

The Flexible Balance of Evolutionary Novelty and Memory in the Face of Environmental Catastrophes

Andrew Buchanan, Mark Triant and Mark A. Bedau
Reed College, 3203 SE Woodstock Blvd., Portland OR 97202 USA
bedau@reed.edu

Abstract

We study the effects of environmental catastrophes on the evolution of a population of sensory-motor agents with individually evolving mutation rates, and compare these effects in a variety of control systems. The evolution of mutation rates must balance (i) the need for evolutionary “novelty,” which pushes mutation rates up, and (ii) the need for evolutionary “memory,” which pushes mutation rates down. We observe that an environmental catastrophe initially shifts the balance toward evolutionary novelty and causes mutation rates to evolve upwards. Then, as the population adapts to the new environment, the balance shifts back toward evolutionary memory and the mutation rate falls. These observations support the hypothesis that second-order evolution of the mutation maintains a flexibly shifting balance between evolutionary novelty and memory.

Introduction

The evolution of life on Earth has repeatedly been forced to adapt to environmental catastrophes (Raup, 1991), and evolution through such catastrophes has been suggested as a mechanism behind adaptive radiation (Raup, 1986). As artificial evolving systems increasingly interact with unpredictable environments, their robustness in the face of environmental catastrophes will also become increasingly critical. Some (e.g., (Stanley, 1973; Stanley, 1990; Sepkoski, 1997)) account for diversity trends in the fossil record by appealing to the tendency of ecologically unspecialized taxa to better survive drastic environmental change. This may likewise figure into an explanation of a taxon’s resilience in the face of environmental catastrophe. In any case, understanding the nature and source of evolutionary robustness in the face of environmental catastrophes is important for understanding both natural and artificial evolution.

The ability of a system to adapt to environmental catastrophes is related to its evolvability, i.e., the capacity for evolution to create new adaptations. Evolvability has received considerable attention recently in both the biological and evolutionary computation communities. There is general agreement that evolvability is crucial for understanding the origin of complex adaptations (Conrad, 1982; Wagner

and Altenberg, 1996) as well as the process of open-ended evolution (Taylor, 1999; McMullin, 2000), a central open problem in artificial life (Bedau et al., 2000). This paper treats the issue of how evolvability evolves and specifically focuses on the evolution of evolvability in a finite population adapting to a dynamic endogenous fitness function in a spatial environment.

A system’s evolvability depends on its ability to produce adaptive phenotypic variation, and this hinges on both the extent to which the system’s phenotype space contains adaptive variation and the ability of evolutionary search to locate it while avoiding maladaptive traps. Two main factors control the effectiveness of the evolutionary search process: the way in which genetic operators traverse genotype space, and the way that genotypes are phenotypically expressed (the genotype-phenotype mapping). For evolutionary search to explore a suitable variety of viable evolutionary pathways, genetic operators must generate sufficient amounts of the right kind of genetic *novelty*. At the same time, since evolutionary adaptations are built through successive improvements, genetic *memory* is required for the evolutionary process to retain incremental improvements.

Genetic operators like mutation rate simultaneously affect the demands for both evolutionary novelty and memory, so evolvability requires genetic operators to balance these competing demands successfully. (In evolutionary computation this principle is known as the tension between “exploration” and “exploitation” (Holland, 1975).) Furthermore, the appropriate balance between evolutionary novelty and memory can continually change as evolution progresses, so evolvability requires genetic mechanisms to adjust the balance flexibly. Thus, understanding evolvability involves understanding how the balance between evolutionary novelty and memory shifts during the course of evolution, and what general mechanisms allow this balance to flexibly shift. One way to address these questions is to let second-order evolution control the genetic mechanisms that structure first-order evolution, enabling us to study the evolution of evolvability. There are a variety of such genetic mechanisms, but perhaps the simplest is the mutation rate. This paper addresses the

Copyrighted Material

evolution of evolvability in this context.

