
Parameter Perturbation Mechanisms in Binary Coded GAs with Self-Adaptive Mutation

Jim Smith

Intelligent Computer Systems Centre
University of the West of England, Bristol, U.K.
www.cems.uwe.ac.uk/~ jsmith
email: James.Smith@uwe.ac.uk

Abstract

Although the use of Self-Adaptation to control Mutation Rates within Genetic Algorithms is now well documented and accepted, there remains a variety of opinions as to what exactly “Self-Adaption” entails. While most authors agree on the encoding of strategy parameters within the genotype of the individual, and the use of selection at the phenotype level to implicitly control the relative frequency of different strategies, there is less agreement as to how these strategies should be perturbed. Given recent empirical and theoretical results concerning premature convergence with Self-Adaptive mutation, the issue of how best to generate variation within the space of strategy parameters clearly warrants further investigation.

In this paper we consider models of a generational mutation-selection genetic algorithm, with self-adaptation of the mutation rates using two different strategy perturbation mechanisms. Considering the Eigenvectors of these models, we derive predictions of the converged behaviour, and of the time taken to find the global optimum for two classes of problem, which differ in the nature of the barriers between regions of search space corresponding to different fitness classes. These predictions are compared to experimental results, and used to explain some anomalous observations. On the basis of these results we draw some conclusions about the behaviour of these two different types of algorithm on the two classes of problem.

Introduction

Although the notion of using Self-Adaptation to control Mutation Rates within Genetic Algorithms (GAs) is now well documented and accepted, there remain a variety of opinions as to what exactly “Self-Adaption” entails. While most authors agree on the encoding of strategy parameters within the genotype of the individual, and the use of selection at the phenotype level to implicitly control the relative frequency of different strategies, there is less agreement as to how these strategies should be perturbed, i.e. as to what variation operators should be used on the strategy encoding.

In this paper we consider this issue of perturbation, or equivalently, of the choice of variation operators within the search space of strategy parameters. We examine the behaviour of two different paradigms. In the first, perturbation is controlled by an exogenously provided parameter, which is equivalent to the use of a fixed mutation rate within strategy space. In the second, perturbation is endogenously controlled, i.e. the bitwise probability with which a strategy parameter’s encoding is varied is given by the value of that parameter. The approach taken is to examine the algorithms’ behaviour on two types of problem, differentiated by whether the regions of progressively higher fitness are separated from those below by an entropic barrier (typified by the Royal Road functions) or a fitness barrier (typified by a deceptive function).

We do this by first deriving infinite population “dynamical systems” models of the systems, which we use to predict the steady state distributions and the attractors of the system. We then derive approximations for the “Hitting Times” i.e. the number of evaluations required to locate the global optimum with finite populations. These are compared for the two paradigms and for fixed mutation rates, and verified experimentally.

The rest of this paper proceeds as follows:

- In Section 1 we provide more background on the issues at stake and the different perturbation methods in use.
- In Section 2, we briefly discuss some of the approaches that have been taken to modelling the behaviour of adaptive GAs, and introduce the dynamical systems models that we will use. We then develop models for the two algorithms and the two different types of evaluation function. We present predictions of behaviour and compare these with empirical results.
- In Section 4 we describe the derivation of approximations for the hitting times and calculate these for the Royal Road problems. We compare these with experimental observations, use the models to explain some of the observed behaviour and comment on some factors which affect the quality of the approximations.
- In Section 5 we contrast the behaviour of the two different perturbation mechanisms and draw some conclusions.

1 Strategy Perturbation Mechanisms

Any consideration of Self-Adaptive mutation probabilities naturally takes its cue from Evolution Strategies (ES) and Evolutionary Programming, where such methods have been

successfully applied for over two decades [11, 3]. For these algorithms, where a real-valued strategy parameter encodes the mutation step-size, it is typically perturbed by multiplication. This is either with a fixed amount (as in for example the “1:5 Success Rule” [8]), or with a random value from a distribution governed by a fixed global rate. Rudolph [10] has shown that unsuitable meta-learning rates can lead to premature strategy convergence and the loss of ability to escape from local optima. Glickman and Sycara [5] have observed similar phenomenon in real-coded GAs, and posited a number of situations under which this can occur.

The Self Adaptation of bit-wise mutation probabilities for a binary genotype was proposed by Bäck [1], who used a binary encoding for the mutation rate within a generational GA. This rate is perturbed in a two stage process. First the strategy parameter is decoded to give an initial rate, and then bit-wise mutation is applied to the strategy encoding at this rate. The updated strategy parameter is then decoded to find the rate to be applied to the problem encoding. This idea was expanded by Smith and Fogarty [13] who examined a number of different encodings (binary, Gray, and an exponential scheme due to [6]) within the context of a Steady State GA. They showed empirically that the choice of encoding, and hence of perturbation operator, has a significant effect on optimisation performance. These schemes have the property that ¹ changes are most likely to be to a similar rate, and that the probability that a rate remains the same (or very nearly so) increases as the encoded rate decreases. These properties are of course very similar to the schemes described above for mutation step sizes. They differ in that they do not require an external learning rate, but rather are endogenously controlled. Another significant difference is that the use of a fixed length encoding restricts the mutation rate to take one of a fixed set of values.

Stephens et al. [14] showed theoretically why the optimal choice of strategies is not only problem dependent, but also population dependent. Observing that the addition of strategy parameter genes created a non-injective mapping from genotype to phenotype, and considering endogenously controlled variation, they showed that adaptation of mutation rates can arise from a breaking of the symmetry in the genotype to phenotype redundancy. They use the term “effective fitness” to denote the fact that in general genotypes encoding for good solutions with low mutation rates are more likely to have similar (and hence fit) offspring than genotypes with the same problem encoding but high mutation rates.

