. (The probability that an offspring will have mutation somewhere in its genome, i.e., the mutation rate per individual is $m_i = 1 - (1 - m_l)^l$.) The parameters together determine the model's generic behavior. The genotype space is a hypercube of dimension l and size a^l (number of possible genotypes), with each location in this space corresponding to a given genotype. The current state of the model is given by the distribution of N individuals in genotype space. The population wanders through the space stochastically, spreading and clustering at random.

The individuals in the initial population are assigned genotypes at random. Time is discrete, and moves forward each time step by iterating the following two-step algorithm: (1) r individuals (selected at random, with replacement) each produce a child that is genetically identical to itself except for mutations. Mutant alleles are chosen at random from the set of possible alleles. (2) r individuals (selected at random, without replacement) die and are removed from the population and are replaced by the r children produced at step (1).

A specific instance of the generic neutral model corresponds most closely to systems in which the generic model parameters are relatively constant. However, Packard's Bugs model with Heap, Archipelago or Ridge environments typically produce widely varying population sizes and reproduction rates. So, to create a neutral model for such target systems, we divide the target system behavior into a sequence of periods in which the population size and reproduction rate are relatively constant. Then we create a parallel sequence of neutral models with parameters that match those periods.

Definition of evolutionary activity

The evolutionary activity statistics we visualize are computed from data obtained by observing an evolving system. Here we give a brief overview and explain how to apply the method to the cases at hand. Extensive details about evolutionary activity statistics are available elsewhere (Bedau & Packard 1992; Bedau 1995; Bedau *et al.* 1998; Bedau & Brown 1999; Rechtsteiner & Bedau 1999; Skusa & Bedau 2002).

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 evolutionary activity statistics is to identify innovations that persist and continue to be significant. The fundamental idea behind evolutionary activity of a component of an evolving system is to identify which components make a difference in the evolutionary process. Generally we consider a component to "make a difference" if it continues to be active in the evolving system. To implement these ideas, counters are attached to components for bookkeeping purposes, to reflect each component's current activity throughout its lifetime. If components are inherited during reproduction, the corresponding counters are inherited along with them, maintaining an increasing activity count for an entire lineage.

To collect activity statistics, one must settle two questions: (1) What components of the system are to be examined? (2) How should a component's activity be incremented over time? The relative adaptive significance of different components is reflected and so can be operationalized in different ways, always pursuing the goal of having activity measure a component's adaptive value. Here we examine various components of Packard's Bugs model and increment their activity as follows:

- **Sensorimotor tokens and types.** The adaptive value of a sensorimotor gene is reflected by the number of times the gene has been used in the history of its lineage. The more a gene has been used, the more evidence we have that it is adaptive and persisting because of its adaptive value. So we define the activity of a token of some sensorimotor allele as the count of each time the allele is used. I.e., each time a sensory state is detected and the corresponding behavior is produced, the activity counter for that allele at that locus is incremented by one. If a child has a mutation at a given locus, then the activity counter for the allele token at that locus is set to 0. If a child inherits a sensorimotor allele from a parent, they will have two tokens of the same type of allele at some locus, and each allele token will have its own activity value. The initial activity value of the child's allele token is inherited from its parent's activity value, so an allele token's activity is the number of times that allele has been used in the lineage tracing from the child to its parent, its parent's parent, etc., all the way back to the ancestor in which that allele arose as a mutation.

  The adaptive value of a *type* of sensorimotor gene is best reflected by the adaptive value of its tokens. So we define the activity of a type of a sensorimotor allele at a given locus as the sum of the activity of all the tokens of that type at that locus. Think of the family tree showing the relations among all allele tokens of the same type at a given locus. To collect activity for a sensorimotor *token* one has a separate activity counter for each branch of the family tree. To collect activity for a sensorimotor *type* one has one counter covering all of the family trees for the same type of allele at each locus. The activity value of a given type of sensorimotor allele at a given locus reflects the total number of times that type of allele at that locus has been in continuous use. If all tokens of an allele type at a locus vanish from the population at some time, the activity counter for that allele type is reset to 0. If the same type of allele later reappears at the same locus through a new mutation, the activity counter restarts anew.

- **Sensory threshold tokens and phenotypic equivalence classes.** An agent's sensory threshold gene is used all the time, since agents sense their neighborhoods all the time. So we define the activity of a token of a sensory-threshold allele simply as the age of the allele. When sensory threshold tokens are phenotypically equivalent in a given environment we group them into equivalence classes and define the activity of a phenotypic equivalence class of sensory threshold allele tokens as the sum of the activity of all the tokens of phenotypically equivalent threshold allele tokens.

