Adaptation of Mutation Rates in a Simple Model of Evolution*

Mark A. Bedau   Robert Seymour
Reed College
3213 SE Woodstock Blvd., Portland, OR 97202, USA
FAX (505) 777-7769
`mab@reed.edu`   `rseymour@reed.edu`

August 11, 1995

Abstract

We have studied the adaptation of mutation rates in a simple model of evolution. The model consists of a two-dimensional world with a periodically replenished resource and a fluctuating population of evolving agents whose survival and reproduction are an implicit a function of their success at finding resources and their internal metabolism. Earlier work suggested that mutation rate is a control parameter that governs a transition between two qualitatively different kinds of complex adaptive systems, and that the power of adaptive evolution is maximized when the mutation rate is around this transition. This paper provides evidence that evolving mutation rates adapt to values around this transition. Furthermore, the mutation rates adapt up (or down) as the evolutionary demands for novelty (or memory) increase.

1 The Evolution of Evolvability

Complex adaptive systems generate order through adaptive evolution. This process is strongly affected by the availability of a suitable variety of viable

---

*In *Complex Systems: Mechanism of Adaptation*, edited by R. J. Stonier and X. H. Yu, IOS Press, Amsterdam, 1994, pp. 37-44

evolutionary pathways. Adaptation to continually changing environments requires "novelty". Yet evolutionary adaptations are built through successive improvements, which requires "memory" of incremental improvements. So, evolutionary adaptation must balance these competing demands for novelty and memory. Furthermore, this balance must change as the context of evolution evolves. Thus, to understand how evolution generates order one must address the following sorts of questions:

- Where is the point at which novelty and memory are balanced?
- How does this balance shift during the course of evolution?
- Can the mechanisms that structure evolution adapt to this shifting balance?

One way to address these questions is to study the evolution of evolvability, i.e., to let the mechanisms that structure *first-order* evolution be subject to *second-order* evolution. Since perhaps the simplest mechanism that controls the variety of viable evolutionary alternatives is the mutation rate, we have studied the evolution of mutation rates.

The evolution of mutation rates bears on basic issues in both evolutionary biology and evolutionary programming. A fundamental open question in evolutionary biology, according to John Maynard Smith [13], is whether mutation rates are as low as physically possible or whether they are at an optimal positive value. The premise that mutations are generally harmful has been used to argue that natural selection of mutation rates can only go in one possible direction: toward zero (e.g., see [16]). On the other hand, mathematical analyses of certain one-locus, two-allele models with a separate modifier (mutation rate) locus have shown that evolution yields optimal, positive mutation rates under some conditions (see, e.g., [9, 8, 10, 12]). But it remains unclear what, if anything, such analytical results and methods could ever reveal about the evolution of mutation rates when there are many loci and many alleles per loci, and when a continually changing context of evolution keeps evolutionary dynamics far from equilibrium.

Evolutionary programming has started to study evolving mutation rates in an effort to automate control of evolutionary search for function optimizations. Preliminary efforts proved that automated control is feasible (e.g., [7, 1]), and continuing research is fine-tuning this process (e.g., [6, 15]). The potential technological value of evolutionary programming is obvious, but all the above-cited work presumes that evolution is driven by a fixed and externally-specified fitness function. Thus, its bearing on the most interesting and difficult cases of evolution—those that involve implicitly-specified fitness functions that change unpredictably in the course of evolution [14]—is uncertain.

The present work—empirical observations of evolving mutation rates in a computer model with many loci and many alleles per loci, and with an implicit fitness function that can continually change as the population evolves—should complement existing work in evolutionary biology and evolutionary programming. Our work is most akin to a study of evolving mutation rates in a host–parasite model [11], except that our fitness function is more implicit and we directly connect our work to a fundamental earlier result about evolving systems.