Optimal and evolving mutation rates are discussed in both the biological and evolutionary computation literatures (Kimura, 1960; Eiben et al., 1999); further references are cited elsewhere (Bedau and Packard, 2003). But it remains unclear what, if anything, such results might reveal about the evolution of mutation rates in a finite spatial population with many loci and many alleles per locus, with heavy neutral evolution, and when the context of evolution continually changes and thus the fitness function unpredictably varies—the context of the present study. The discussion of evolving mutation rates in the evolutionary computation literature mainly focuses on the issue of automated control of evolution, in the context of efforts to minimize the time required to solve function optimization problems. But this work typically presumes that evolution is driven by a fixed and externally-specified fitness function, whereas the theoretical issues that interest us concern evolution with implicitly-specified fitness functions that continually change in the course of evolution.

An earlier treatment of these issues provided preliminary evidence that second-order evolution of mutation rates allows a flexible response to exogenously shifting demands for evolutionary novelty and memory (Bedau and Packard, 2003). The present paper systematically explores this earlier result, in order to determine how robust the process is and to identify the mechanisms behind it. In particular, we study the robustness of this process by observing hundreds of catastrophes caused by transitions to a variety of different kinds of environments. In addition, we isolate the mechanisms behind what we observe by comparing these results with a variety of control systems.

Our experiments with environmental catastrophes also have the potential to address a further issue: whether and how evolving through a series of environmental catastrophes increases a population's ability to adapt to novel environments in the future. Earlier work has shown that autonomous agents with the ability to learn during their lifetime adapt to changing environments (Menczer and Belew, 1996), choose from a variety of environmental niches (Walker, 1999), and evolve increased genomic and neural complexity when subjected to noisier environments (Seth, 1998). Environmental catastrophes have also been shown to affect the diversity dynamics of evolving communicative systems (Green et al., 2000). Because our system includes second-order evolution, we can look at similar issues in a simpler and more fundamental context, and thus perhaps explain what happens in more complex settings.

The Model

Our evolutionary system is composed of many agents that could be called organisms, on analogy with biological systems. It has been used in various studies of the evolution of sensory-motor functionality; see the references in earlier

work (Bedau and Packard, 2003). The system consists of many agents that exist together in an environment, in this case a toroidal lattice. The lattice has a binary field defined on it, $\mathcal{E}(x)$, which may be interpreted as a resource field. The pattern of the resource field is static and resources are not depleted by agents. Thus, agents are constantly gathering resources and expending them through metabolism and reproduction, but the resource pattern in the environment remains fixed. The exception to this are the periodic environmental catastrophes, as described below.

The resource patterns in our system consist of three qualitatively distinct types. The first, B , maps a square onto the lattice. The absolute amount of resources available to the agents across different simulations in B is fixed, but the location is chosen at random. The second, C , maps resources onto the environment in numerous pairs of equivalent, mutually bisecting, perpendicular line segments. The placement of each pair of line segments in C is chosen at random independently of other pairs leading to partial overlap. The absolute amount of resources available is thus an inverse function of the degree of overlap in that particular instantiation of C . The third resource pattern, R , maps individual resource sites onto the environment at random. As with C , the sites are chosen independently of one another leading to a degree of overlap and variable absolute amounts of food. However, in all but improbable cases of extreme overlap, the total resources, $\Sigma \mathcal{E}$, are such that $\Sigma \mathcal{E}_R > \Sigma \mathcal{E}_C > \Sigma \mathcal{E}_B$.

In analogy with biological systems, the dynamics of the population as a whole are comprised of all the birth-life-death cycles of the agents. Births occur when agents accumulate enough resources to reproduce (see below), deaths occur when agents run out of resources, and the lives of agents consist of their interactions with the environment. We label each agent with the index i , let I^t be the set of agents existing at t . Time is discrete. One unit of time is marked by each agent acting based on their local environment.