Although the details of evolutionary activity implementations vary from context to context, the overall motivation is always the same: to have some measure of the extent to which the item under investigation has been subjected to selection pressure. Individual alleles become visible to natural selection only when they are phenotypically expressed, which in the present context means when they have some impact on an agent's behavior. A sensorimotor gene in Packard's model impacts an agent's behavior each time the agent uses the rule for which that allele (and it alone) codes. Hence we increment sensorimotor allele activity by allele use. A sensory threshold allele impacts behavior each time *any* rule is used, i.e., at each time step. So we increment sensory threshold allele activity at each moment when the allele continues to exist.

Evolutionary activity plots

We collected evolutionary activity data both of sensorimotor allele tokens and types, and of sensory threshold tokens and their phenotypic equivalents, in a variety of Block, Archipelago, and Ridge environments. When these kinds of evolutionary activity are graphed, adaptively significant components show up as a salient line or "wave." Allele token activity plots provide a detailed picture of the myriad adaptive phenomena that occur in the course of evolution. Plots of the activity of phenotypically equivalent allele types provide clear and concise pictures of just the key adaptive phenomena. The following subsections illustrate these points with typical examples of evolution in Block, Archipelago, and Ridge environments.

Evolutionary activity in a Heap

The top of Figure 1 shows the population size over the first 16,000 time steps of a typical Heap run. Three dramatically different epochs are discernible. The first consists of the first 7000 time steps during which the population level rises to about 100. The second lasts from 7000 to about time 12,000, during which the population shoots up to 400 and then climbs over 600. The third epoch lasts for the rest of the time period shown and starts when the population level shoots up to about 800.[4]

The middle of Figure 1 shows the activity of all the tokens of the sensorimotor alleles during the same time period. The graph consists of a bunch of (mostly) straight lines or "waves" of different duration and slope, showing the dynamics of activity accumulation of the sensorimotor alleles. Since the activity of a sensorimotor allele is incremented each time that allele is used, the activity waves show the patterns of repeated use of the alleles in the population. The most heavily used alleles create waves with the highest slope. Slope directly reflects frequency of use, and change in slope reflects change in frequency of use.

Comparison with the activity generated in the neutral model (green waves along the bottom) confirms that the significant waves in the target system accrue much more activity than would ever be expected in a similar neutral model driven by random selection. So those significant waves are produced by adaptations. The long waves along the bottom with very low slope are *also* clearly produced by adaptations. The key property of a neutral model wave is not its height but its duration (projection onto the x-axis), and the long but low waves last much longer than any in the neutral model. Blowing up the bottom of the allele token activity plots would reveal a swarm of very short-lived waves caused by maladaptive genes that are quickly driven from the population. (It would also reveal some short-lived waves caused when agents are bumped into the resource desert.)

The edge strategies that evolve in Heap environments involve a pair of genes: one for moving across the middle of the Heap until hitting one of the edges and the other for returning to the middle of the Heap upon detecting the edge. Thus, an edge strategy creates a pair of waves, one for each allele. The slopes of the two waves sum to 1, and their respective slope values reflect how far back into the middle of the Heap the agents jump at the edge and their step size as they move across the middle of the Heap. For example, if they jump back 3 sites and they move across the middle one site at a time, then their respective slopes are $\frac{3}{4}$ and $\frac{1}{4}$. First, the V-shaped pair of waves that arise from the origin of the activity token waves in the middle of Figure 1 is just such a pair of waves. The lower wave shows the use of the edge-jump gene, and the higher wave shows the use of the middle-traverse gene. This pair of waves signals the first epoch, and it shows it to be dominated by an edge strategy. Since the slopes of the two waves are $\frac{4}{5}$ and $\frac{1}{5}$, we can see that this edge strategy jumps back once every five time steps. (There is additional wave phenomenology during the first epoch involving waves splitting off from the main pair of waves as new waves arise. We omit detailed analysis of most of these waves here, but we will return to them in a moment.)

The start of the second epoch is signaled in the sensorimotor activity plot by the termination of most of the waves in the first edge strategy coinciding with the initiation of a huge new wave at about time step 7000. This wave has slope $\frac{9}{10}$. Less visible at the bottom of the plot is the simultaneous initiation of a second wave with very low slope of $\frac{1}{10}$. This new, wider V-shaped pair of waves shows the origin of a second edge strategy, one which jumps back once every ten time steps. By jumping farther into the middle of the block this second strategy can extract more of the resources in the Heap, so it drives the first edge strategy to extinction by competitive exclusion.

However, note that the wave for the edge-jump gene continues past 7000. This is because a variant of the first edge strategy employing that gene—specifically, the one that is used with the gene behind the new middle-traverse wave that starts about time step 5000—can coexist with the second edge strategy. Coexistence of V-pairs of waves signals coexistence of edge strategies. The second variant of the first edge strategy finally goes extinct a little after time step 10,000.

