Optimal Mutation Rates and Selection Pressure in Genetic Algorithms

Gabriela Ochoa, Inman Harvey and **Hilary Buxton**
Centre for Computational Neuroscience and Robotics
COGS – University of Sussex, Brighton BN1 9QH, UK
E-mail: {gabro, inmanh, hilaryb}@sussex.ac.uk

Abstract

It has been argued that optimal per-locus mutation rates in GAs are proportional to selection pressure and the reciprocal of genotype length. In this paper we suggest that the notion of *error threshold*, borrowed from molecular evolution, sheds new light on this argument. We show empirically the existence of error thresholds in GAs running on a simple abstract landscape; and then investigate a real-world industrial problem, demonstrating comparable phenomena in a practical application. We study the correspondence between error thresholds and optimal mutation rates on these two problems, and explore the effect of different selection pressures. Results suggest that error thresholds and optimal mutation rates are indeed correlated. Moreover, as the selection pressure increases, both error thresholds and optimal mutation rates increase. These findings may have practical consequences, as heuristics for measuring error thresholds in real-world applications will provide useful guidelines for setting optimal mutation rates.

1 INTRODUCTION

The performance of a GA heavily depends on the choice of its main control parameters: population size, mutation rate, and recombination rate. Despite research so far, there are no general heuristics on how to set them. It has been suggested that the mutation rate is the most sensitive GA control parameter (Schaffer et al., 1989; Bäck, 1996). Several studies in the literature look for "optimal" mutation rates (Hesser & Männer, 1992; Mühlenbein, 1992; Bäck, 1992, 1993), and optimal schemes for varying the mutation rate over a single run (Fogarty, 1989; Bäck & Schütz, 1996). The studies by Hesser and Männer (1992), Mühlenbein (1992), and Bäck (1992, 1993, 1996) coincide in that optimal per-locus mutation rates depend mainly on $1/L$ (the reciprocal of the genotype length). Moreover, Bäck (1996) suggests that as the selection pressure increases the optimal mutation rate also increases.

In this paper we suggest new foundations for the dependency of optimal mutation rates on both the selection pressure and the reciprocal of the genotype length. Also, the sensitivity of the mutation rate is explained by this new viewpoint. This knowledge comes from the field of molecular evolution, in particular from the notion of *error thresholds* (Section 3). The error threshold is the critical mutation rate beyond which structures created by an evolutionary process are destroyed more frequently than selection can reproduce them. The existence of this phenomenon in GAs and its relationship with the more familiar notion of optimal mutation rates, has been introduced in Ochoa et al. (1999) for simple landscapes. Here, we explore empirically optimal mutation rates and error thresholds on a real-world engineering problem — the *Wing-Box* design optimization problem. In particular, we study the dependence of both optimal mutation rates and error thresholds on the selection pressure. Also, we are interested in the relationship between error thresholds and optimal mutation rates. These issues are initially explored on abstract toy problems (e.g. the Royal Staircase functions, described in section 4.1), thus the following step would be to explore whether similar phenomena occurs in a real-world problem. The Wing-Box problem was formulated originally in the framework of the GAME (Genetic Algorithms in Manufacturing Engineering) project at COGS, University of Sussex. British Aerospace provided industrial realistic data for the definition of this problem (Section 4.2).

The remainder of this document is organized as follows. Section 2 discusses the notion of selection pressure in GAs, describes ranking and how this selection scheme allows control over the selection pressure. Section 3 introduces the notions of *quasispecies* and error thresholds from the field of molecular evolution, it also discusses the hypothesized relationship between error thresholds and optimal mutation rates. Section 4 describes the test problems used in this paper: the Royal Staircase, and the Wing-Box problem. Sections 5 and 6 describe our methods and results respectively, and, finally, Section 7 summarizes our findings.

2 SELECTION PRESSURE

In GAs, selection allocates reproductive opportunities for each organism in the population. The fitter the organism, the more times it is likely to be selected for reproduction. Selection has to be balanced with variation from mutation and recombination — the *exploitation-exploration* balance. Several selection mechanisms have been suggested in the literature, although there are no general guidelines on which to use on a given circumstance. There is, however, the notion of different *selection pressures* associated with each selection mechanism. Selection pressure is an informal term that indicates the strength of a selection mechanism. Loosely, the selection pressure measures the ratio of maximum to average fitness in the population.