During its lifetime, each agent extracts information from the environment by sensing and reacting to the local resource field. There is a discrete set of different possible sensory states, $s \in S = \{s_1, \dots, s_{N_S}\}$. The agents exchange no information with each other directly (although this would be an easy generalization). Each agent (labeled with the index i) has certain information associated with it: (i) a current location, x_i^t , (ii) a current sensory state, s_i^t , (iii) a current reservoir of resources, E_i^t , and (iv) a sensory-motor map, ϕ_i , that yields a behavior given sensory information as input. In this model, an agent's behavior is a vector, \vec{b} , denoting the agent's movement in the environment, $x_i^t \rightarrow x_i^{t+1} = x_i^t + \vec{b}_i^t$ where $\vec{b}_i^t = \phi_i(s_i^t)$. Every \vec{b}_i^t is a member of a discrete set of different possible behaviors, $\vec{b}_i^t \in B = \{\vec{b}_1, \dots, \vec{b}_{N_B}\}$. (Another easy generalization would be to include other kinds of possible behaviors.) Apart from reproduction, the agent's supply of resources may be changed in one of two manners: an augmentation from extracting resources at the agent's

new location, or a constant-sized universal reduction, $E_i^t \rightarrow E_i^{t+1} = E_i^t + \alpha(\mathcal{E}(x_i^{t+1})) - \beta$, where $\alpha(z)$ is a resource extraction function and β is a constant metabolic cost of surviving.

The sensory-motor map, φ_i , operating on a sensory input, s_i^t , has a particularly simple form because the sensory input is discrete. Since $s_i^t \in \{s_1, \dots, s_{N_S}\}$, we may identify the function φ_i with its graph, a set of N_S behavior values, $\{\varphi_{is}\}$. Pursuing the biological analogy, we will consider the sensory-motor strategy elements $\{\varphi_{is}\}$ as i 's *genome*, with N_S loci, and each particular element $\bar{b} = \varphi_{is}$ as the trait (or allele) at the s^h locus of i 's genome.

Reproduction occurs when an agent's resource supply exceeds a threshold, E_r . The parent splits its resources with its child, and the child inherits its parent's strategy elements $\{\varphi_{is}\}$, apart from changes introduced by mutation. The genome of each agent i contains a special gene, μ_i , which controls the rate at which i 's strategy elements mutate when i reproduces, i.e., the probability that a strategy element of i 's children is chosen (with equal probability) from the set of possible behaviors, B . We also introduce a *meta-mutation* rate parameter, μ_μ —the probability that i 's children's mutation-rate gene is chosen (with equal probability) from the interval $[\mu_i - \epsilon, \mu_i + \epsilon]$. The value of μ_μ is fixed during the course of a given simulation, and the value of μ_i is fixed during the course of i 's lifetime (as is φ_i).

This model provides a simple setting for empirical study of the evolution of evolvability. Agents' immediate environments produce sensory states that then trigger actions by means of the agents' sensory-motor maps. Since the agents' survival and reproduction depends directly on their ability to continually find resources in their environment, the implicit fitness functions in this model are constantly buffeted by the contingencies of natural selection and, thus, unpredictably change. This first-order evolution is structured by the sensory-motor maps actually compared and tested by natural selection. One especially simple mechanism that regulates the variety of maps available for evolutionary exploration is the mutation rate; the higher the mutation rates, the greater the variety. Thus, by allowing mutation-rate genes to evolve, we can study second-order evolution of evolvability.

Methods

We collected data from the model in an experimental situation and three kinds of controls. In the experimental condition, the agents' mutation rates were allowed to evolve and the environment cycled through a sequence of qualitatively different patterns, B , C and R , at a regular interval. The effect of these environmental catastrophes on the agents in this phase and all subsequent phases was measured in terms of population and average mutation rate across the entire population. The Same-Environment (SE) control was identical to the experimental situation except that the catastrophes




Figure 1: Top-down pictures of a part of the three environments B , C , and R .

rophes did not involve a qualitative change in environment; rather, whichever pattern was in place was redistributed so that local environments were changed, but the overall environment was qualitatively identical. In the Population-Only (PO) control, the environmental catastrophe was replaced by an artificial population catastrophe. Ninety percent of the population was selected at random and killed, mimicking the effect that environmental catastrophes had on the population. We conducted 80 trials in the experimental situation with counterbalanced sequences, and these data were averaged for analysis. In addition, we generated a total of 48 runs in the SE control and 29 in the PO control.