2 A Simple Model of Evolution

The model used here is designed to be simple yet able to capture the essential features of a population of agents with evolving sensorimotor functionality [14, 4, 5, 3, 2]. The system consists of many agents that exist together in an environment, in this case a toroidal lattice. The lattice has a resource field, which is driven by periodically adding (from an external source) a pyramidal pile of resources at a randomly chosen location. The agents are constantly extracting resources from their location and expending them through their behavior, so the agents function as the system's resource sinks and the whole system is dissipative.

In analogy with biological systems, the dynamics of the population as a whole is comprised of all the birth-life-death cycles of the agents. Births occurs when agents accumulate sufficient resources (see below), deaths occur when agents run out of resources, and the lives of agents consist of their interactions 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 t is marked by each agent interacting with the environment.

During its lifetime, each agent exchanges information with the environment by sensing and affecting the resource field in its local neighborhood. We assume that there is a discrete set of different possible sensory states,

$s \in S = \{s_1, ..., s_{N_S}\}$. The agents exchange no information with each other directly. Each agent 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 sensorimotor 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, $x_i^t \rightarrow x_i^{t+1} = x_i^t + \vec{b}_i^t$; in general, $\vec{b}_i^t = \varphi_i(s_i^t)$ (an exception is explained below). We assume that $\vec{b}_i^t$ is a member of a discrete set of different possible behaviors, $\vec{b}_i^t \in B = \{\vec{b}_1, \ldots, \vec{b}_{N_B}\}$. Each behavior causes a change in the agent's supply of resources: an augmentation from extracting resources at the agent's new location, a constant-sized reduction, and a reduction proportional to the distance moved.

The sensorimotor map φ_i operates on discrete sensory input $s_i^t \in \{s_1, \ldots, s_{N_S}\}$, so we may identify the function φ_i with its graph, $\{\varphi_{is}\}$, a set of N_S behavior values. Pursuing the biological analogy, we will consider the sensorimotor strategy elements $\{\varphi_{is}\}$ as i's *genome*, with N_S loci, and each particular element $\vec{b} = \varphi_{is}$ as the trait (or allele) at the s^{th} locus of i's genome. The genome of each agent i contains one additional gene, μ_i—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 are all of i's genes). When an agent's resource supply exceeds a threshold, it reproduces and splits its resources with its child, and its child inherits the parent's strategy elements $\{\varphi_{is}\}$, except for mutations.

In order to investigate how natural selection affects the evolutionary dynamics of behavioral strategies and mutation rates, we introduce a *behavioral noise* parameter, B_0, defined as the probability that $\vec{b}_i^t$ is chosen at random from B rather than determined by $\varphi_i(s_i^t)$. If $B_0 = 1$, then agents survive and reproduce differentially, and children inherit their parents' strategy elements (except for mutations), but the inherited strategies $\{\varphi_{is}\}$ reflect only random genetic drift.

This model provides a simple setting for empirical study of the evolution of evolvability. Agents' immediate environments produce sensory states, which then trigger actions by means of the agents' sensorimotor maps, and these actions subsequently change the environment. In this way, the agents' evolving sensorimotor maps influence the conditions of their own subsequent evolution, and thus their implicit fitness functions are constantly buffeted by the contingencies of natural selection. This first-order evolution is structured by the variety of sensorimotor maps actually compared and tested by natural selection, and this variety is regulated by the mutation rates: the higher the mutation rates, the greater the variety. Thus, evolving mutation rates amount to second-order evolution—the evolution of evolvability.

3 Adaptation of Mutation Rates

Earlier we extensively studied the present model—but with *fixed* mutation rates. One global mutation parameter μ governed all agents; during reproduction, a parent's strategy elements $\{\varphi_{is}\}$ were copied over to its child, with probability μ of being chosen randomly from the set of possible strategy elements B. We measured several macroscopic quantities, including various types of population diversity. Since this model is resource-driven, we estimated the population's overall fitness by its efficiency at extracting resources from the environment. And since we make all our simulations pump resources into the model at the same rate, the population's overall fitness is inversely reflected by the residual resource in the environment, i.e., the resource in the world that is not contained in the agents. The basic picture that emerged from this work is that mutation rate is a control parameter governing a transition in the range $10^{-3} \leq \mu \leq 10^{-2}$ between two qualitatively different kinds of evolving systems—meta-stable, quasi-clonal systems, and randomly fluctuating systems—and that the capacity for effective adaptive evolution is maximized around this transition. We suggested [3] that these results could be explained by the hypothesis that mutation rates around the transition optimally balance the competing evolutionary demands for novelty (high mutation) and memory (low mutation). (More detailed discussion of these results can be found elsewhere [5, 3].)