For the experiments in this paper we used *rank selection*. This selection scheme is parameterized in such a way that allows control over the selection pressure. It also eliminates the need for fitness scaling mechanisms.

2.1 RANK SELECTION

In rank selection, individuals in the population are ranked according to fitness. The expected value of each individual depends on its rank rather than on its absolute fitness. The *linear ranking* method proposed by Baker (1985) works as follows: Organisms in the population are ranked in increasing order of fitness, from 1 to M (the population size). The user chooses the expected value Max ($Max \geq 0$) of the individual with rank M. The expected value of each individual i in the population at time t is given by:

$$ExpVal(i,t) = Min + (Max - Min)\frac{rank(i,t) - 1}{M - 1} \tag{1}$$

Where Min is the expected value of the individual with rank 1. Given the constraints $Max \geq 0$ and $\sum_i ExpVal(i,t) = M$, it is required that $1 \leq Max \leq 2$ and $Min = 2 - Max$. At each generation the individuals in the population are ranked and assigned expected values according to equation 1, Baker recommended $Max = 1.1$ and showed that this scheme compared favorably to proportional selection on some selected test problems. The selective pressure of linear ranking can be varied by tuning the maximum expected value Max (see Equation 1), which controls the slope of the linear function. The value recommended by Baker (1985) of $Max = 1.1$ means that, on average, the best individual is expected to be sampled 1.1 times, this is a rather moderate selective pressure, close to the extreme case of a random walk ($Max = 1.0$). The maximum possible expected value for linear ranking is $Max = 2.0$.

3 QUASISPECIES AND ERROR THRESHOLDS

Quasispecies theory was derived in the 70s by Eigen and Schuster (1979) to describe the dynamics of replicating nucleic acid molecules under the influence of mutation and selection. This theory was originally developed in the context of pre-biotic evolution, but in a wider sense it describes any population of reproducing organisms. An important concept in quasispecies theory is the notion of error threshold of replication. If replication were error free, no mutants would arise and evolution would stop. On the other hand, evolution would also be impossible if the error rate of replication were too high (only a few mutations produce an improvement, but most will lead to deterioration). The notion of error threshold allows us to quantify the resulting minimal replication accuracy (i.e. maximal mutation rate) that still maintains adaptation.

This can be seen at its clearest in an extreme form of a fitness landscape which contains a single peak of fitness $\sigma > 1$, all other sequences having a fitness of 1. With an infinite population there is a phase transition at a particular error rate p, the mutation rate at each of the ν loci in a sequence. In Eigen and Schuster (1979), this critical error rate is determined analytically (Equation 2), and it is defined as the rate above which the proportion of the infinite population on the peak drops to chance levels.

$$p = \frac{ln(\sigma)}{L} \tag{2}$$

In Equation 2, σ represents the selective advantage of the master sequence over the rest of the population (i.e. the selection pressure), and L the chromosome length. In the simplest case, σ is the ratio of the master sequence reproduction rate (fitness) to the average reproduction rate of the rest.

As stated originally, the quasispecies model considered infinite and asexual populations. Later on extensions were developed that consider finite populations and recombination. Most quasispecies studies considered simple landscapes, including single peak landscapes, double peak landscapes and flat fitness landscapes. The work of Bonhoeffer and Stadler (1993), described below, studied error thresholds on more complex landscapes.