In contrast to these approaches, Smith [12] examined a system which also used a fixed set of permissible values for the mutation rate, but which had an external parameter z controlling the probability of perturbation of the encoded strategy. On perturbation a new value was selected from a uniform random distribution rather than from a distribution which favoured similar values. In [7] the results are given of experiments which demonstrate that the mechanism described here is sufficient to permit adaptation of the mutation rates to optimal values on a range on NK landscapes. Subsequent work [15] confirmed that, even with $z = 1$, the algorithm found the global optimum more reliably on a range of binary coded test functions than an otherwise identical algorithm with ES-style log-normal perturbation of continuously varying strategy parameters.

The nature of search on a binary problem encoding, means that the effects of mutation (in terms of the distance moved in space) are effectively “binned” compared to variations

¹allowing for the presence of Hamming cliffs within the decoding function

in continuous space, which coupled with the greatly reduced strategy space may account for the improved performance seen with a fixed set of values. By using the same set of permissible mutation rates in the two cases, we can use this model to examine the question of whether the perturbation should be endogenously or exogenously controlled.

2 A Model of Self-Adaptation

2.1 Perturbation Probabilities

In this model a generational model is used similar to that of [1, 14, 12]: an individual is deemed to have a single mutation rate (m) attached to it which takes one of a fixed number (q) of values. Mutation is a two phase process, where first the value of m is varied to yield a new value m' , then the problem representation is mutated with this new bit-wise mutation probability.

The two strategy perturbation paradigms are thus defined by the perturbation probability distribution $P_p(m, m')$ governing the transition $m \rightarrow m'$.

For endogenously controlled variation, this distribution is a function of the length of the strategy encoding l_m , the Hamming distance $d_{m,m'}$ between the two rates, and the initial rate m :

$$P_p(m, m') = m^{d_{m,m'}} (1 - m)^{l_m - d_{m,m'}} \quad (1)$$

By contrast, for the simpler exogenously controlled scheme a new value m' is chosen at random with probability z , and the value is left unchanged with probability $(1 - z)$, i.e:

$$P_p(m, m') = \begin{cases} z/q & m' \neq m \\ (1 - z) + z/q & m' = m \end{cases} \quad (2)$$

We will refer to z as the Perturbation Rate. Note that by setting $q = 1$, we can use either model for fixed mutation rates as well.

2.2 Problem Classes

As mentioned above, in this paper we will consider problems of unitation, i.e. where the fitness of an individual solution i depends solely on u_i , the number of 1's in its binary representation. For such a problem with a representation of length L there are $L + 1$ equivalence classes of solutions with different representations but equal fitnesses. Although this at first may seem like a restriction to a small set of classes, we can in fact provide very different barriers to incremental increases in fitness by varying the way in which the fitness is given as a function of unitation. Firstly of course, we can provide no barriers, i.e. the One Max fitness function: $f(i) = u_i$.

This can be seen as the simplest example ($K = 1$) of a type R1 Royal Road fitness function [4] where the fitness is given by the number of blocks "aligned" to the target string in a problem with L blocks, each of length K . Assuming (without loss of generality) that the target string is all 1s for each block:

$$f_{R1}(i) = \sum_{l=1}^L \prod_{j \in block_l} i_j \quad (3)$$

A well known property of these functions is that for $K > 1$ they possess “plateaus” of equal fitness, and as discussed in [18] these represent entropic barriers to evolutionary search. As has been extensively documented, search on these problems typically proceeds via a series of “epochs”, and during transitions the entropy of the population is reduced as the correct alignment is found for the next block, and fixated through the population.

In order to construct a mapping from the R1 class of problems onto functions of unitation, we simply replace the unitation of a string by the number of aligned blocks it contains, and define the probability of mutating the number of aligned blocks as a function of the parameter K . In other words, we can treat them as OneMax functions where rather than $P(0 \rightarrow 1) = P(1 \rightarrow 0) = m$, we use the notation that a block is aligned with probability A_i , and unaligned with probability D_i , for a member of class i . Under the Maximum Entropy assumption the block alignment and destroying probabilities for a class with mutation probability m are [17]:

$$A_i = \frac{1 - (1 - m)^K}{2^K - 1} \quad D_i = 1 - (1 - m)^K \quad (4)$$

An alternative class of problems which present a fitness barrier, rather than an entropic one to evolutionary progress to the global optimum is characterised by so-called “Trap” or deceptive functions. We will consider functions with a single deceptive partition whose fitness varies as a function of the unitation u_i (i.e. $K = 1$ in the notation above) as given by Deb [2]:

$$f(i) = \begin{cases} \frac{a}{c}(c - u_i) & u_i \leq c \\ \frac{b}{l-c}(u_i - c) & otherwise \end{cases} \quad (5)$$

Figure 1 illustrates this function, using the parameters from section 3.2 with length 16. Note that by considering a single partition with varying size, we can look at different orders of deception, without the need to take into account the position of different bits in the fitness calculation.

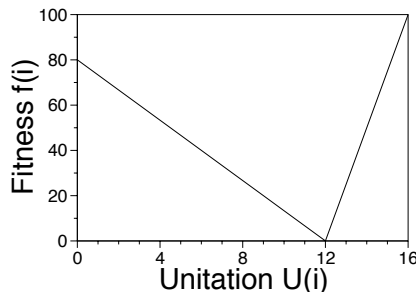


Figure 1 Deceptive function with parameters $\{a, b, l, c\} = \{80, 100, 16, 12\}$

2.3 The Dynamical Systems Model

The advantage of functions of unitation is that the fitness equivalence classes provide a “natural” coarse-graining for modelling. Whilst such coarse-grained approaches are not

without their restrictions, they do enable the construction of predictive models which are more tractable than exact models which consider all the different possible strings. In order to deal with the different mutation rates we will extend this model so that each of the fitness equivalence classes is subdivided into q further classes according to the mutation rate attached to the individual solutions, giving a total of $N = q * (L + 1)$ states. For an individual state i the fitness is $f'(i) = f(i/q)$ and the mutation rate is indexed by $m = i\%q$, where the % symbol has its usual modulus meaning, and the division i/q is taken to be rounded down to an integer value.