Our present work investigates whether second-order evolution can adapt mutation rates to the transitional region, $10^{-3} \leq \mu \leq 10^{-2}$, at which first-order evolution is most effective. We extensively simulated evolving mutation

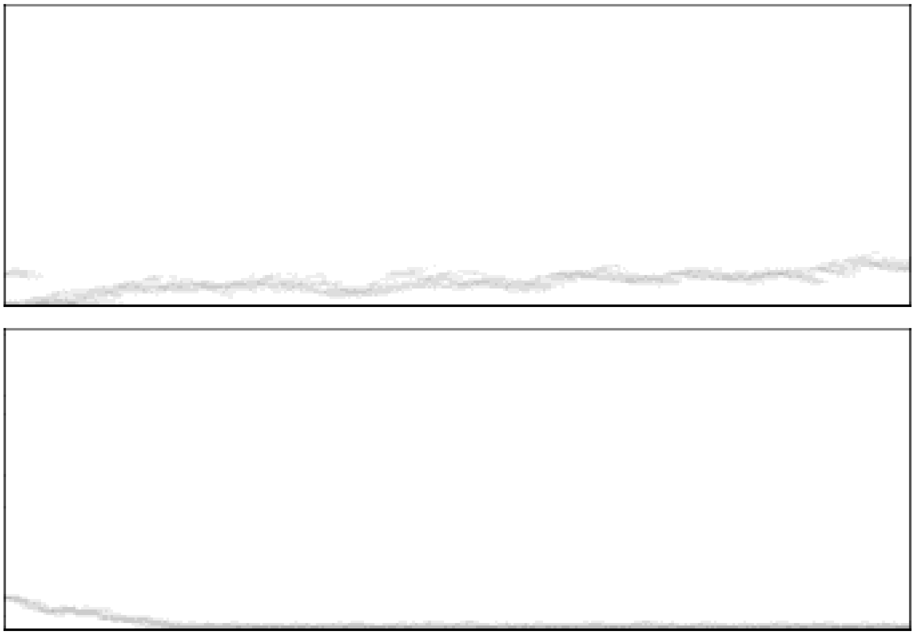

Figure 1: Typical dynamics of $M^t(\mu)$, with $\mu_\mu = 1.0$. Time on x-axis, for 10^5 timesteps. The y-axis depicts 100 equal-sized bins, $1.00 \le \mu_i < 0.99, \ldots, 0.01 \le \mu_i \le 0.00$ (top to bottom). Gray-scale shows a bin's population frequency (darker for greater frequency). Above: no adaptive evolution ($B_0 = 1$). Below: adaptive evolution ($B_0 = 0$).

rates, holding all model parameters constant except for systematically varying the meta-mutation rate, μ_μ, and switching adaptive evolution on and off ($B_0 = 0$ and $B_0 = 1$). In the initial population sensorimotor genes were assigned randomly and initial mutation-rate genes were either assigned randomly or all set to an arbitrary value. We then observed various macroscopic properties of the model, including the distribution of mutation rates,

$$M^t(\mu) = \sum_{i \in I^t} \delta(\mu_i - \mu), \tag{1}$$

where $\delta(z)$ is the Dirac delta function, which yields one if z is zero and yields zero otherwise.