3.1 ERROR THRESHOLDS ON COMPLEX LANDSCAPES

Bonhoeffer and Stadler (1993) studied the evolution of molecular quasispecies on two different complex fitness landscapes, the Sherrington Kirkpatrick spin glass and the Graph Bipartitioning landscape. They described an empirical approach for locating error thresholds on these highly correlated landscapes. In order to locate the error threshold empirically, they simulated the evolution of a population at a constant error rate for 200,000 cycles, which proved long enough to reach equilibrium on several parameters of the population (the maximal fitness, the average fitness among others). The error threshold may be approached from below and above, with both methods producing similar results. To approach it from below, the simulation starts with a homogeneous population at the global optimum. Then the population is allowed to reach equilibrium at a constant mutation rate of 0.0. Afterwards, the mutation rate is increased by a fixed, small step and the computation continues with the current population. This process continues until a predefined maximum for the mutation rate, p_{max}, is reached. To approach the error threshold from above, the simulation starts with a random population. Then the population is allowed to reach equilibrium at a constant mutation rate of p_{max}. Afterwards, the mutation rate is decreased by a fixed step and the computation continues with the current population. This process is repeated until the mutation rate is 0.

For both approaches, the *consensus sequence* in the population is calculated at the end of each simulation cycle for each mutation step. The consensus sequence is defined as the sequence of predominant symbols (bits) in each position; it is plotted as follows: if the majority of individuals has a '1' or '0' in a position i the field is plotted white or black, respectively. The field is plotted grey if the position is undecided (see Figures 3 - 5, for similar plots). The error threshold is characterized by the loss of the consensus sequence, i.e. the genetic information of the population. Beyond the error threshold the consensus sequence is no longer constant in time.

The empirical method described above was developed using the quasispecies equations as the underlying model of evolution. In this paper, we borrow this approach but use a GA instead of the quasispecies model for simulating evolution. The resulting method can be applied to locate error thresholds in GAs running on general complex landscapes.

3.2 ERROR THRESHOLDS AND OPTIMAL MUTATION RATES

The notion of error threshold seems to be intuitively related to the idea of an optimal balance between *exploitation* and *exploration* in genetic search. Too low a mutation rate implies too little exploration; in the limit of zero mutation, successive generations of selection remove all variety from the population, and once the population converges to a single point in genotype space all further exploration ceases. On the other hand, mutation rates can be too excessive; in the limit, where mutation places a randomly chosen allele at every locus on a genotype, the evolutionary process degenerates into random search with no exploitation of the information acquired in preceding generations.

Any optimal mutation rate must lie between these two extremes, but its precise position will depend on several factors including, in particular, structure of the fitness landscape. It can, however, be hypothesized that where evolution proceeds through a successive accumulation of information, then a mutation rate close to the error threshold is an optimal mutation rate for the landscape under study. This mutation rate should maximise the search done through mutation subject to the constraint of not losing information already gained.

4 TEST PROBLEMS

4.1 THE ROYAL STAIRCASE FAMILY OF FUNCTIONS

The *Royal Staircase* family of functions was proposed by van Nimwegen and Crutchfield (1998) for analyzing epochal evolutionary search. These functions are related to the *Royal Road* functions (Mitchell et al., 1992). Although simple, Royal Staircase functions capture some essential elements found on complex problems, namely, the existence of highly degenerate genotype-to-phenotype maps (i.e. the mapping from genetic specification to fitness is a many-to-one function). Next, we present a description of the Royal Staircase class of fitness functions:

1. Genotypes are specified by binary strings $s = s_1 s_2 \ldots s_L, s_i \in \{0,1\}$, of length $L = NK$.

2. Starting from the first position, the number $I(s)$ of consecutive 1s in a string is counted.

3. The fitness $f(s)$ of string s with $I(s)$ consecutive ones, followed by a zero, is $f(s) = 1 + \lfloor I(s)/K \rfloor$. The fitness is thus an integer between 1 and $N+1$, corresponding to 1 plus the number of consecutive fully-set blocks starting from the left.

4. The single global optimum is $s = 1^L$; namely, the string of all 1s.

Fixing N (number of blocks) and K (bits per block) determines a particular problem or fitness landscape. For the experiments in this paper we selected $N = 3$ and $K = 10$, that is a string length of 30.

4.2 THE WING-BOX PROBLEM

The Wing-Box problem was formulated as part of the Genetic Algorithms in Manufacturing Engineering (GAME) project at COGS, University of Sussex [1]. An industrial partner, British Aerospace, provided data from a real Airbus wing box.

