The Puzzle of the Persistent Question Marks: A Case Study of Genetic Drift

Inman Harvey

CSRP 278, April 1993

Cognitive Science Research Paper

Serial No. CSRP 278

The University of Sussex
School of Cognitive and Computing Sciences
Falmer
BRIGHTON BN1 9QH
England, U.K.

This paper will appear in:

S. Forrest, (ed.),
Genetic Algorithms: Proceedings of the Fifth International Conference (GA93),
Morgan Kaufmann, San Mateo, CA.

he Puzzle of the Persistent Question Marks: A Case Study of Genetic Drift

Inman Harvey*
School of Cognitive and Computing Sciences
University of Sussex
Brighton BN1 9QH, England

Abstract

In an important paper, Hinton and Nowlan (1987) demonstrate the Baldwin effect in a simple Genetic Algorithm. The ability of the phenotype to adapt, coupled with the evolutionary process, allows behavioural goals to become over time genetically specified; this seems Lamarckian but is not. In that paper, as a subsidiary point, the slowness of fixation of the last few goals is commented on, and a later paper by Belew (1989) attempts an analysis. In this paper I show that genetic drift is the explanation for this slowness phenomenon. Using a diffusion equation approach, I give an analysis of genetic drift for genetic algorithms, where it is too often ignored. Critical relationships between mutation rate, population size, and forces of selection are given which decide whether genetic drift will be of significance or not.

1 Introduction

In an important and elegant paper, Hinton and Nowlan (1987) demonstrate with a deliberately simple example the Baldwin effect, wherein the ability of a phenotype to adapt in its lifetime (ability to 'learn') alters the fitness landscape of the corresponding genotype. This has the consequence that selection within a population moves the genotypes towards the region where the adaptations, that were originally made in the lifetime of the phenotypes, are genetically fixed. This has the appearance of Lamarckism, but is not so, as there has been no direct flow of information from the adapted phenotype to the genotype.

The model chosen as an example uses genotypes with a number of genes that can be specified as incorrect, correct, or open to adaptation during the lifetime of the phenotype. The evaluation function only favours those phenotypes that, within a finite lifetime, find a perfect solution through a combination of 'correct' genes, and 'adaptive' genes which successfully adapt. It is demonstrated that with the application of a standard genetic algorithm (GA) to the population as specified, the number of incorrect alleles on the genotype rapidly decreases to zero; the number of correct alleles increases at first rapidly and then slows down; the number of undecided (adaptive) alleles decreases slowly. If the same experiment is tried out only with correct and incorrect genes, and no adaptive ones, then the 'needle in a haystack' nature of the single perfect solution means that only random search works, and takes an unreasonably long time.

The main thrust of Hinton and Nowlan's paper is endorsed here, but a subsidiary matter that is mentioned as an aside there is taken up as the main point for investigation here in this paper:

> One interesting feature of [the figure] is that there is very little selective pressure in favor of genetically specifying the last few potential connections, because a few learning trials is almost always sufficient to learn the correct settings of just a few switches.

The figure in question indicates that there could be an asymptote at a relative frequency of about 0.45 below which the number of undecided alleles will not fall.

My own re-implementation of the model usually shows an asymptote at between 0.05 and 0.2. A typical run is shown in figure 1, showing the dramatic changes in the first 50 generations, and the longer term behaviour over 500 generations. The variations between runs is indicated in table 1, showing the values at the end of 20 runs of 500 generations each. The re-implementation by Belew (1989) shows 'an almost steady-state' at about 0.3. He asserts that the curve is 'in fact asymptotically approaching ...0.0'. This I will demonstrate to be false, in the general case; the analysis of what is really happening shows that the combination of genetic drift and the hitch-hiking effect so completely swamps the selective pressures that some of the genes are com-