The size of the torroidal lattice was held fixed at 141×141 sites and all trials had 120,000 iterations with catastrophes every 10,000 iterations. When the mutation rate was allowed to evolve, the initial mutation rates μ were set to 0.3 (30% of the loci mutated each reproduction event). In order to prompt quick evolution of mutation rates, μ_μ was set to 0.5 ϵ to 0.1 (e.g., 50% chance of the mutation rate mutating up or down by as much as 0.1). The ratio of metabolic cost to the resource intake rate and to the reproduction threshold was 10:11:50 in all conditions. Initial populations always consisted of 1,000 agents. Reproduction was always asexual and always produced exactly one new agent. Population and mutation rate data were sampled every 30 model updates. Resource pattern B was 45×45 sites for a total of 2,025 resource sites. Pattern C consisted of 70 pairs of lines with 101 total resource sites for a total of 7,070 resource sites less any overlapping sites. 12,500 resource sites were distributed in the R pattern (less some overlap). See Figure 1.

Results

Figure 2 shows the population level response to the following regular sequence of environmental catastrophes: R , B , C , R , B , etc. The results from 80 runs were averaged, and the data before the first catastrophe was discarded. The catastrophes occur in such quick succession that the population level never reaches equilibrium in an environment, thereby enabling us to study the population's immediate response to catastrophe. The catastrophes have a dramatic impact on population level, truncating the population to less than 10%

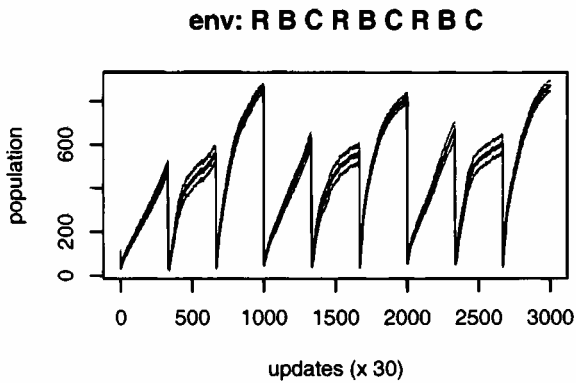


Figure 2: Dynamics of the population level (dark) plus and minus the standard error (light), averaged across 80 runs. Every 10,000 model updates the environment suffers an environmental catastrophe.

of its pre-catastrophe size, but the population recovers as it adapts to its new environment. By comparison, the population bottleneck in the SE control is never even as small as 50% of the population size (data not shown). Evidently, when the population experiences a new instance of the same kind of environment in the SE control, most agents can successfully use the same genes in the new environment. This implies that after an environmental catastrophe in the experimental condition, the population cannot rely on genes it already has but must readapt. Even if the population experienced that same kind of environment earlier in its history, adaptive behavioral patterns are no longer in genetic memory; it remembers little if anything about how to behave in that environment since it has most recently been in two very different kinds of environments.

When the population adapts to a given environment and never experiences environmental catastrophes, the mutation rate evolves down to an equilibrium value (data not shown). Figure 3 shows how the mutation rate responds to a series of environmental catastrophes. After each kind of catastrophe, the mean mutation rate rises significantly, and the population level starts to recover from the catastrophe. While the population level continues to rise, after some time the mean mutation rate then starts to fall and is still falling when the next catastrophe strikes. This starts the same mean mutation rate dynamics again.

Figure 4 shows the mean mutation rate recovery from the three different kinds of environmental catastrophes. Note that each kind of environment produces a characteristic mutation rate dynamic, and that the *R* dynamic differs significantly from the characteristic *B* and *C* dynamics, which are similar but distinguishable. Comparing the mutation rate dynamic in the PO control (light) shows that the mutation rate response to environmental catastrophes is

env: R B C R B C R B C




Figure 3: Dynamics of the mean mutation rate (dark) plus and minus the standard error (light), averaged across the same runs as in Figure 2. Every 10,000 model updates the environment suffers an environmental catastrophe.

result of population bottleneck the catastrophes cause. Furthermore, the mutation rate dynamic in the SE control (data not shown) is like that in the PO control. Thus, the mutation rate response to environmental catastrophe is not caused by any type of environmental disruption. Rather, it is caused by experiencing a new kind of environment, even if that same kind of environment had been experienced at some point in the past.