A common problem faced in the design of aircraft structures, is to define structures of minimum weight that can withstand a given load. Fig. 1 sketches the elements of a wing relevant to this problem. The wing is supported at regular intervals by slid ribs which run parallel to the aircraft's fuselage. On the upper part of the wing, thin metal panels cover the gap separating adjacent ribs. The objective is to find the number of panels and the thickness of each of these panels while minimizing the mass of the wing and ensuring that none of the panels buckle under maximum operational stresses. More details, and the equations for calculating the fitness function, can be found in McIlhagga et al. (1996).

4.2.1 GENETIC REPRESENTATION - THE DELTA ENCODING

A full description of a potential solution to the Wing-Box problem requires the definition of the number of ribs N and the thickness of the $N-1$ panels. There is a constraint on the thickness of these panels which is that adjacent panels should not differ in thickness by more than 0.25 mm. The simplest way to accomplish this, is to encode the differences in thickness between adjacent panels rather than the absolute thickness of the panels. If we know the difference in thickness $\delta th(i)$ between panels i and $i+1$ for $i \in (1, N-1)$, the absolute thickness of the first panel is enough to define everything else.

[1] http://www.informatics.sussex.ac.uk/projects/game/

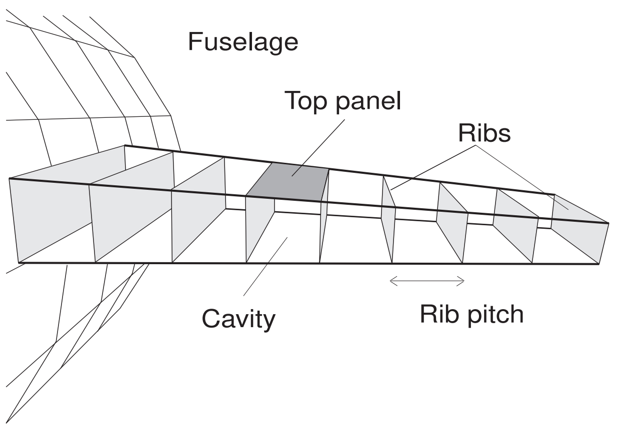

Figure 1: Relevant elements of a wing. Wing dimensions are fixed. The variable elements are the number of ribs and the thickness of the top panels.

N	th(1)	Δ th(1)= th(2)-th(1)	...	Δ th(i)= th(i+1)-th(i)	...	Δ th(N-2)= th(N-1)-th(N-2)

N: Number of ribs
th(i): Thickness of ith panel

Figure 2: Genetic representation of the wing parameters.

Originally, the Wing-Box parameters were encoded following the order described by Fig. 2. For the experiments in this paper we fixed the number of panels in 50 (i.e $N = 51$ ribs, since the number of ribs is 1 + the number of panels), thus our genetic encoding is the same, but excluding the first gene. The thickness of the first panel was allowed to vary between 10 and 15 mm by steps of 10^{-3} mm. This requires 5×10^3 values which can be represented with a minimum of 13 bits. For all subsequent $N-2$ panels the difference in thickness with the previous panel is encoded. According to manufacturing tolerance considerations, only five values were allowed for these differences in thickness: $\{-0.25, -0.125, 0.0, 0.125, 0.25\}$. Three bits are needed to encode these five values. Notice that a change in $\delta th(i)$ leads to changes in the thickness of panel $i+1$, and of all subsequent panels up to the tip of the wing. Notice also that in both the encoding of the first section, and the remainder $N-2$ sections, there is an amount of redundancy in the genotype to phenotype mapping. To sum up, the number of bits needed for encoding an individual is 13 for the first panel, and 3 for each of the others 49 panels, that is $13 + 3 \times 49 = 160$.