*Email: inmanh@cogs.susx.ac.uk

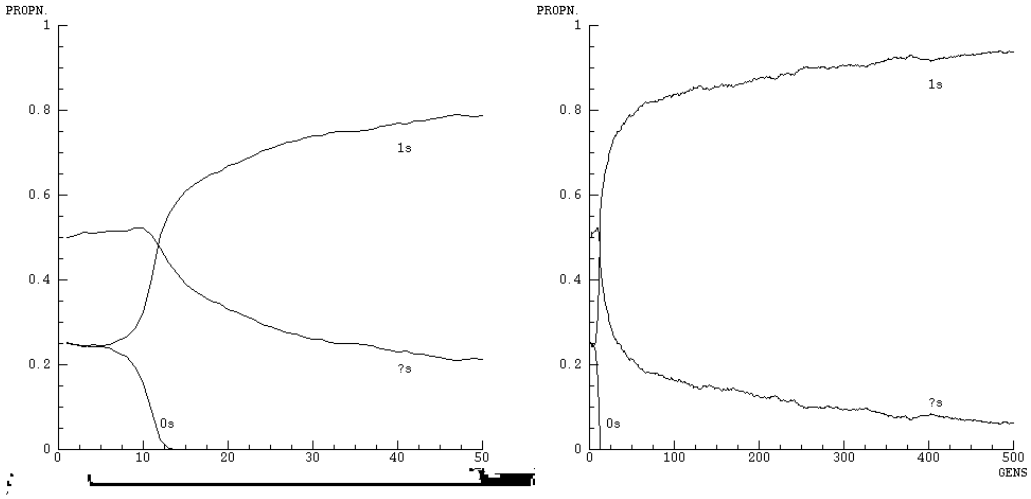

Figure 1: The proportions of incorrect, correct, and undecided (adaptive) alleles (**0**s, **1**s, **?**s) in the whole population, against generations. On the left, the first 50 generations of a run, and on the right the same continued for 500 generations.

pletely converged to the undecided value, rather than the 'correct' one.

Table 1: The final proportions of undecided alleles after 20 runs each of 500 generations, with no. of loci converging or converged on **?**. 4 runs† have in fact completely converged at all 20 loci, only one run‡ does not yet have a locus with **?** fixed.

Propn of ?s at 500 gens.	Loci having >50% ?s	Loci having 100% ?s	Propn of ?s at 500 gens.	Loci having >50% ?s	Loci having 100% ?s
0.063	1	1	0.108	2	2
0.109	2	2	0.093	2	1
0.082	1	1	0.150	3	3†
0.123	2	2	0.150	3	3†
0.118	2	2	0.112	2	2
0.074	1	1	0.100	2	2†
0.107	2	2	0.093	2	1
0.200	4	4†	0.121	2	1
0.134	3	2	0.115	2	2
0.092	2	0‡	0.115	1	1

bility of that member contributing to the reproductive pool for the next generation. In the early stages, virtually all the members will have the same minimum fitness. Something similar will happen also at the later stages, after the incorrect (**0**) alleles have been eliminated; virtually all members will have small q-values, and hence, because of the flatness of the curve for $F(q)$ at small q, nearly identical fitnesses. At both these stages there is very little selective pressure.

However, as Figure 1 indicates, typically around generations 5 to 15 successful members emerge with a fitness nearly 20 times as great as that of the original random members. This enormous selective differential operates near-exponentially for a few generations, giving the sharp swings indicated in the figure. If the

q	F(q)	q	F(q)
0	20.000		
1	19.962	11	4.965
2	19.924	12	3.140
3	19.848	13	2.113
4	19.696	14	1.568
5	19.392	15	1.287
6	18.784	16	1.144
7	17.569	17	1.072
8	15.233	18	1.036
9	11.649	19	1.018
10	7.868	20	1.009

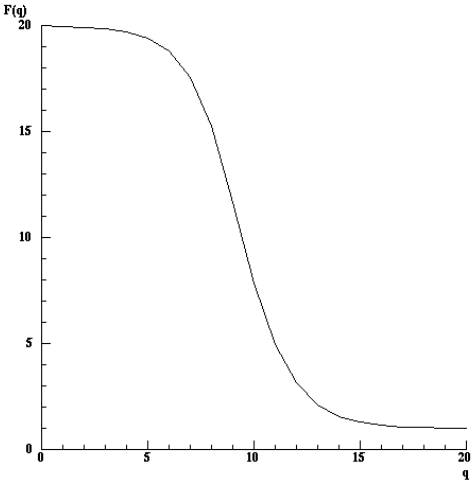

Figure 2: The expected fitness $F(q)$ of a gene with q undecided alleles and $(20-q)$ correct ones. For $q = 0$ it