We can therefore define a population vector $\bar{p} = (p_1, \dots, p_N)$ such that the components p_i represent the proportions of the population in class i , subject to the restriction $\sum_i p_i = 1$.

Following Vose's model [19], we can model the effect of the GA on this population vector as:

$$\bar{p}' = G\bar{p} = MF\bar{p} \quad (6)$$

where the functions M and F represent the mixing (mutation and crossover) and selection operators respectively. The outcome (\bar{p}') represents the probability distribution from which the next population will be sampled, which is equivalent to the next generation in the *Infinite Population Model* (the reader is referred to [19] for a more detailed discussion).

In order to concentrate on the effects of different perturbation operators, we restrict ourselves to the case where there is no crossover, and note that the i^{th} component of the probability distribution vector for the next generation, is simply the sum over all classes j , of the probability of selecting a member of class j , multiplied by the probability of mutating that individual from class j to class i , i.e.

$$Gp_i = \sum_{j=1}^N P_m(ij) * P_s(j) \quad (7)$$

where $P_s(i) = Sp_i$ denotes the probability of selecting an individual of class i , and $P_m(ij)$ denotes the probability of creating a member of class i by mutation from one of class j .

For fitness proportional selection, the selection operator F can be modelled by using a diagonal matrix S with elements

$$S_{ij} = \begin{cases} f'(i) & i = j \\ 0 & i \neq j \end{cases} \quad (8)$$

and the operation of the selection operator is given by

$$Fp = \frac{Sp}{\langle f \rangle(p)} \quad \text{where} \quad \langle f \rangle(p) = \sum_{j=0}^N p_j f'(j) \quad (9)$$

We can simplify the specification of the mutation operator by noting that there is a two-phase process : first the correct mutation rate must be obtained, and then after that the number of aligned blocks must correctly changed. If a of the j/q aligned blocks become

unaligned and b of the $L - j/q$ unaligned blocks get aligned, the mutation from class j to class i happens with probability :

$$P'_m(ij) = \sum_{a=0}^{j/q} \sum_{b=0}^{L-j/q} \delta_{j/q-a+b, i/q} * P_p(m, m') \quad (10)$$

$$\left(\begin{array}{c} L - j/q \\ b \end{array} \right) \left(\begin{array}{c} j/q \\ a \end{array} \right) A_i^b (1 - A_i)^{L-j/q-b} D_i^a (1 - D_i)^{j/q-a}$$

where $\delta_{x,y}$ is the Kronecker delta function, and the appropriate choice of perturbation probabilities are used.

Nimwegen et al. have shown that for fitness proportionate selection with mutation, the combined operator $G = MF$ can be turned into a linear form, and derived equations for the calculation of the mean fitness as a function of time without the need for iterated matrix multiplication [17]. Smith [12] showed that this could be extended to cover his discrete exogenous model.

Given a linear form for M and F we can revisit (6) and consider the case for fixed points (if they exist) of the algorithm. If v is a fixed point of the system (i.e. $GV = v$) then we have $MSv = \langle f \rangle(v)v$, i.e. v is an eigenvector of MS with eigenvalue equal to its average fitness. For non-zero fitness and mutation G is positive, so the Perron-Frobenius theory tells us that there will only be one eigenvector of the exogenous system corresponding to a ‘‘real population’’ (i.e. which lies within the simplex), and that this corresponds to the population whose mean fitness is equal to the biggest eigenvalue. The predictive value of this analysis for finite population GAs has been illustrated in [17, 18, 9, 12]. For the case of the endogenous strategy, if the set of permissible mutation rates includes the value zero, then the $L + 1$ vertices of the simplex corresponding to the equivalence classes with zero mutation will also be fixed points of the system.

3 Behaviour of Mean Fitness

In order to examine the two different paradigms we implemented models for systems with $q = 64$, equivalent to a six bit strategy encoding. This was interpreted as a binary coded integer i and the mutation probability was scaled over the range $[0.0, 0.1]$ as $m = i/640$. Having defined the genetic operator G we solved for the eigenvectors and their corresponding eigenvalues in order to predict the fixed points of the systems.

Recalling that these are infinite population models, the predictive accuracy was verified experimentally, by recording the mean population fitness as a function of time for 500 generations of a GA with a popsize 1000, and calculating the mean and standard deviation of these values over 50 runs.

3.1 Entropic Barriers: Royal Road functions

We considered first Royal Road problems of 32 bits with different K (i.e. the number of blocks $= 32/K$). The fitness of a solution with i aligned blocks was $1.0 + iK$ (the constant allows us to avoid problems with degenerate matrices). For each paradigm we calculated

Mutation Rate		K 1	K 2	K 4	K 8
Static					
0.0005	Predicted	32.484	32.478	32.476	32.476
	Observed	32.446 (0.075)	31.257 (0.461)	31.210 (2.083)	23.319 (4.348)
0.001	Predicted	31.992	31.971	31.962	31.960
	Observed	31.968 (0.078)	31.948 (0.070)	31.942 (0.072)	27.157 (4.363)
0.05	Predicted	20.214	15.027	10.554	7.960
	Observed	20.220 (0.120)	15.331 (0.186)	11.308 (0.318)	8.588 (0.336)
0.1	Predicted	18.640	12.181	6.821	4.138
	Observed	18.630 (0.097)	12.397 (0.140)	7.328 (0.189)	4.378 (0.136)
Endogenous					
Predicted		33.000	33.000	33.000	33.000
Observed		31.677 (1.505)	32.000 (1.455)	30.920 (3.223)	24.200 (6.039)
Exogenous					
Z 0.05	Predicted	31.377	31.377	31.377	31.377
	Observed	30.730 (0.772)	30.870 (0.845)	30.334 (1.993)	26.827 (4.294)
Z 0.1	Predicted	29.758	29.758	29.758	29.758
	Observed	28.813 (0.642)	29.383 (0.837)	29.117 (1.384)	24.826 (3.667)

Table 1 Predicted and experimentally observed (with std. dev. for mean in braces) steady-state fitness values for Royal Road functions with varying block sizes.

the steady state population mean fitness, and we also performed these calculations for static mutation rates in the set $\{0.0005, 0.001, 0.05, 0.1\}$.