The start of the third epoch around time step 12,000 is signaled by the termination of many of the waves and the *increase in slope* of one of the fat waves that arises a little before time step 10,000. At the same time a new low slope wave starts. These phenomena signal the origin of a third major edge strategy. Since the waves for the third strategy coexist with those for the second strategy, we can see that the second and third edge strategies are compatible and can coexist.

These three dominant adaptive events are revealed much more clearly in the bottom of Figure 1, which shows the activity of phenotypically equivalent sensorimotor alleles. The first epoch shows significant activity only for the two waves for the genes in the first edge strategy. The origin of the high-slope wave (and the simultaneous origin of a significant low-slope wave) at time step 7000 shows the start of the second epoch. And the third high-slope wave that arises around time step 12,000 shows the start of the third epoch. The periods when these pairs of waves overlap shows when these dominant strategies are coexisting.

It is striking how sharply the *type* activity plots put their finger on the main adaptive events. The type activity plots clean up the myriad little waves that fill out the *token* activity plots. The little token activity waves that are consolidated in the type activity waves are evidently

---

[4] Two less dramatic shifts in population level occur at the start of the second epoch, around time step 8000. We omit analysis of these events.

genetic drift among neutral variants. Since they do not cause new allele type waves, they involve the origin and proliferation of tokens of alleles that are phenotypically equivalent to those already present in the population.

So, type and token activity plots are *each* informative in their own way. Type activity focuses attention cleanly on the central adaptive events. Token activity rounds out the central adaptive story with myriad adaptive sideshows and subplots that are also in play. Each type of activity provides useful and distinctive information.

Evolutionary activity in an Archipelago

The top of Figure 2 shows the population dynamics generated in a typical Archipelago environment. The first evolutionary epoch is clearly visible. It consists of the first fifth of the run and is marked by a relatively small population level, and it ends with the population explosion at time step 200,000. The second epoch lasts for the second fifth of the run, but its end is not so easy to discern. There does seem to be a distinct rise in population level at about time step 450,000, followed by another rise at around time step 650,000, and another at around time step 850,000. Thus the population level dynamics seems to suggest a succession of five distinct evolutionary epochs, but sharp epochal distinctions are hard to identify.

The sensorimotor allele token activity shown in the middle of Figure 2 helps to clarify the picture. The sensorimotor strategies in Archipelago environments consist of periodic cycles and the central form of adaptation is replacing one cycle with another that has a higher period (recall the discussion above). Cyclic strategies with period 2 will show up as a pair of waves (perhaps on top of each other) with slope of $\frac{1}{2}$, and in general period n strategies show up as n waves (some perhaps on top of others) with slope $\frac{1}{n}$. A period n strategy that evolves into a period $n+1$ strategy often reuses a number of the genes used in the period n strategy. In that case, when the period $n+1$ strategy arises the waves for those genes will have a kink as the slope of their waves changes from $\frac{1}{n}$ to $\frac{1}{n+1}$. These kinks will coincide with the start of a new wave for the new two (or more) genes that fill out the higher-period strategy.

We can see these patterns in the allele token activity waves shown in the middle of Figure 2. The first wave with slope $\frac{1}{2}$ signals a cycle with period 2 (or 2-cycle). The new wave that starts at time step 200,000 and that coincides with a kink in one of the two overlapping initial waves shows when the 2-cycle is driven extinct by a 3-cycle that arises at time step 200,000. A variety of different size waves arise over the next fifth of the run, and it is not obvious when the next major adaptive event occurs. With some careful analysis, though, the next major adaptive event can be identified as the origin of a significant wave at about time step 450,000 that coincides with a kink in the slope of the third long wave in the diagram. This signals the origin of a higher-period cycle, and analysis of slopes reveals it to have period 4. Note that some of the long waves do *not* change slope at time step 450,000, which shows that the period-3 strategy remains alive on some of the islands in the archipelago. During the middle of the run the analysis again is ambiguous, and there is some evidence of another adaptive event at around time step 500,000. But close inspection shows a fairly clear adaptive event at around time step 650,000 when the start of a new significant wave coincides with a change in slope of three of the long waves that exist then. Calculation of slopes reveals this to be a strategy with period 5. Another series of kinks in waves around time step 850,000 that coincide with the start of another new wave and the termination of some old waves signals the origin of a longer-period strategy. Slope calculation shows it to have period 6.