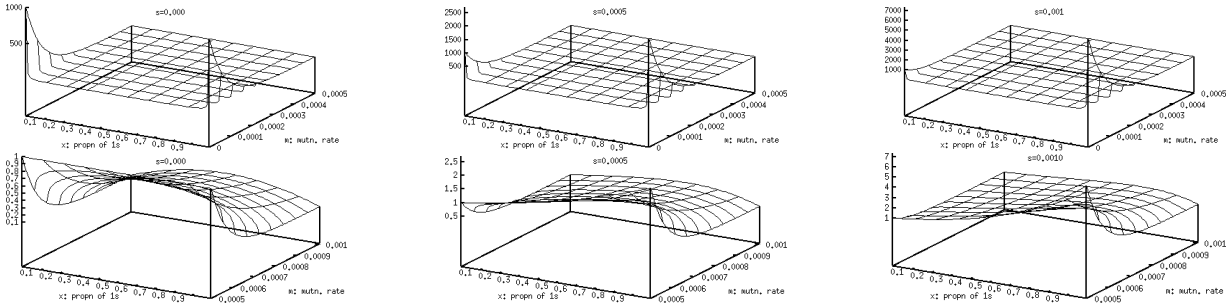

Figure 3: Equilibrium distributions, varying m for particular values of s. The horizontal scale is the proportion x of the allele being selected for, in the range $x = 0.001$ to 0.999. the vertical scale varies from graph to graph, as the constant c in eqn. 11 has here been set to 1; whereas it should normalise the graph so that the area underneath is unity. Hence for the U-shaped curves, only the general shape is indicative.

$$\Delta x_{sel} = x\left(\frac{f_1}{\overline{f}} - 1\right) = \frac{sx(1-x)}{1+sx} \qquad (7)$$

Using the population size N, we now convert Δx to a new time scale where N generations equals one unit of time.

$$\begin{aligned} M(x) &= N(\Delta x_{mut} + \Delta x_{sel}) \\ &= mN(1-2x) + \frac{sNx(1-x)}{1+sx} \end{aligned} \qquad (8)$$

On calculating $V(x)$ we use the fact that the variance of Δx over one generation is $x(1-x)/N$. Converting to the same time scale as above we have

$$V(x) = N\frac{x(1-x)}{N} = x(1-x) \qquad (9)$$

Substituting (8) and (9) into (4) we have the following (the constants on integration can be assimilated into the normalizing constant c):

$$\hat{\rho}(x) = \frac{c}{x(1-x)}\exp\left(W(x)\right)$$

where we defin0.24Tf11.03980Tdertts

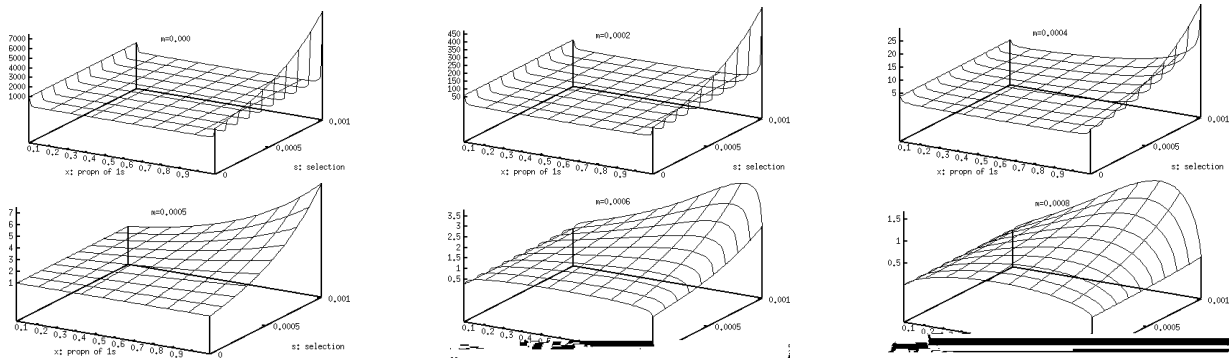

Figure 4: Equilibrium distributions, varying s for particular values of m. See caption to figure 3

of the relative proportions of the converged population that settle at $x = 0$ or $x = 1$, the numerator should be considered for these values; although not much re-

A Appendix: The Hinton & Nowlan model

A.1 Expected fitness of potential winner

To calculate the expected fitness of