Table 1 shows the predicted and observed steady state fitnesses for the different paradigms and the static mutation rates. As can be seen there is close agreement between prediction and observations, except for the case of $K = 8$, where the observations are still within experimental error (2 standard deviations) of the predicted values, but have lower mean values. Closer inspection reveals that this is because not all runs have converged onto the same point in this case.

For the static mutation rate the mean performance tails off with increasing K and for increasing mutation rates. The application of higher mutation rates causes a corresponding increase in genotypic diversity (and hence entropy) and the preservation of fewer aligned blocks, since these correspond to a decrease in genotypic entropy. The extent to which lower fitness solutions are allowed to propagate depends of course on the selection pressure, which is relatively low in this case compared to e.g. ES's. However we note that the absolute and relative decrease in fitness is greater for the high K cases, where the selective difference is greater between fitness classes, pointing to the entropic effects as dominating.

Considering the Endogenous scheme, the predicted mean fitness lies on the global optimum, i.e. the population is predicted to fully converge to the single genotype corresponding

to a zero-mutation population at that point². In practice the mean fitness remains below the predicted rate within our experimental timescale. For $K < 8$ this appears to result from the fact that it takes a relatively long time for chance endogenous perturbations to drive mutation out of the system, especially when those perturbations are only happening with low probability. For $K=8$, some runs prematurely converged onto suboptimal solutions. As noted above, in addition to the feasible solution at the global optimum, there are attractors corresponding to lower “epochs”, i.e. lower numbers of aligned blocks. The fact that some runs converged to these attractors explains the high standard deviations observed.

Considering the exogenous scheme, we note that the use of a uniform random distribution prevents convergence to zero mutation at the global optimum. However we also note the observed results for this scheme with low z are similar to those for the endogenous scheme. In Figure 2 we plot the predicted and observed values for different K values as a function of the perturbation rate z . These are in very close agreement. The principal feature to note is that as the perturbation rate is increased so the mean fitness in the steady state decreases. This can be explained by the observation that as the perturbation rate increases, so too will the proportion of high mutation rates. Beyond a given value of Z (which increases with K) the rate of decrease of mean fitness falls off rapidly (and indeed the mean fitness is roughly constant for $K \in \{1, 2\}$) as the constructive and destructive effects of the high mutation rates cancel each other out.

Analysis of the eigensystems does not show the existence of low-fitness attractors that was seen for the endogenous case, and this explains the closer fit between prediction and observations as evidenced by Figure 2.

Overall for this class of problems the endogenous scheme appear to show behaviour which would correspond to better on-line performance

3.2 Fitness Barriers : Deceptive Partitions

In order to investigate the effect of fitness (rather than entropic) barriers between most regions of above-average fitness, and the global optimum, we consider a single function of the form (5) with varying orders of deception. This was achieved by varying l with $c = 0.75 * l$ and fixing $a = 100, b = 80, .$

Table 2 shows the predicted and observed behaviours for the orders of deception in the set $l = \{4, 8, 12, 16, 20\}$. From this we can make the following observations:

1. The model predicts (and this is verified experimentally) that above a certain static mutation rate selection is unable to maintain the population on the peak corresponding to the global optimum and there is a dramatic decrease in mean fitness. The decrease in fitness is more marked for higher orders of deception.
2. By contrast, the model predicts that both of the Self-Adaptive algorithms should be able to preserve a population of high mean fitness. This comes about as a result of locating the optimum and then selection “weeding out” the high mutation rates which cause more exploration and a loss of fitness.

²Note that unlike other fitness levels, the equivalence classes at the global optimum are only represented by one genotype for these problems

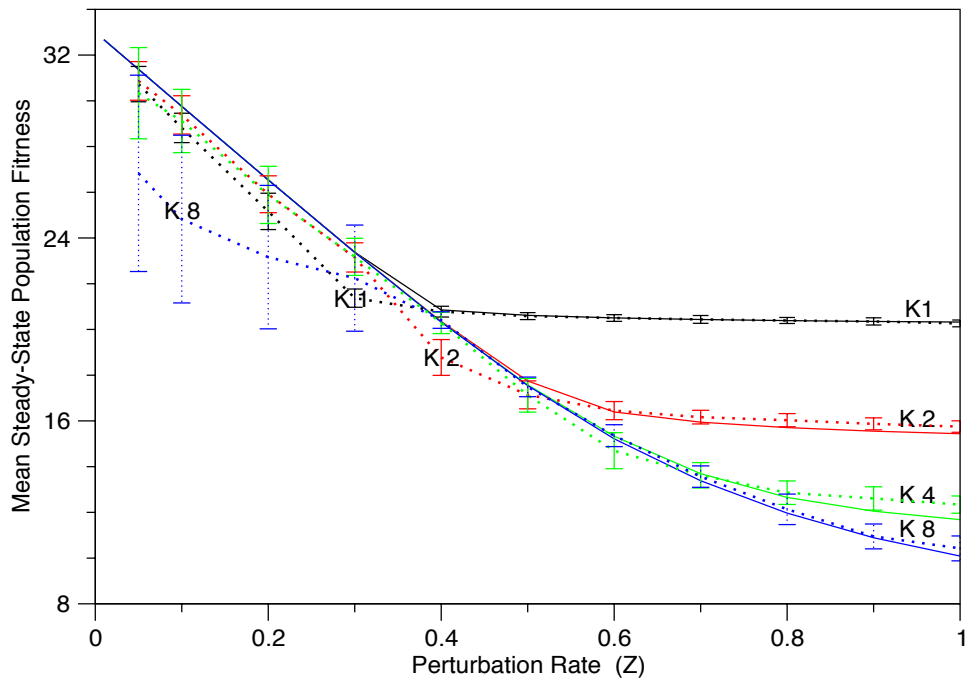


Figure 2 Mean Population Fitness at Steady-State for the Exogenous Scheme. Solid lines are predictions. Dotted lines are mean (with error bars at 1 standard deviation) of data from 50 runs with a GA of population size 1000 after 500 generations