In this way, the allele token activity reveals the timing and nature of the key adaptive events in this run much more clearly than the population level dynamics. This plotline is disclosed still more sharply by the allele type activity shown in the bottom of Figure 2. The 2-cycle epoch is shown by the initial period with two allele type activity waves (note that they initially overlap). The 3-cycle epoch clearly starts with the pair of waves that originate at time step 200,000. Their coexistence with one of the waves from the 2-cycle shows that the 3-cycle includes one of the alleles from the 2-cycle. The third epoch clearly starts with the start of the next pair of (initially overlapping) big waves at time step 450,000. This strategy is a 4-cycle because it consists of four waves, two of which are from the 3-cycle. (The other wave with slope falling to near zero shows that the 3-cycle persists in a small subpopulation somewhere in the archipelago.) The fourth epoch clearly starts with the two new waves originating at about time step 650,000. Counting waves shows these to be part of a 5-cycle which uses three of the genes from the 4-cycle. (Note again that the wave whose slope falls to near zero shows that the 4-cycle persists for a while in a subpopulation.) The fifth epoch clearly starts with the origin of the two new waves at time step 850,000. Counting waves shows this strategy to be a 6-cycle which incorporates four of the genes from the 5-cycle. (Other period-lengthening adaptations happen very near the end of the run.)

As in the Block run analyzed above, the Archipelago run shows the relative merits of token and type activity data. The few big waves of type activity clearly and sharply identify the main branches of the adaptive story. Genetic drift of neutral variants is filtered out and only phenotypically novel adaptations appear. The wealth of little and large token activity waves discloses a variety of details of the adaptive subplots and sideshows that criss-cross the main branches of the adaptive story. (For example, the distinctive “bushy” character of the second main token activity wave in the middle of Figure 2 is the signature of a certain distinctive kind of adaptation the analysis of which is left as an exercise for the reader (see (Bedau *et al.* 1999) for a hint). Also, a blow-up of the very low activity waves would disclose some caused when agents are knocked out of their cycle by another agent occupying the same grid location.) Taken together, token and type activity provide a wealth of information about the adaptive dynamics of the components under analysis.

Evolutionary activity in a Ridge

In the Ridge environments we set a positive sensory threshold mutation, $p_{\text{th}} = 0.2$, so the evolution of sensorimotor strategies often coincides with the evolution of sensory thresholds. This creates a new and different kind of evolutionary dynamics, and it affords the opportunity to examine evolutionary activity of two different kinds of genes: sensorimotor genes and the sensory threshold genes.

The top of Figure 3 shows the population level in a typical Ridge run, along with the mean value of the threshold genes in the population. We can clearly see at least three epochs. The first consists of the first 10,000 time steps, while the population rises to about 200 and the mean sensory threshold is low. The second epoch covers the period from time step 10,000 to time step 20,000, in which the population rises to about 240 and the mean sensory threshold rises to about 190. The third covers the rest of the run, in which the population falls to about 210 and the mean sensory threshold rises to about 225. There seems to be a little blip in the population level at about time step 35,000 but it is not clear if this is significant.

Plotting the sensorimotor allele token activity clarifies what is happening (middle of Figure 3). The initial wave has slope very close to one, which shows that the agents quickly evolve the sensorimotor strategy of running along the length of the ridge. Since the mean sensory threshold is pretty low, those agents that are almost anywhere on the ridge seem to be in the middle of a wide band of resources. But unbeknownst to them, only those sites near the spine of the ridge actually contain enough resources to sustain existence for long. So the agents that happen to be running along the ridge near its spine survive and proliferate. At around time step 10,000 there is a very slight kink in the slope of the big sensorimotor token wave, and a second very low slope wave starts. This along with the change in mean sensory threshold shown in the top of Figure 3 shows that the agents have evolved two adaptations: First, they can now detect edges near the spine of the ridge (specifically, two steps away from the spine); second, they have evolved a strategy for jumping back onto the spine when they detect the spine's edges. This increases their chances of survival and hence the population rises. New kinks in these two sensorimotor token waves at time step 20,000 and the change in mean sensory threshold at the same time reveals similar adaptation: the agents can now detect the edges of the spine itself. Since their perception of the spine is now narrower, they detect spine edges more frequently than before, and so they spend more time jumping back onto the spine from the edge. Note also the pair of waves that arise at about time step 35,000, and the subsequent termination of the big wave. This confirms our earlier suspicion that another sensorimotor adaptation arises at this point and reveals it to be a new kind of behavior for running along the length of the ridge.