As one could expect, when the agents' sensorimotor genes merely drift ($B_0 = 1$), the mutation distribution drifts up and down (e.g., figure 1, top),

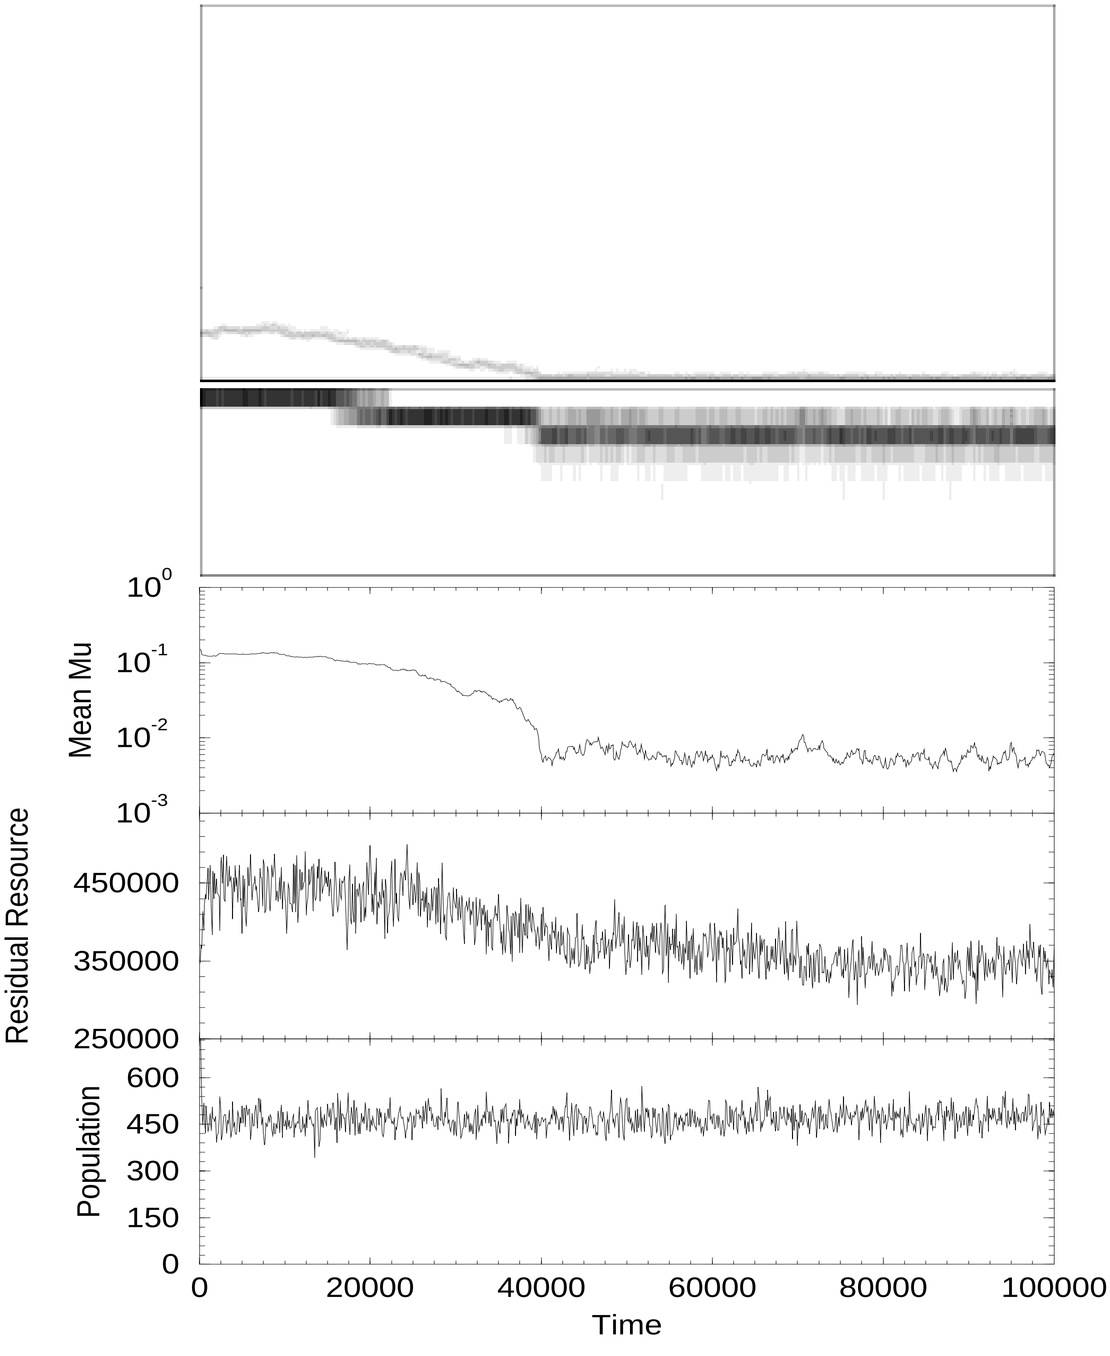

Figure 2: Typical time series (1000 data points) with first- and second-order adaptive evolution ($B_0 = 0$), showing the mutation distribution $M^t(\mu)$ (with two forms of binning), population mean mutation rate $\langle \mu_i \rangle_{i \in I^t}$ (note the normal-log scale), and residual resource; $N_S = 1024$ and $\mu_\mu = 0.66$. At the top, the mutation distribution's y-axis depicts 100 equal-sized bins. Second from the top, the mutation distribution's y-axis depicts ten