3. The results for the exogenous scheme with increasing z clearly show how the increasing z (and hence the probability of re-introducing the higher mutation rates) leads to a corresponding decrease in mean fitness for the same order of deception. Even if selection only propagates copies of the global optimum, perturbation to higher mutation rates, with a correspondingly larger number of bits mutated, creates individuals of low fitness in the next generation.
4. Even with these relatively large population sizes, the effects of stochastic fluctuations on a finite population is sufficient to cause significant variation in performance for the larger problem sizes ($l = 16, 20$), as indicated by the differences between observed and predicted means, and the large standard deviations.

The predictive power of the modelling technique used comes into its own when trying to explain these experimental results.

In the exogenous strategy case, we predict the steady-state mean fitness by solving for the eigenvectors of the genetic operator G , and taking the eigenvalue corresponding to the single non-negative eigenvector. However, we can also analyse the other (infeasible) eigenvectors in order to see which of these lie close to the simplex and are likely to act as attractors for the system. When we perform this analysis we see that the values given

in Table 2 do correspond to the eigenvalues of the feasible eigenvector, i.e. the predicted mean fitness decreases with l and z . Using $z = 0.05$ as an example, shows another factor, for values of $l > 8$ an eigenvector with corresponding fitness 76.08 appears at distance of $7.83 * 10^{-10}$ for $l = 8$ decreasing to $6.58 * 10^{-13}$ for $l = 20$. These have corresponding mean mutation rates of ≈ 0.015 , in other words individuals with very low mutation probabilities are predicted to dominate the populations. Similar low fitness attractors occur for other values of z , and this explains the lower fitness observed empirically for the exogenous scheme. Interestingly, analysis of the eigensystem for $l = \{4, 8\}$ does not indicate the presence of near-feasible attractors close to the deceptive optimum, and the empirical results show a very low standard deviation, suggesting that all runs end up in the same place.

By contrast, for the endogenous scheme we know that the equivalence classes corresponding to a zero mutation rates are fixed points of the system, and that one of these exists for every different fitness level. Muhlenbein and Mahnig have shown (theorem 2.3 of [?]) that for $m = 0$, (with Gene Pool recombination which ensures linkage equilibrium) only those vertices of the simplex where the corresponding binary string is fitter than its Hamming neighbours are stable, i.e. those at the deceptive peak and the global optimum. For these mutation only algorithms, the assumption of Maximum Entropy takes the place of Linkage Equilibrium, and intuitively one can see that any displacement from the zero mutation vertices will create a non-zero probability of creating a population member with a different unitation. This suggests that all but the global optimum are unstable attractors, although the probability of mutation from the deceptive to the global optimum is small, making this much more stable than the $L - 1$ other vertices. Closer analysis confirms that the GA runs do indeed converge to one of these two places.

3.3 Summary of behaviour of mean fitness

Overall, when we compare the results for the two types (entropic and fitness-based) of barriers we see that they are broadly similar. Looking at the empirical results for the higher orders or deception/block sizes, we see that the endogenous scheme performs better than the exogenous on the deceptive functions, and worse on the Royal Road functions.

Analysis of the eigenvectors of the Genetic operator G in each case indicates the reason for this behaviour. For the entropic barriers with $K > 4$ there are low-fitness attractors close to the simplex for endogenous mutation and not for exogenous mutation. Empirically, with endogenous mutation we observe that the population gets “trapped” at points on sub-optimal fitness plateaus with zero mutation.

For the fitness barriers, the situation is reversed : there are fixed points on the simplex centred on both peaks for the endogenous case, regardless of the order of deception. For the exogenous case, as the order of deception increases so attractors appear close to the simplex but at a lower fitness than even the deceptive peak. Empirically we observe that the endogenous scheme tends to rapidly converge to one or other of the peaks, whereas the exogenous scheme tends to converge to the population centred around the sub-optima.

Mutation	l 4	l 8	l 12	l 16	l 20
Static					
0.0005 Pred.	99.80	99.60	99.40	99.20	99.07
Observed	99.81 (0.13)	99.60 (0.12)	96.21 (7.26)	85.71 (9.25)	83.17 (7.98)
0.001 Pred.	99.60	99.20	98.81	98.42	98.03
Observed	99.60 (0.22)	99.24 (0.19)	96.84 (5.90)	86.98 (9.69)	81.59 (7.14)
0.05 Pred.	81.45	67.78	57.50	49.80	47.14
Observed	81.47 (1.17)	67.87 (1.18)	56.69 (1.93)	49.53 (0.51)	47.10 (0.72)
0.1 Pred.	65.64	48.51	43.93	41.064	39.08
Observed	65.66 (1.47)	48.53 (0.67)	43.91 (0.57)	41.12 (0.54)	39.08 (0.51)
Endogenous					
Predicted	100.00	100.00	100.00	100.000000	100.00
Observed	100.00 (0.00)	100.00 (0.00)	97.60 (6.50)	87.61 (9.70)	83.50 (7.76)
Exogenous					
Z 0.01 Pred.	99.02	99.02	99.02	99.02	99.02
Observed	98.99 (0.33)	99.00 (0.22)	95.51 (7.64)	83.95 (8.47)	80.69 (5.46)
Z 0.05 Pred.	95.12	95.11	95.10	95.10	95.10
Observed	95.22 (0.73)	95.15 (0.47)	92.03 (6.99)	83.33 (9.29)	77.64 (5.14)
Z 0.1 Pred.	91.07	90.35	90.25	90.22	90.20
Observed	91.01 (0.87)	90.24 (0.80)	87.46 (6.64)	78.65 (8.80)	74.44 (5.88)

Table 2 Predicted and experimentally observed (with std. dev. for mean in braces) steady-state fitness values for deceptive partitions of different orders.