Figure 5 shows how the population level responds to different kinds of environmental catastrophes. Note that different kinds of environment produce characteristically different kinds of population dynamics. Also note that in the *R* environment the population is adapting more quickly to later *R* catastrophes (the thick line has lower slope than the thin line). We see a different but related effect in the *B* environment (higher maximum population levels after experiencing more catastrophes).

Discussion

The mutation rate dynamics seen in Figure 3 fit the hypothesis that second-order evolution can fluidly balance evolutionary memory and novelty. The low population level immediately after an environmental catastrophe indicates that the population is not well adapted to its new environment, so the balance shifts toward evolutionary novelty. The observed rise in mutation rate achieves this. As the population becomes adapted to its new environment, the population level climbs, and the balance shifts back toward evolutionary memory. The observed fall in mutation rate brings this about.

What is the mechanism by which the mutation rate adapts? Having a higher or lower mutation rate does not affect an individual's chances of surviving or reproducing; its effects would be felt only by an individual's offspring. If an




Figure 4: The response of mean mutation rate (dark) in three different kinds of environmental catastrophes, averaging data from 80 runs. The three different curves in a given environment correspond to the first (thick dark line), second (medium dark line), and third (thin dark line) time that kind of environment was encountered in the series of catastrophic environmental changes. This is compared with the mean mutation rate response in the same environment to 6 catastrophes in the PO control (light), averaging 11 runs in the *C* environment, 11 in the *B* environment, and 7 in the *R* environment.

individual has the appropriate mutation rate in a given context, it is more likely that the individual's offspring will survive and flourish. So, if an individual's fitness is measured also by the survivorship and fecundity of its offspring, then having a contextually appropriate mutation rate increases an individual's fitness. Just as "inclusive fitness" extends an individual's fitness over space through kin (Hamilton, 1964), an individual's fitness also can be extended over time through a lineage. Thus, the mechanism behind the adaption of mutation rates can be viewed as a kind of group selection.

The different characteristic mutation rate dynamics elicited by different kinds of environmental catastrophes shown in Figure 4 further indicate the fluidity of this balancing process. When and how the balance shifts depends on the nature of the environment and the extent to which the population has already adapted to it.

We would explain the characteristic mutation rate dynamics in the different environments by reference to their different local niche structure. The balance remains shifted towards evolutionary novelty (higher mutation rates) longer in *R* than in *B* and *C* because *R* contains a greater number of different local niches that call for different local behavioral strategies, and, thus, adaptation takes more time. By contrast, it is relatively easy to find the local strategies needed in *B* and *C* so the balance quickly tips back toward evolutionary memory. In addition, we hypothesize that the faster shift back toward memory in *B* than in *C* is a result of resources in *B* occurring in one contiguous space, while the resources in *C* are distributed across space, so it takes the agents longer to physically explore all the niches.

The local niche structure in the different environments also explains the characteristic population dynamics seen in Figure 5. The population in *R* continues to adapt to new local niches throughout its time in the *R* environment, so the

population is continually rising. In both the *B* and *C* environments the population eventually starts to saturate the available niches, and the population starts to level off.

The fact that the population adapts to new *R* environments better after it has experienced more *R* catastrophes is intriguing (left panel of Figure 5). This provides evidence for the hypothesis that experiencing environmental catastrophes increases a population's ability to adapt to new environments, a hypothesis suggested by the fossil record (Raup, 1986). The way in which the final population level in *R* and *B* environments increases with the number of catastrophes experienced tends to further corroborate this hypothesis.

Topics for future work include getting experimental evidence about the robustness of the flexible balance of evolutionary memory and novelty by studying a wider range of models subject to a wider range of environmental catastrophes, and constructing a quantitative theory of the flexible balance of evolutionary memory and novelty that enables the mutation dynamics to be predicted, and that explains the connections of this phenomenon with other theories about mutation rate such as that concerning quasispecies (Eigen and Schuster, 2001; Eigen et al., 1988).