5 METHODS

All experiments were run using a generational GA with linear ranking and stochastic universal sampling. Three different selection pressures – strong, medium, and weak, were tested. These qualitative magnitudes correspond to setting the Max parameter in Equation 1 to 2.0, 1.5, and 1.1, respectively. The genetic operators were the standard bit mutation, and two-point recombination with a rate of 0.6. The mutation rate was expressed as mutations per genotype. Several mutation rates were explored. The population size was always 100. The string lengths were 30 for the Royal Staircase function, and 160 for the Wing-Box problem (see Section 4.2.1)

Two types of experiments were run. First, for calculating and producing the error thresholds plots (Figs. 3-8), we used the empirical described in section 3.1. We approached the error threshold from "above", that is, the simulation started from a random population. Then the population was allowed to reach equilibrium [2] at a constant high mutation rate ($5.0/L$ for the Staircase problem and $6.0/L$ for the Wing-Box Problem). Afterwards, the mutation rate was decreased by a small step (0.2 for the Staircase problem and 0.1 for the Wing-Box Problem) and the computation was continued with the current population. This process was repeated until the mutation rate was 0.0.

Second, for estimating optimal mutation rates, we calculated the number of evaluations before finding the optimum string on the Royal Staircase problem (averaged over 100 runs). For the more complex Wing-Box problem, where the optimum string is not known before-hand, the approach was to calculate the best-so-far fitness reached after a fixed number of function evaluations [3] (averaged over 20 runs).

[2]5,000 generations proved to be long enough to reach equilibrium on population best and average fitness on our test problems.

[3]5×10^5 evaluations proved to be enough to reach equilibrium on best-so-far average fitness.

6 RESULTS

6.1 ERROR THRESHOLDS

6.1.1 Royal Staircase Problem

Figure 3 shows results for strong and medium selection pressures on a Royal Staircase function with $N = 3$, $K = 10$. The pictures illustrate the existence of a stable consensus sequence for mutation rates below the error threshold. The error threshold is characterized by the loss of the consensus sequence. Beyond the error threshold the consensus sequence is no longer constant in time. In this case the consensus sequence below the error threshold is the single optimum string in the landscape (the string of all ones, displayed white below). The error thresholds for each fitness level or step can be observed. Error thresholds were shown to be lower for medium selection pressure.

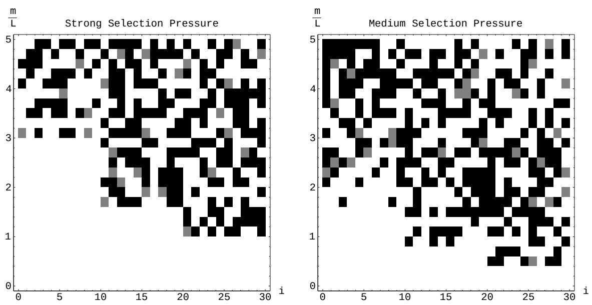

Figure 3: The consensus sequence on the Royal Staircase function ($N = 3, K = 10$), for strong and medium selection pressures. The X-axis shows the consensus bit (0 = black, 1 = white) for each position i, the Y-axis shows the mutation rate.

6.1.2 Wing-Box Problem

Figure 4 shows results for a strong selection pressure on the Wing-Box problem. Figure 4(Left) shows the existence of a stable consensus sequence for mutation rates below the error threshold. The error threshold is again visualized as the transition from a stable consensus sequence to a random sequence of bits. Notice that there is not a clear and single transition, from approximately bit 75 to bit 125 the error threshold looks higher than for the rest of the bits. Otherwise the transition seems to occur around 1.5 to 2.0 mutations per genotype. Figure 4(Right) was produced with the aim of highlighting the results, it plots the same data as figure 4(Left) but mapping the consensus sequence onto the string of all ones. Thus the error threshold is distinguished as the transition from a white to a random pattern of bits.

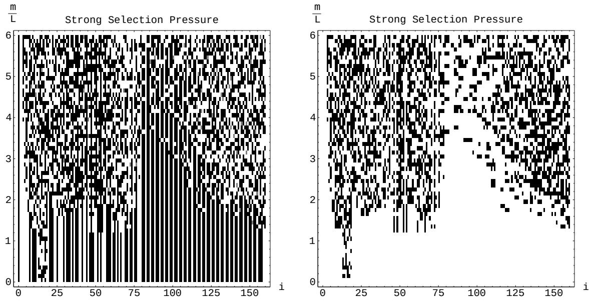

Figure 4: The consensus sequence on the Wing-Box problem for a strong selection pressure. Left figure shows the standard consensus sequence plot, whereas Right figure highlights the error thresholds by mapping the consensus sequence onto the string of all ones.