4 Approximation of Hitting Times

In the previous section, analysis of the principal eigenvector of the linearised function G was used to predict the convergence behaviour of GAs with finite populations. It was noted that the endogenous strategy had a number of of feasible fixed points corresponding to the vertices of the simplex with zero mutation for each fitness equivalence class. For the positive matrices arising with the exogenous strategy only one eigenvector lay within the simplex, and represented a “feasible” population, but that under certain circumstances other eigenvectors could lie very close to the simplex and act as attractors for the system. The empirical results showed that with finite populations it was possible for “real” GAs to get “stuck” close to (or on) these attractors.

This effect has been noted by other authors [17, 9] using a single fixed mutation rate. In particular Nimwegen [18] shows that with finite populations, in many cases an “epochal” type of behaviour is observed, and that the mean fitness values of the population during these epochs is closely related to the eigenvalues of the infeasible attractors. In [17] it is shown that the population vectors during epoch i ($0 \leq i \leq N$ with epoch N containing the highest fitness class) are centred around the principal eigenvector v^i of the restricted operator G^i . In [16] a mutation and population size dependent prediction of times to

optimum is derived, based on the assumption that the population stabilizes on every epoch of a ‘‘Royal Staircase’’ function. It was found that this correctly predicted the form of the behaviour, although the precise values were consistently underestimated.

There are a variety of reasons the predictions derived in [16] are not transferrable to the model of self-adaptation used here, arising from the fact that we have extended the number of equivalence classes to include a q -fold degeneracy at each fitness level. Instead we will derive initial estimates of the hitting time for the Royal Road function with $K > 2$ under the following assumptions:

1. For small finite populations, search proceeds in an epochal fashion, such that the time spent in each epoch searching the neutral basin for the portal to the next fitness class occupied is far greater than that spent in transition between epochs.
2. During each epoch i we can assume that the finite population at each generation is drawn from a distribution given by the principal eigenvector of the corresponding restricted Genetic Operator G^i
3. The Maximum Entropy condition holds, that is to say that when we consider all of the representatives of an equivalence class, for each of the unaligned blocks there is a uniform distribution of all the possible genotypes corresponding to that class.
4. Transitions between epochs occur when the action of mutation on a member of the currently fittest class i changes one unaligned block into an aligned one.

Under these assumptions we can estimate the hitting time as the sum of the expected times to locate the next fitness class for each epoch.

In the following we will use the notation e_j^i to denote the j^{th} component of the principal eigenvector in the i^{th} epoch. The expected time to locate the next class is the inverse of the probability $P(i \rightarrow i + 1)$ of achieving this, which under our assumptions is given by:

$$P(i \rightarrow i + 1) = \sum_{j/q=i} e_j^i \sum_{k/q=i+1} M_{kj} \quad (11)$$

where M is the mutation matrix defined in (10)

The assumptions used will have different effects on the quality of our estimates.

- The first assumption effectively assumes that as soon as a member of the next fittest class is discovered, the whole population jumps to a state close to the next attractor. This ignores both the time taken for the population to move to the next epoch, and also the possibility that the new best individual will be lost by stochastic selection errors before it has a chance to propagate. The net effect is to cause an underestimation of the hitting time. In [16] it was found that with static mutation rates, the estimates of solution time were consistently below the observations for these reasons.
- The effect of assumption two is harder to quantify. We know that the populations corresponding to the v^i will have high proportions of individuals with low mutation rates attached. Thus one might expect that they will be ‘less explorative than a more uniform distribution of mutation rates and take longer to discover the next fitness class, leading to an overestimation of the hitting time.

- The principal effect of the third assumption (maximum entropy) is in the definition of A_i , where we effectively say that given we could be in *any* unaligned state and so each mutation has equal probability of aligning a block. Although there may be particular GA run where all unaligned blocks only require a single mutation to align them, this would get lost in the averaging process.
- The last of the assumptions above rules out the possibility that a member of class $i - 1$ gets mutated to class $i + 1$ in epoch i , which is possible since the steady-state population in epoch i may contain members of lower fitness class. We could allow for this situation by changing the limits in the first summation of (11) from a strict equality to $j/q \leq i$. However, we could permit a further relaxation by permitting transitions from an epoch to *any* higher epoch.

In this case, we can derive estimates for the hitting time by considering the epochs to be the states in a Markov chain in which the final epoch is an absorbing state and all others are transient. We calculate the entries in the associated transition matrix by changing the condition in the first summation as indicated above, and that in the second summation to generalise over states $j \in \{i + 1, i + 2, \dots, l\}$. The entries for the lower triangle are zero by assumption, and the leading diagonal entries are calculated so that the matrix is stochastic, i.e. $P_{ii} = 1.0 - \sum_{j>i} P_{ij}$. Using standard techniques, we can then obtain the mean absorption time from each of the transient states. Finally we calculate the fitness probability distribution for a randomly created individual and use this to weight the times to give a prediction of the mean hitting time for a randomly created population ³.

We calculated estimated hitting times for the exogenous algorithm on a Royal Road functions with $K \in \{4, 8\}$, and compared these to empirical observations for the algorithms and settings used above.

4.1 Exogenous behaviour

Figure 3 shows predicted hitting times according to both methods, and empirical results for a population size of 1000, for the exogenous scheme for the two higher values of K . Comparing the results from the two different sets of assumptions in Figure 3, we see that in both cases the second method gives lower predicted times, but for $K = 8$ the difference is less marked than for $K = 4$, i.e. making allowances for the possibility of “leapfrogging” over an epoch to the next but one, has less effect for high K . This arises from the exponential factor K in the alignment probability (4).