Since the Ridge runs involved the evolution of sensory thresholds, we collected activity for sensory threshold alleles. These alleles can be set at integer values between zero and the maximum resource level (here, 250). But since the ridge is only fifteen sites wide and it rises steadily from its edges to its maximum height in the middle site, there are only eight different resource levels. So we collected the sensory threshold alleles into equivalence classes according to which of the eight resource levels in the Ridge environment they could discriminate. These classes of threshold types collected phenotypically equivalent threshold alleles. The bottom of Figure 3 shows the activity of these phenotypically equivalent threshold allele tokens. The three epochs of the threshold evolution are very clearly indicated by the three rising threshold activity waves. (The adaptive event at time step 35,000 does not appear on this plot because it is a sensorimotor, rather than a sensory threshold, adaptation.)

As with sensorimotor type and token activity, sensory threshold activity provides a clear picture of key threshold adaptations, at least when the activity is collected into phenotypic equivalence classes. (Threshold activity that is not binned into phenotypic equivalence classes is much less clear. Data not shown.)

Conclusions

We have added new weight to the evidence that shows how evolutionary activity highlights significant adaptive evolutionary phenomena such as competitive exclusion, cooperation, and neutral genetic drift. We also show the first measurements of evolutionary activity of genes controlling the evolution of evolvability. But perhaps most significantly, we have shown how to extend the evolutionary activity method to a new and especially revealing component of evolving systems: phenotypically equivalent types of alleles. The simultaneous analysis of evolutionary activity of different components of an evolving systems enables us to tease apart the various adaptive phenomena interwoven through the evolutionary process. Our methods easily generalize to many evolving systems with other kinds of alleles, as well as to non-allelic components of such systems.

The one-to-one mapping between genes and behaviors in Packard's model makes it especially easy to identify phenotypically equivalent alleles. It will be more difficult to identify phenotypically equivalent alleles in systems with more complex genotype to phenotype mappings. Some genotype-phenotype mapping are context sensitive because many genes together affect single characters (e.g., epistasis or polygenes) or many characters are affected by single genes (e.g., developmental programs). For instance, a gene coding for a given connection strength in a neural network might influence behavior to different ways depending on how activation is flowing through all the other nodes. It does not take much context sensitivity before phenotypically equivalent alleles are impossible to identify in practice. So genetic context sensitivity can limit the application of evolutionary activity analysis. But it can still be applied. For example, one could measure the activity of individual genes coding for neutral network controllers as we do here for sensorimotor genes: by noting their age. And one could measure the activity of phenotypic behaviors in such systems by attaching activity counters directly to those behaviors and incrementing the counters every time the behavior was observed.

There are many ways to measure evolutionary activity, many components of evolving systems that can be observed, and many ways in which the evolutionary activity of those components can be incremented. We have shown here that measuring evolutionary activity of allele tokens, allele types, and phenotypically equivalent classes of alleles provides a useful visualization of the adaptive evolutionary dynamics in Packard's Bugs model. This implementation of the evolutionary activity method was made especially easy because of the simplicity of Packard's model. Indeed, this simplicity is precisely why we used Packard's model, for it makes illustrating evolutionary activity so straightforward.

Our interpretation of the evolutionary activity graphs depended heavily on our understanding of the details of the niches afforded in Heap, Archipelago, and Ridge. We used our prior knowledge of Packard's model to propose and evaluate hypotheses about the meaning of specific activity wave phenomena. But this prior knowledge is not necessary for understanding the most salient features of activity plots. In particular, even if we knew nothing about Packard's model, glancing at the allele token and type activity plots indicate the timing of the main adaptive innovations: these correspond to the start of major new waves. The diagrams indicate something about the nature the innovations and how they are related.

For example, the major new activity waves in the Archipelago data shown in Figure 2 coexist with the previous activity waves. By contrast, the first major new wave in the Heap data shown in Figure 1 coincides with the extinction of the previous waves. This suggests that the Archipelago innovations involve the cooperation (or at least coexistence) of some aspects (alleles) of previous innovations while the first innovation in the Heap drives the initial adaptation to extinction via competitive exclusion. The point is that these and other lessons can be gleaned from activity plots without knowing any specific information about the nature of adaptation in the system.

Learning how to interpret activity diagrams is a little like learning how to interpret the tracks in cloud chambers. Experience helps a lot. The more you know about the system, the better you can understand the wave phenomena that you observe. At the same time, the more you analyze activity waves and compare wave phenomena created by incrementing activity of different kinds of components in different ways, the more you learn about the underlying adaptive dynamics in those systems.

Acknowledgments

Thanks to Norman Packard for a decade of collaboration on visualizing and quantifying evolutionary activity. For collaboration on visualizing evolutionary activity of various components of various systems, thanks to Shareen Joshi, Ben Lillie, Titus Brown, Terry Jones, Simon Frasier, Andre Skusa, Tom Ray. And thanks to the anonymous reviewers for helpful comments.