Figures 5 shows error thresholds for medium and weak selection pressures. The consensus sequence was again mapped onto the string of all ones to highlight the transition. Hence, the error threshold is distinguished as the transition from a white to a random pattern of bits. It can be noticed that the error threshold decreases as the selection pressure decreases. For the medium selection pressure the transition for most bits occurs around 1.0 mutations per genotype, whereas for the weak selection pressure it occurs around 0.1 mutations per genotype. Again there is no single transition for all bits in the genotype.

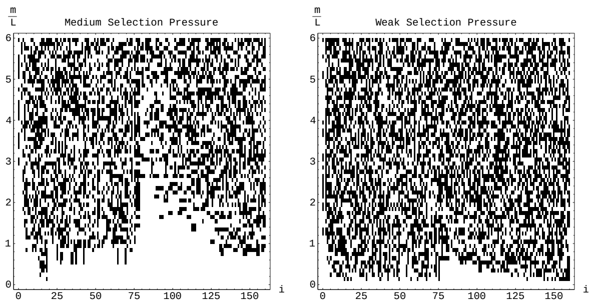

Figure 5: Error thresholds on the Wing-Box problem for medium ($Max = 1.5$) and weak ($Max = 1.1$) selection pressures.

6.2 OPTIMAL MUTATION RATES

6.2.1 Royal Staircase Problem

Curves in Figure 6 show the number of evaluations to reach the global maximum as a function of the mutation rate, for medium and strong selection pressures. Each data point gives the average of 100 runs. The standard deviations (not shown for the sake of clarity) were of the same order of magnitude as the average. Thus, there were large run-to-run variations in the time to reach the optimal string. Optimal mutation rates are those which find the peak with the least number of evaluations. Notice that there is not a single critically precise optimal mutation rate, but instead a range of mutation values producing near-optimal results. Optimal mutation rates were shown to be lower for the medium selection pressure. In the plots we indicate, with an arrow, the empirically estimated error thresholds. Error thresholds were found to be within the range of optimal mutation rates for both medium and strong selection pressures. Thus, these results support the hypothesized relationship between these two measures.

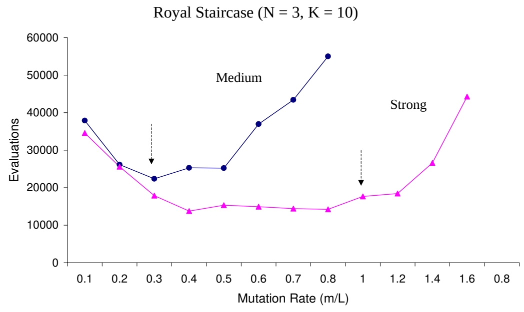

Figure 6: Number of evaluations for finding the optimum string as a function of the mutation rate on a Royal Stair case function. Arrows indicate the approximate magnitude of the error threshold in each case.

6.2.2 Wing-Box Problem

Curves in figures 7-9 show the average best-so-far fitness attained after 5×10^5 evaluations as a function of the mutation rate for strong, medium and weak selection pressures. Because we are dealing with a minimization problem, optimal mutation rates are those that produce the lower fitness value (the lower wing structure weight). Each data point gives the average of 20 runs. Notice that different mutation rate ranges and step sizes were used for each selection pressure. In each case a single mutation value can be distinguished as the one producing the minimal average weight. These optimal mutation rates are shown to decrease in magnitude as the selection pressure decreases.

Again we indicate in the plots, with arrows, the empirically estimated error thresholds. For the three selection pressures explored, error thresholds were found to be close to the empirically estimated optimal mutation rates. Thus, results for this real-world application also support the hypothesized relationship between error thresholds and optimal mutation rates.