Broadly speaking the predictions show that as the innovation rate is increased so the time to solution decreases (up to a certain point). This prediction is borne out by the experimental results, although with some differences.

We note that for $z < 0.6$ [K4] or 0.7 [K8] the models consistently underestimate the hitting times suggesting that the errors introduced by assumption dominate.

For $0.4 < z < 0.8$ the models are far more accurate, and correctly predict that for $z \geq 0.52$ the times for $K = 8$ become less than those for $K = 4$. Given the errors in the predictions, this

³in practice the absorption times were observed to be similar from each state.

changeover point is predicted with surprising accuracy (error < 0.01). We ran experiments with a populations of 250 and noted the same effect.

For higher values of z we see that the model consistently overestimates the hitting time, quite dramatically so for $K = 4$, less so for $K = 8$. The hitting times observed continue to reduce with increasing z , whereas the models predict a sharp upturn for higher z , especially with $K = 4$.

These errors arise from either the whole model being invalid (which we can discount from the accuracy of the earlier predictions), or one of the assumptions in our calculation of the hitting times being invalidated. From the discussion above we turn to the second assumption as the source of errors for the high z regime. This appears to draw into question the epochal nature of search assumed for Self-Adaptive algorithms with extreme perturbation probabilities, suggesting rather that the next fitness level is discovered before the population has converged to the eigenvector v^i . Clearly further investigation is

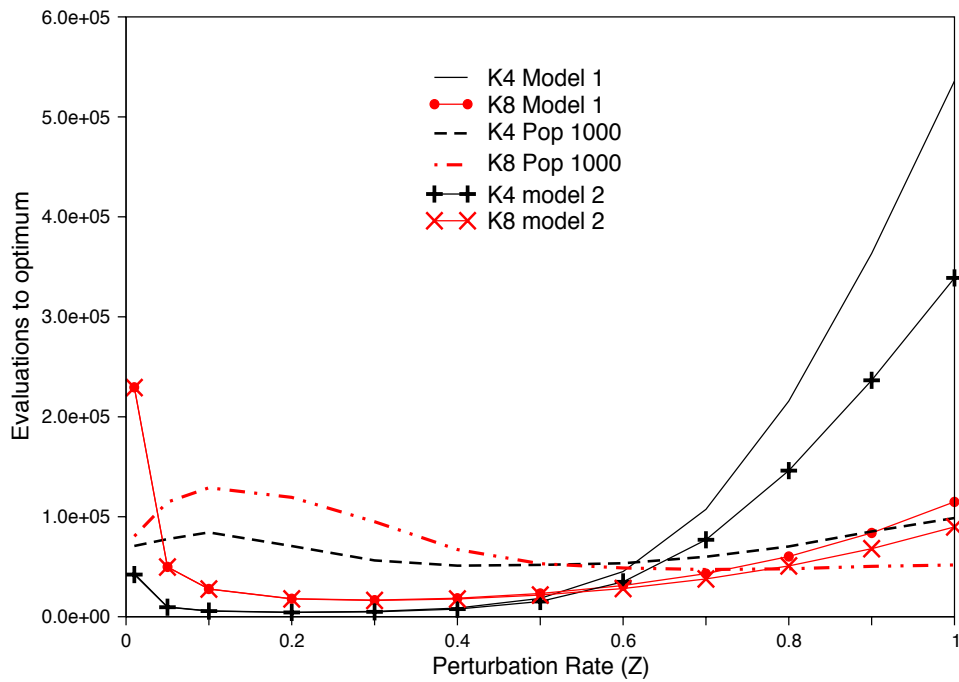


Figure 3 Mean time to solution as function of perturbation rate for Royal Road functions with $K = 4$ and 8 . Figures in braces show number of successful runs where this is less than 50.

4.2 Endogenous behaviour

It was found empirically for $K > 1$ that the exogenous scheme found the global optimum far more reliably than endogenous scheme, and that both systems were preferable to the static versions. For $K = 8$ the endogenous scheme only found the optimum in 14 of the 50 runs, but when it did so, it took slightly less time than the exogenous scheme with the “best” perturbation rate.

We cannot directly apply the analysis above to the endogenous strategy, since we know that the fixed points of the system corresponding to fitness classes below the optimum contain only individuals of the same fitness with zero probabilities of mutation and hence of escape to a higher fitness class. Overall it is clear that the endogenous scheme will get “trapped” in an epoch if random mutations do not discover a superior fitness class before the population converges to the fixed point.

We proceed as follows: first we construct an estimate for the time taken between discovering a new fitness class, and converging to the corresponding fixed point. Considering a population of iso-fitness with transitions between mutation classes, we used a Markov chain approach to derive estimates for the time to convergence onto the zero-mutation class for a single individual. We constructed a transition matrix according to (1) and used it to calculate the expected absorption time to zero mutation from each different mutation state. These varied between 465.8 and 616.7 iterations.

Next we construct a lower bound for the time to find epoch $i + 1$ from epoch i . The first discovered member of fitness class i can belong to any non-zero mutation class and under the conditions here, will have a high selective advantage. Thus we consider that the population has not converged to zero mutation, but rather that each mutation class is equally represented. Examining the predicted numbers of generations between epochs for the exogenous algorithm with equally represented mutation classes ($z = 1$), we see values of $\{21, 51, 152, 579, 2992, 24141, 423033\}$ for $K = 4$ and $\{696, 5410, 106985\}$ for $K = 8$.

From this we can see that, even while a range of mutation rates are present, the expected transition times to higher epochs are an order of magnitude longer than the expected times for absorption to the zero mutation class. These observations help to explain why the endogenous strategy is observed to prematurely converge onto sub-optimal solutions with a loss of the ability to explore further. They also explain why the algorithm is fast when it is successful, since escaping from the traps at each epoch requires that the discovery of the next fitness level before the ability to mutate is lost.