Acknowledgements

Thanks to Mike Raven and Albyn Jones for help with code, to Jeffrey Fletcher and John Huss for helpful discussion, and to the ALife IX referees for helpful comments.

References

- Bedau, M. A., McCaskill, J. S., Packard, N. H., Rasmussen, S., Adami, C., Green, D. G., Ikegami, T., Kaneko, K., and Ray, T. S. (2000). Open problems in artificial life. *Artificial Life*, 6:363–376.




Figure 5: The response of the population level to three different kinds of environmental catastrophes, averaging data from 80 runs. The three curves in a given environment show the response to the first (thick line), second (medium line), and third (thin line) environmental catastrophe of a given kind.

- Bedau, M. A. and Packard, N. H. (2003). Evolution of evolvability via adaptation of mutation rates. *Biosystems*, 69:143–162.
- Conrad, M. (1982). Natural selection and the evolution of neutralism. *Biosystems*, 15:83–85.
- Eiben, A. E., Hinterding, R., and Michalewicz, Z. (1999). Parameter control in evolutionary algorithms. *IEEE Transactions on Evolutionary Computation*, 3:124–141.
- Eigen, M., McCaskill, J., and Schuster, P. (1988). Molecular quasiespecies. *Journal of Physical Chemistry*, 92:6881–6891.
- Eigen, M. and Schuster, P. (2001). *The hypercycle: a principle of natural self-organization*. Springer-Verlag.
- Green, D. G., Newth, D., and Kirley, M. G. (2000). Connectivity and catastrophe: Towards a general theory of evolution. In Bedau, M. A., McCaskill, J. S., Packard, N. H., and Rasmussen, S., editors, *Artificial Life VII*, page 153. Cambridge, MA: MIT Press.
- Hamilton, W. D. (1964). The genetical evolution of social behavior. *Journal of Theoretical Biology*, 7:1–52.
- Holland, J. H. (1975). *Adaptation in Natural and Artificial Systems: An Introductory Analysis with Applications to Biology, Control, and Artificial Intelligence*. University of Michigan Press, Ann Arbor. (Second Edition: MIT Press, 1992.).
- Kimura, M. (1960). Optimum mutation rate and degree of dominance as determined by the principle of minimum genetic load. *Journal of Genetics*, 57:21–34.
- McMullin, B. (2000). John von neumann and the evolutionary growth of complexity: looking backwards, looking forwards. *Artificial Life*, 6:347–361.
- Menczer, F. and Belew, R. K. (1996). From complex environments to complex behaviors. *Adaptive Behavior*, 4:317–363.
- Raup, D. M. (1986). Biological extinction in earth history. *Science*, 231:1528–1533.
- Raup, D. M. (1991). *Extinction: bad genes or bad luck?* Norton.
- Sepkoski, J. J., J. (1997). A kinetic model of phanerozoic taxonomic diversity ii: early phanerozoic families and multiple equilibria. *Paleobiology*, 5:222–251.
- Seth, A. K. (1998). The evolution of complexity and the value of variability. In Adami, C., Belew, R. K., Kitano, H., and Taylor, C. E., editors, *Artificial Life VI*, page 209. Cambridge, MA: MIT Press.
- Stanley, S. M. (1973). An explanation of cope's rule. *Evolution*, 27:1–26.
- Stanley, S. M. (1990). Delayed recovery and the spacing of major extinctions. *Paleobiology*, 16:401–414.
- Taylor, T. (1999). On self-reproduction and evolvability. In Floreano, D., Nicoud, J.-D., and Mondada, F., editors, *Advances in Artificial Life*, pages 94–103. Springer-Verlag.
- Wagner, G. P. and Altenberg, L. (1996). Complex adaptations and the evolution of evolvability. *Evolution*, 50:967–976.
- Walker, R. (1999). Niche selection and the evolution of complex behavior in a changing environment – a simulation. *Artificial Life*, 5:271–289.