exponentially-smaller-sized bins (from top, $10^0 \leq \mu_i < 10^{-1}$, $10^{-1} \leq \mu_i < 10^{-2}$, $10^{-2} \leq \mu_i < 10^{-3}$, etc.).

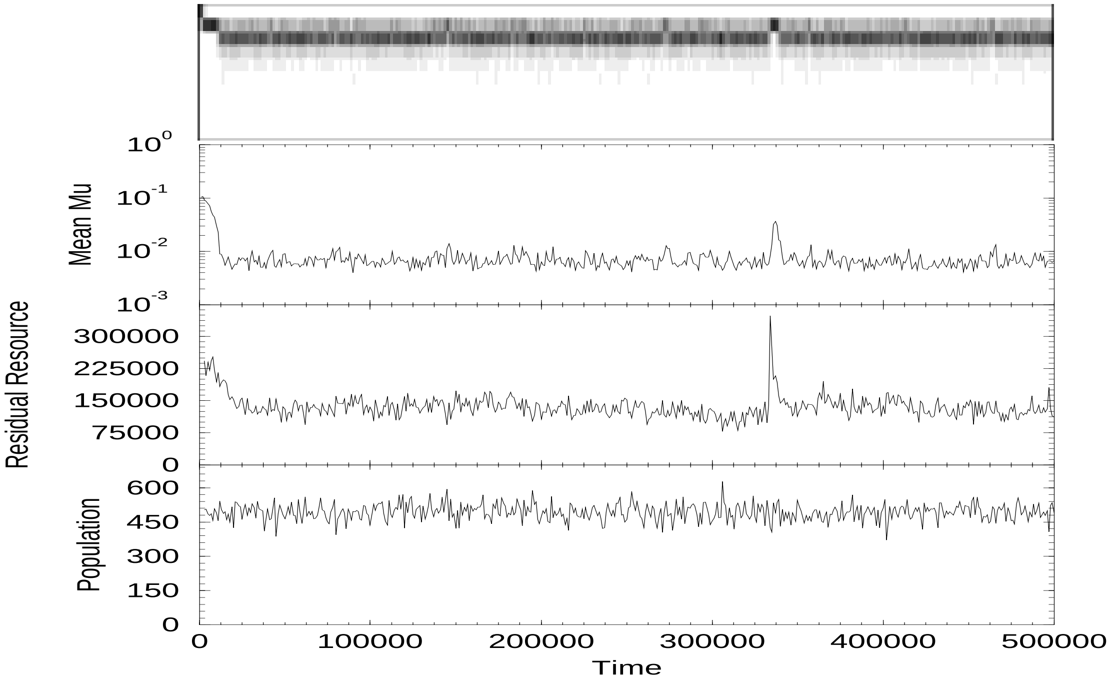

Figure 3: Typical perturbed time series (500 data points) with adaptive evolution ($B_0 = 0$), showing the mutation distribution $M^t(\mu)$ (with ten exponentially-smaller-sized bins, $10^0 \leq \mu_i \leq 10^{-1}$, $10^{-1} < \mu_i \leq 10^{-2}$, etc.), mean mutation rate $\langle \mu_i \rangle_{i \in I^t}$ (note the normal-log scale), residual resource, and population level; $N_S = 32$ and $\mu_\mu = 0.66$. At timestep 333334 all genes—except those governing mutation-rates—are randomized.