References

Bedau, M. A. 1995. Three illustrations of artificial life's working hypothesis. In W. Banzhaf and F. Eeckman, eds., *Evolution and Biocomputation—Computational Models of Evolution*, pp. 53–68. Berlin: Springer.

Bedau, M. A. 1996. The nature of life. In M. Boden, ed., *The Philosophy of Artificial Life*, pp. 332–357. New York: Oxford University Press.

Bedau, M. A. 2001. Dynamics of the environment for adaptation in static resource models. In J. Keleman and P. Sosik, eds., *Advances in Artificial Life, Proc. 6th European Conference on Artificial Life*, pp. 76–85. Heidelberg: Springer-Verlag.

Bedau, M. A. 2002. Evolution of evolvability via adaptation of mutation rates. *BioSystems.* To appear.

Bedau, M. A.; and Brown, C. T. 1999. Visualizing evolutionary activity of genotypes. *Artificial Life* 5: 17-35.

Bedau, M. A; Joshi, S.; and Lillie, B. 1999. Visualizing waves of evolutionary activity of alleles. In A. Wu, ed., *Proceedings of 1999 Genetic and Evolutionary Computation Conference Workshop Program*, pp. 96-98.

Bedau, M. A.; and Packard, N. H. 1992. Measurement of evolutionary activity, teleology, and life. In *Artificial Life II*, C. Langton, C. Taylor, D. Farmer, S. Rasmussen, eds., pp. 431-461. Redwood City, CA: Addison-Wesley.

Bedau, M. A.; Snyder, E.; Brown, C. T.; and Packard, N. H. 1997. A comparison of evolutionary activity in artificial evolving systems and in the biosphere. In *Proceedings of the Fourth European Conference on Artificial Life*, P. Husbands and I. Harvey, eds., pp. 125-134. Cambridge, MA: MIT Press.

Bedau, M. A.; Snyder, E.; and Packard, N. H. 1998. A classification of long-term evolutionary dynamics. In *Artificial Life VI*, C. Adami, R. Belew, H. Kitano, and C. Taylor, eds., pp. 228-237. Cambridge, MA: MIT Press.

Burian, R. M. 1992. Adaptation: historical perspectives. In E. F. Keller and E. A. Lloyd, eds., *Keywords in Evolutionary Biology* pp. 7–12. Cambridge, MA: Harvard University Press.

Dawkins, R. 1979. *The Selfish Gene,* new edition in 1989. New York: Oxford University Press.

Dawkins, R. 1987. *The Blind Watchmaker: Why the evidence of evolution reveals a universe without design.* New York: Norton.

Fletcher, J. A., Zwick, M., Bedau, M. A. 1996. Dependence of adaptability on environment structure in a simple evolutionary model. *Adaptive Behavior* **4**, 283–315.

Gould, S. J., Lewontin, R. C. 1979. The spandrals of San Marco and the Panglossian paradigm: a critique of the adaptationist programme. *Proceedings of the Royal Society B* 205: 581–598.

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.

Holland, J. H. 1995. *Hidden Order: How Adaptation Builds Complexity*. Reading, MA: Addison-Wesley/Helix Books.

Maynard Smith, J. 1975. *The Theory of Evolution*, 3rd edition. New York: Penguin.

Mayr, E. 1988. *Towards a New Philosophy of Biology*. Cambridge, MA: Harvard University Press.

Packard, N. H. 1989. Intrinsic adaptation in a simple model for evolution. In C. G. Langton, ed., *Artificial Life*, Sante Fe Institute Studies in the Sciences of Complexity (pp. 141–155). Redwood City, CA: Addison-Wesley.

Rechtsteiner, A.; and Bedau, M. A. 1999. A generic model for quantitative comparison of genotypic evolutionary activity. In *Advances in Artificial Life*, D. Floreano, J.-D. Nicoud, F. Mondada, eds., pp. 109-118. Heidelberg: Springer-Verlag.

Skusa, A.; and Bedau, M. A. 2002. Towards a comparison of evolutionary creativity in biological and cultural systems. In *Artificial Life VIII*, R. Standish, H. Abbass, and M. A. Bedau, eds. Cambridge, MA: MIT Press.

West-Eberhard, M. J. 1992. Adaptation: current uses. In E. F. Keller and E. A. Lloyd, eds., *Keywords in Evolutionary Biology*, pp. 13–18. Cambridge, MA: Harvard University Press.