5 Conclusions

We have identified two contrasting paradigms for the perturbation of mutation rates in Genetic Algorithms with Self-Adaptive Mutation. We have derived dynamical systems models of these paradigms and shown how these can be used to predict and explain the behaviour of “real” GAs with finite populations on two different classes of functions.

Calculating the eigenvectors and associated eigenvalues of the models yields predictions of the steady-state mean fitness of the algorithms. It also gives insight into where, and

why GAs with finite populations get “trapped” in sub-optimal solutions, a result of the presence of “attractors” of the system lying on or very close to the simplex.

We showed that for the endogenous strategy perturbation scheme, with entropic barriers to the discovery of the global optimum, the presence of fixed points corresponding to populations where every member has the same fitness and zero probability of mutation led to the system frequently getting trapped as it lost the ability to explore the iso-fitness plateau before it discovered higher fitness levels. Whilst the exogenous perturbation scheme did not suffer from the effects of attractors on or near the simplex for entropic barriers, it did on deceptive problems where there are fitness barriers to the discovery of the global optimum.

For the exogenous scheme we have derived estimates for the time to solution, which show that as the order of the problem increases, so higher perturbation rates reduce this time, even though the mean fitness decreases. Experimental results confirmed some of these findings, but lead us to question that search occurs in an epochal fashion for exogenous strategies, and this clearly requires further investigations.

Although we have considered the two best known approaches, it would be interesting to compare variants, for example an endogenous scheme which used the encoded probability to decide whether to choose a new value from a uniform distribution, or the use of a Gaussian probability distribution with an exogenous scheme. This will be examined in future work.

In conclusion, it would appear that (subject to the usual provisos about the choice of performance metric) the exogenous scheme offers a more robust problem solving mechanism than the endogenous scheme, despite the undoubted attractions of the latter’s algorithmic elegance.

References

- [1] Thomas Back. Self adaptation in genetic algorithms. In F. Varela and P. Bourguine, editors, *Towards a Practice on Autonomous Systems: Proceedings of the First European Conference on Artificial Life*, pages 263–271. MIT Press, 1992.
- [2] Kalyanmoy Deb and David E. Goldberg. Analyzing deception in trap functions. In L. Darrell Whitley, editor, *Foundations of Genetic Algorithms 2*, pages 93–108, 1992.
- [3] David Fogel. *Evolving Artificial Intelligence*. PhD thesis, University of California, 1992.
- [4] S. Forrest and M. Mitchell. Relative building block fitness and the building block hypothesis. In L. Darrell Whitley, editor, *Foundations of Genetic Algorithms 2*, pages 109–126. Morgan Kaufmann, 1992.
- [5] Matt Glickman and Katie Sycara. Reasons for premature convergence of self-adaptating mutation rates. In *Proceedings of the 2000 Congress on Evolutionary Computation CEC00*, pages 62–69, 2000.
- [6] K. Kaneko and T. Ikegami. Homeochaos: Dynamic stability of a symbiotic network with population dynamics and evolving mutation rates. *Physica-D*, 56:406–429, 1992.
- [7] Natalio Krasnogor and Jim Smith. Emergence of profitable search strategies based

- on a simple inheritance mechanism. In Lee Spector, Erik Goodman, Annie Wu, William B. Langdon, Hans-Michal Voigt, M. Gen, S. Sen, Marco Dorigo, S. Pezeshk, M. Garzon, and E. Burke, editors, *Proceedings of the Genetic and Evolutionary Computation Conference, GECCO-2001*, pages 432–439. Morgan Kaufmann, 2001.
- [8] Ingo Rechenberg. *Evolutionsstrategie: Optimierung technischer Systeme nach Prinzipien der biologischen Evolution*. Frommann-Holzboog, Stuttgart, 1973.
- [9] Jonathan Rowe. Population fixed-points for functions of unitation. In W. Banzhaf and C. Reeves, editors, *Foundations of Genetic Algorithms 5*, pages 69–84. Morgan Kaufmann, 1999.
- [10] G. Rudolph. Self-adaptive mutations may lead to premature convergence. *IEEE Transactions on Evolutionary Computation*, 5:410–414, August 2001.
- [11] Hans Paul Schwefel. *Numerical Optimisation of Computer Models*. John Wiley and Sons, New York, 1981.
- [12] Jim Smith. Modelling gas with self-adaptive mutation rates. In Lee Spector, Erik Goodman, Annie Wu, William B. Langdon, Hans-Michal Voigt, M. Gen, S. Sen, Marco Dorigo, S. Pezeshk, M. Garzon, and E. Burke, editors, *Proceedings of the Genetic and Evolutionary Computation Conference, GECCO-2001*, pages 599–606. Morgan Kaufmann, 2001.
- [13] Jim Smith and T.C. Fogarty. Self adaptation of mutation rates in a steady state genetic algorithm. In *Proceedings of the Third IEEE International Conference on Evolutionary Computing*, pages 318–323. IEEE Press, 1996.
- [14] C.R. Stephens, I. Garcia Olmedo, J. Moro Vargas, and H. Waelbroeck. Self-adaptation in evolving systems. *Artificial life*, 4:183–201, 1998.
- [15] Christopher Stone and Jim Smith. Strategy parameter variety in self-adaption. In *Proceedings of Gecco 2002*, pages XX–xX, 2002.
- [16] Erik van Nimwegen and James Crutchfield. Optimizing epochal evolutionary search: Population-size dependent theory. *Machine Learning*, 45(1):77–114, 2001.
- [17] Erik van Nimwegen, James P. Crutchfield, and Melanie Mitchell. Statistical dynamics of the royal road genetic algorithm. *Theoretical Computer Science*, 229:41–102, 1999.
- [18] Erik Jan van Nimwegen. *The Statistical Dynamics of Epochal Evolution*. PhD thesis, Utrecht University and Santa fe Institute, 1999.
- [19] Michael D. Vose. *The Simple Genetic Algorithm*. MIT Press, 1999.