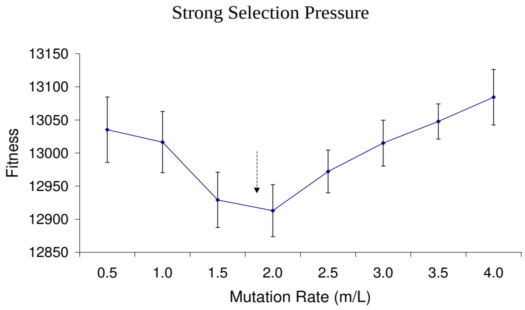

Figure 7: Average best-so-far fitness after 5×10^5 evaluations on the Wing-Box problem for a strong selection pressure ($Max = 2.0$). Error bars show $\pm$ the standard deviation. The arrow indicates the approximate error threshold.

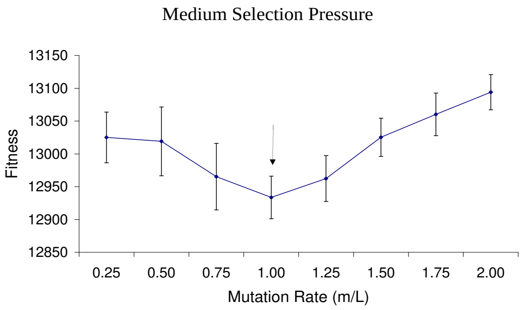

Figure 8: Average best-so-far fitness after 5×10^5 evaluations on the Wing-Box problem for a medium selection pressure ($Max = 1.5$). The arrow indicates the approximate error threshold. Error bars show $\pm$ the standard deviation

7 DISCUSSION

Our results suggest that error thresholds and optimal mutation rates are indeed correlated. Moreover, this relationship carried over from simple toy problems such as the Royal Staircase to a real-world application, the wing-box problem. The main implications of these finding are two-fold. First, theoretically, in helping

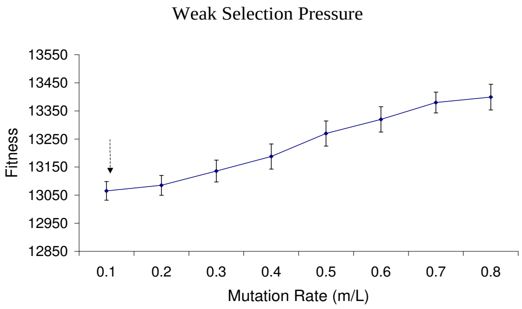

Figure 9: Average best-so-far fitness after 5×10^5 on the Wing-Box problem for a weak selection pressure ($Max = 1.1$). Error bars show $\pm$ the standard deviation. The arrow indicates the approximate error threshold

to understand GAs' behavior, as insights about error thresholds will shed light on our understanding of optimal mutation rates. Second, practically, as heuristics for finding error thresholds will provide useful guidelines for setting optimal mutation rates, thus improving the performance of GAs.

The consensus sequence plots (Figures 3 - 5), borrowed and adapted from theoretical biology (Bonhoeffer & Stadler, 1993), are new to the GA community. They represent a novel way to visualize the structure of fitness landscapes, since features such as the "step-ness" of the Royal Staircase function can be clearly noticed (Figure 3). They may serve as a tool to differentiate critical (and less critical) areas in the genotype, which may have practical implications when tackling real-world problems. First, it may be possible to infer important knowledge about an applied problem. Second, it may be possible to refine the genotype representations and optimal schedules for mutation rates as discussed below.

The strength of selection was shown to have a significant effect on the magnitude of both error thresholds and optimal mutation rates. The stronger the selection pressure, the higher these magnitudes. This suggests that the selection scheme used has to be taken into consideration when setting the mutation rate. This is particularly true when empirically comparing the performance of different selection mechanisms on a given problem. To be fair, comparisons should be made choosing the optimal mutation rate for each selection scheme.