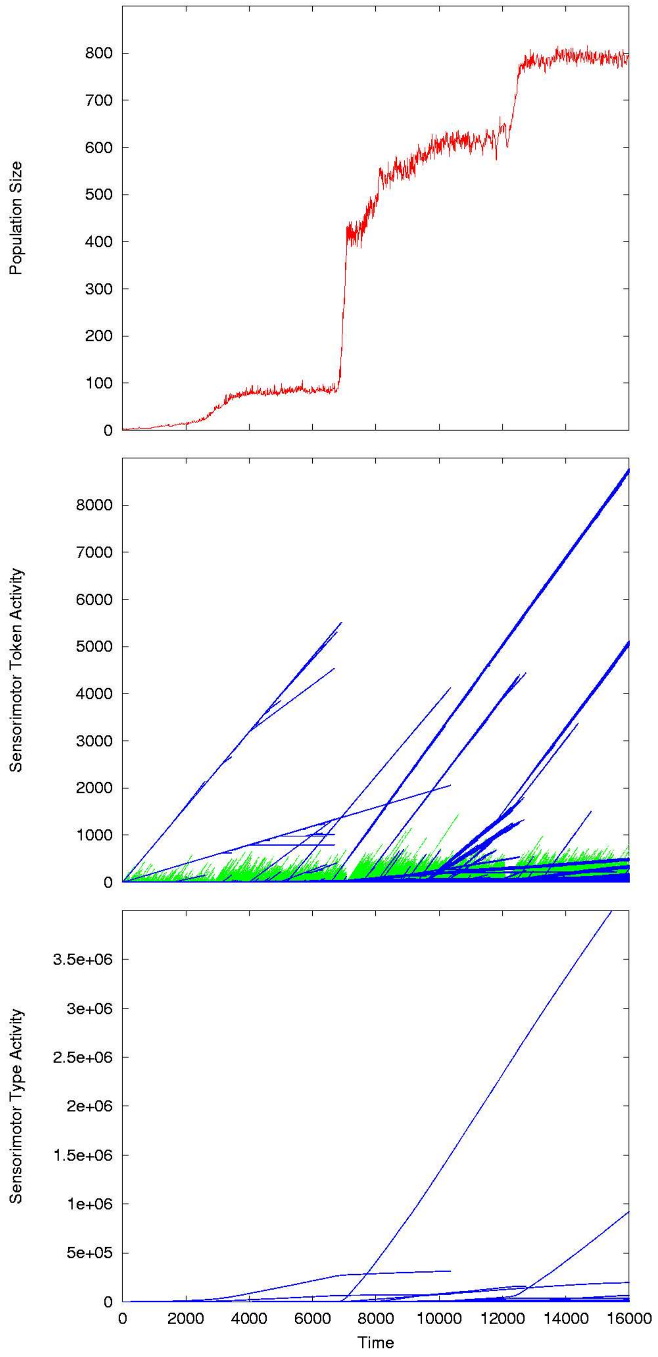

Figure 1: Top: Population level in a typical Heap. Middle: The activity of sensorimotor allele tokens (blue) and their analogues in a neutral model (green). Bottom: The activity of phenotypically equivalent types of sensorimotor alleles (the activity of their neutral analogues is too low to be visible).