with the width of the distribution and the rate of change typically being proportional to μ_μ. (Note that $M^t(\mu)$ does not spread out indefinitely because the population is continually truncated due to the random contingencies of finding resources.) On the other hand, when there is first- and second-order adaptive evolution ($B_0 = 0$), the mutation distribution eventually becomes clustered at low mutation rates (figure 1, bottom), and the width of the distribution and the rate at which the distribution falls are typically proportional to μ_μ, as one would expect. When we resolve the low end of the mutation distribution in exponentially-smaller-sized bins we can clearly see second-order evolution pushing the bulk of the mutation distribution into the bin collecting mutation rates in the range $10^{-3} < \mu_i \leq 10^{-2}$, which exactly corresponds to the previously observed transition. (E.g., in the second mutation distribution shown in figure 2, these transitional mutation rates are represented in the third bin from the top, which we can see are those mutation rates most heavily represented in the population after the mutation rates have adapted.) Population mean mutation rates corroborate this conclusion. These effects are robust across all kinds of initial conditions, even those which are so extreme as to effectively prevent adaptive evolution (e.g., when all agents in the initial population have mutation rates equal to one).

In addition, as figure 2 illustrates, the drop in mean mutation rate corresponds to a marked increase in population fitness (i.e., decrease in residual resource)—clear evidence that the mutation rate drop is associated with increasing functionality of sensorimotor maps. This fitness increase is not associated with any notable increase in population level; rather, the agents are evidently becoming much more efficient at extracting available resources from their environment.

These results further the case for the evolutionary significance of the transition. More important, they present clear evidence that second-order evolution can adapt mutation rates to the region at which first-order evolution is most effective.

If we combine this conclusion with our earlier hypothesis that transitional mutation rates optimally balance the competing demands for novelty and memory, we would predict that the mutation distribution can be *raised* temporally by suddenly increasing the demands for evolutionary novelty. To test this hypothesis, we occasionally externally perturbed the model after first- and second-order evolution had substantially progressed, forcing evolutionary adaptation to start from scratch. For example, one kind of perturbation we employed was to randomize each agents' sensorimotor genes but leave their mutation-rate genes untouched.

These perturbation data confirm our prediction: the populations' mutation distribution adapts upward if we raise the mutation rate at which novelty and memory are presumably balanced. Figure 3 shows data from a typical perturbed simulation. The spike in residual resource shows the dramatic drop in population fitness caused by the the external perturbation.[1] Notice that the perturbation is quickly followed by a sharp rise in the mutation distribution and the mean mutation rate, when the demands for novelty increase. When the demands for memory increase shortly thereafter, the mutation distribution falls back to around its earlier level. Population level fluctuations play no evident role in these effects.

4 Conclusion

The present results confirm our earlier hypothesis that maximal fitness around transitional mutation rates is due to an optimal balance of the demands for evolutionary novelty and memory. In addition, they show that second-order evolution of evolvability can adapt the mechanisms of first-order evolution as the demands for novelty and memory unpredictably shift. These results provide an especially simple illustration of how the evolution of evolvability creates and tunes the capacity of complex adaptive systems to generate order through adaptive evolution.

Acknowledgements

Thanks to Norman Packard for helpful discussion.

---

[1] Notice that the residual resource levels in this simulation are significantly lower than those in figure 2, apparently due to the significantly smaller genome ($N_S = 32$ in figure 3, $N_S = 1024$ in figure 2). Evidently, on these time scales, fitness can become significantly higher if evolution has a smaller genome to adapt. For all size genomes studied ($34 \leq N_S \leq 1024$), the mean mutation rate after perturbation adapts to the same transition region.