As discussed above, there are differences in the error threshold magnitude across the genotype, which are more clearly observed for the Royal Staircase problem (Figure 3) than the wing-box problem. This supports the idea that a time-varying scheme for the mutation rate would be optimal. This idea was originally proposed by Fogarty (1989), who found that varying the mutation rate over time and across the bit representation of individuals (or both), significantly improved the performance of the GA. More recently, similar findings were reported by Bäck (1992) and Mühlenbein (1992). A clear implication of the findings here are that not only can useful estimates of optimal mutation rates be inferred from error thresholds but also that a systematic method of setting a non-fixed schedule of such rates can be devised for families of real world application problems. This, then, deserves future investigation.

Acknowledgements

Many thanks to A. Meier who kindly produced the Wing-Box figure. Special thanks to M. Sordo for careful corrections and critical reading. The Wing-Box fitness function source code was originally written by M. McIlhagga. The first author is funded by CONICIT, Venezuela.

Reference

Bäck, T., & Schütz, M. (1996). Intelligent mutation rate control in canonical genetic algorithms. In Zbigniew, R., & Michalewicz, M. (Eds.), *Proceedings of the Ninth International Symposium on Foundations of Intelligent Systems*, Vol. 1079 of *LNAI*. Springer.

Bäck, T. (1992). The interaction of mutation rate, selection, and self-adaption within a genetic algorithm. In und R. Manderick, B. M. (Ed.), *Parallel Problem Solving from Nature 2*. North-Holland.

Bäck, T. (1993). Optimal mutation rates in genetic search. In Forrest, S. (Ed.), *Proceedings of the 5th International Conference on Genetic Algorithms*. Morgan Kaufmann.

Bäck, T. (1996). *Evolutionary algorithms in theory and practice*. The Clarendon Press Oxford University Press. Evolution strategies, evolutionary programming, genetic algorithms.

Baker, J. E. (1985). Adaptive selection methods for genetic algorithms. In Grefenstette, J. (Ed.), *Proceedings of the 1st International Conference on Genetic Algorithms*. Lawrence Erlbaum Associates.

Bonhoeffer, S., & Stadler, P. (1993). Error thresholds on correlated fitness landscapes. *J. Theor. Biol.*, *164*, 359–372.

Eigen, M., & Schuster, P. (1979). *The Hypercycle: A Principle of Natural Self-Organization*. Springer-Verlag.

Fogarty, T. (1989). Varying the probability of mutation in the genetic algorithm. In Schaffer, J. (Ed.), *Proceedings of the 3rd International Conference on Genetic Algorithms*. Morgan Kaufmann.

Hesser, J., & Männer, R. (1992). Investigations of the M-heuristic for optimal mutation probabilities. In und R. Manderick, R. M. (Ed.), *Parallel Problem Solving from Nature 2*. North-Holland.

McIlhagga, M., Husbands, P., & Ives, R. (1996). A comparison of search techniques on a wing-box optimisation problem. *Lecture Notes in Computer Science*, *1141*.

Mitchell, M., Forrest, S., & Holland, J. H. (1992). The Royal Road for genetic algorithms: fitness landscapes and GA performance. In Varela, F. J., & Bourgine, P. (Eds.), *Proceedings of the First European Conference on Artificial Life*. MIT Press, Cambridge, MA.

Mühlenbein, H. (1992). How genetic algorithms really work: I. mutation and hillclimbing. In Männer, R., & Manderick, R. (Eds.), *Parallel Problem Solving from Nature 2*. North-Holland.

Ochoa, G., Harvey, I., & Buxton, H. (1999). Error thresholds and their relation to optimal mutation rates. In Floreano, J., Nicoud, D., & Mondada, F. (Eds.), *Proceedings of the Fifth European Conference on Artificial Life*. Springer-Verlag.

Schaffer, J., Caruana, R., Eshelman, L., & Das, R. (1989). A study of control parameters affecting online performance of genetic algorithms for function optimization. In Schaffer, J. (Ed.), *Proceedings of the 3rd International Conference on Genetic Algorithms* San Mateo CA. Morgan Kaufmann.

van Nimwegen, E., & Crutchfield, J. P. (1998). Optimizing epochal evolutionary search: Population-size independent theory. Tech. rep. Preprint 98-06-046, Santa Fe Institute.