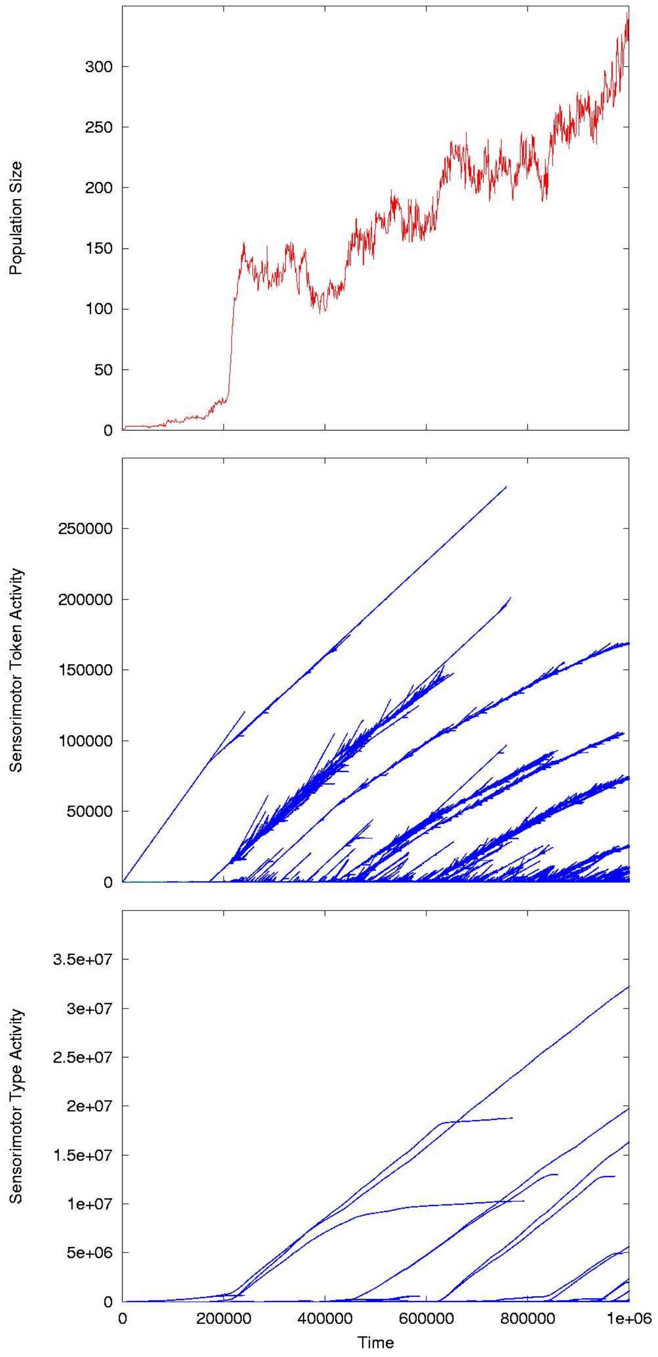

Figure 2: Top: Population level in a typical Archipelago. Middle: The activity of sensorimotor allele tokens (blue) and their analogues in a neutral model (green). Bottom: The activity of phenotypically equivalent types of sensorimotor alleles (neutral analogue activity is too low to be visible).

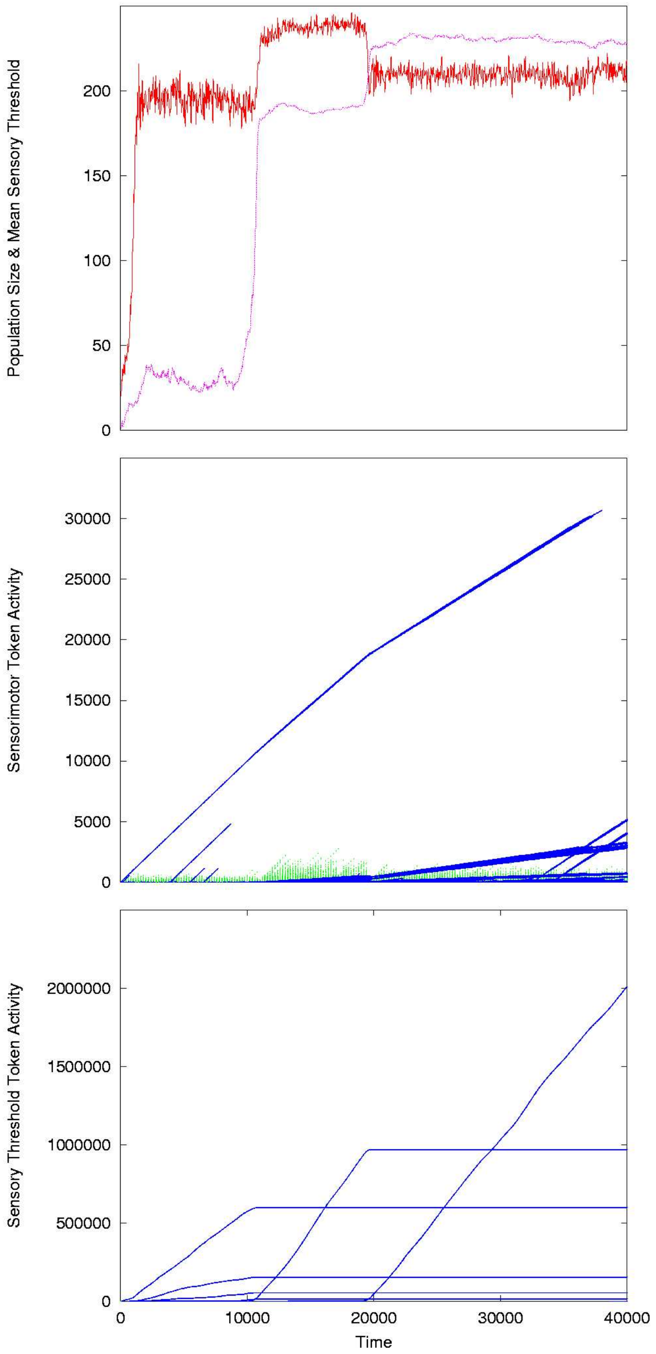

Figure 3: Top: Population level (red) and mean value of the sensory threshold alleles (magenta) in a typical Ridge. Middle: The activity of *sensorimotor* allele tokens (blue) and their analogues in a neutral model (green). Bottom: The activity of phenotypically equivalent types of *sensory threshold* alleles (neutral analogue activity is too low to be visible).