References

[1] Bäch, T. (1992), "Self-Adaptation in Genetic Algorithms". In: F. J. Varela and P. Bourgine (Eds.), *Towards a Practice of Autonomous Systems*, Bradford/MIT Press, Cambridge, MA, pp. 263–271.

[2] Bedau, M. A. (1994), "The Evolution of Sensorimotor Functionality". In: J. -D. Nicoud (Ed.), *From Perception to Action*, IEEE Press, New York, in press.

[3] Bedau, M. A. and Bahm, A. (1994), "Bifurcation Structure in Diversity Dynamics". In: R. Brooks and P. Maes (Eds.), *Artificial Life IV*, Bradford/MIT Press, Cambridge, MA, in press.

[4] Bedau, M. A. and Packard, N. H. (1991), "Measurement of Evolutionary Activity, Teleology, and Life". In: C. G. Langton, C. E. Taylor, J. D. Farmer, and S. Rasmussen (Eds.), *Artificial Life II*, Addison-Wesley, Redwood City, CA, pp. 431–461.

[5] Bedau, M. A., Ronneburg, F., and Zwick, M. (1992), "Dynamics of Diversity in an Evolving Population". In: R. Männer and B. Manderick (Eds.), *Parallel Problem Solving from Nature, 2*, Elsevier, New York, pp. 94–104.

[6] Davis, M. W. (1994), "The Natural Formation of Gaussian Mutation Strategies in Evolutionary Programming". In: A. V. Sebald and L. J. Fogel (Eds.), *Proceedings of the 3rd Annual Conference on Evolutionary Programming*, World Scientific, River Edge, N.J., in press.

[7] Fogel, D. B., Fogel, L. J., and Atmar, J. W. (1991), "Meta-Evolutionary Programming". In: R. R. Chen (Ed.), *Proceedings of the 25th Asilomar Conference on Signals, Systems and Computers*, Maple Press, San Jose, pp. 540–545.

[8] Gillespie, J. H. (1981), "Evolution of the Mutation Rate at a Heterotic Locus", *Proceedings of the National Academy of Science, U.S.A.*, **78**, pp. 2452–2454.

[9] Gillespie, J. H. (1991), "Mutation Modification in a Random Environment", *Evolution*, **35**, pp. 468–476.

[10] Holsinger, K. E., and Feldman, M. W. (1991), "Modifiers of Mutation Rate: Evolutionary Optimum with Complete Selfing", *Proceedings of the National Academy of Science, U.S.A.*, **80**, pp. 6732–6734.

[11] Kaneko, K., and Ikegami, T. (1992), "Homeochaos: Dynamic Stability of a Symbiotic Network with Population Dynamics and Evolving Mutation Rates", *Physica D*, **56**, pp. 406–429.

[12] Liberman, U., and Feldman, M. W. (1986), "Modifiers of Mutation Rate: A General Reduction Principle", *Theoretical Population Biology*, **30**, pp. 125–142.

[13] Maynard Smith, J. (1989), "The Limitations of Evolutionary Theory". In: J. Maynard Smith, *Did Darwin Get It Right?*, Chapman and Hall, New York, pp. 180–191.

[14] Packard, N. H. (1989), "Intrinsic Adaptation in a Simple Model for Evolution". In: C. G. Langton (Ed.), *Artificial Life*, Redwood City, CA, Addison-Wesley, pp. 141–155.

[15] Saravanan, N., and Fogel, D. B. (1994), "Learning Strategy Parameters in Evolutionary Programming: An Empirical Study". In: A. V. Sebald and L. J. Fogel (Eds.), *Proceedings of the 3rd Annual Conference on Evolutionary Programming*, World Scientific, River Edge, N.J., in press.

[16] Williams, G. C. (1966), *Adaptation and Natural Selection*, Princeton University Press, Princeton, N.J.