

An Analysis of the MAX Problem in Genetic Programming

W. B. Langdon

School of Computer Science,
The University of Birmingham,
Birmingham B15 2TT, UK
W.B.Langdon@cs.bham.ac.uk
<http://www.cs.bham.ac.uk/~wbl>

R. Poli

School of Computer Science,
The University of Birmingham,
Birmingham B15 2TT, UK
R.Poli@cs.bham.ac.uk
<http://www.cs.bham.ac.uk/~rmp>

Abstract

We present a detailed analysis of the evolution of genetic programming (GP) populations using the problem of finding a program which returns the maximum possible value for a given terminal and function set and a depth limit on the program tree (known as the MAX problem). We confirm the basic message of [Gathercole and Ross, 1996] that crossover together with program size restrictions can be responsible for premature convergence to a sub-optimal solution. We show that this can happen even when the population retains a high level of variety and show that in many cases evolution from the sub-optimal solution to the solution is possible if sufficient time is allowed. In both cases theoretical models are presented and compared with actual runs.

Price's Covariance and Selection Theorem is experimentally tested on GP populations. It is shown to hold only in some cases, in others program size restrictions cause important deviation from its predictions.

1 Introduction

While genetic programming (GP) has been demonstrated on a wide and increasing range of applications comparatively little work has been devoted to a theory of GP. However such a theory is of great importance to the field as a whole. Work so far has been principally split between extending the Schema Theorem [Holland, 1992] from Genetic Algorithms to GP [Koza, 1992; O'Reilly and Oppacher, 1995; Whigham, 1995; Poli and Langdon, 1997], studying benchmark problems (e.g. Royal Trees [Tackett, 1995; Punch *et al.*, 1996]) and applying population genetics to GP [Altenberg, 1994; Altenberg, 1995]. In this paper we combine the later

two approaches by applying Price's Theorem from theoretical biology, plus other detailed analysis, to the MAX problem.

[Gathercole and Ross, 1996] introduce the MAX problem to GP to highlight deficiencies in the standard GP crossover operator which result from the practical requirement to limit the size of evolved programs. They concentrate on the case where program trees are restricted to a maximum depth but suggest similar effects will be seen when programs are restricted to a maximum number of nodes.

The MAX problem has known optimal solutions which are composed of regularly arranged sub trees or building blocks. Despite this GP finds solving larger versions of the MAX problem difficult.

In this paper we extend [Gathercole and Ross, 1996] by considering bigger trees, different selection pressures, different initialisations of the population, measuring the frequency with which the depth limit affects individual crossovers, measuring population variety and the number of steps required to solve the MAX problem. Qualitative models of crossover and population variety are presented and compared to measurement. We give an improved explanation for the premature convergence noted by [Gathercole and Ross, 1996] and this leads to the realisation that there are two separate reasons why GP finds the MAX problem hard. Firstly the tendency for GP populations to converge in the first few generations to suboptimal solutions from which they can never escape. Secondly convergence to suboptima from which escape can only be made by slow search similar to randomised hill climbing.

Section 2 describes the various MAX problems used in these experiments. Section 3 describes the GP used in our experiments, Section 4 gives the results obtained and provides a detailed analysis and comparison with [Gathercole and Ross, 1996]. Section 5 shows typically the population retains a high level of diversity and presents models of the variety in the initial population and its subsequent evolution. Section 6 considers the role of selection pressure and shows reduced perfor-

mance with small tournament size. Section 7 describes Price’s Theorem, presents experimental evidence on its applicability to GP and shows it can be used to predict the behaviour of subsequent generations of the MAX problem.

2 The MAX Problem

In the MAX problem “the task is to find the program which returns the largest value for a given terminal and function set and with a depth limit, D , where the root node counts as depth 0” [Gathercole and Ross, 1996, page 291]. In this paper we use the function set $\{ +, \times \}$ and there is one terminal 0.5. For a tree to produce the largest value $+$ nodes must be used with 0.5 to assemble subtrees with the value of 2.0. These can then be connected via either $+$ or \times to give 4.0. Finally the rest of the tree needs to be composed only of \times nodes to yield the maximum value of 4^{2^D-3} . It is important to note that this artificial problem forces every component of the evolved programs to contribute to the programs’ fitness and no part of any program can escape from the effect of selection. I.e. there can be no introns.

3 GP Parameters

Our GP system was set up to be the same as given in [Gathercole and Ross, 1996]. The details are given in Table 1, parameters not shown are as [Koza, 1994, page 655]. On each version of the problem 50 independent runs were conducted (with 6 depths and two types of initialisation in the three bigger problems and 7 tournament sizes this makes a total of 3150 runs).

In this paper two sets of experiments are presented. In the first the usual ramped half-and-half method [Koza, 1992, page 93] which creates random trees with depths between 2 and 6 (i.e. $D = 1 \dots 5$) was used. (However the initial trees obeyed the problem specific height restriction). A second set of runs (with $D > 5$) were made with the maximum tree height in the initial population identical to the problem specific limit. (Unless otherwise stated the discussion refers to this second set of experiments). Gathercole and Ross created their initial populations “with no constraint on tree size (other than the overall size limit), i.e. not ramped” [Gathercole, 1997].

4 Results

Figure 1 shows the mean number of generations taken by the GP to solve the MAX problem in the successful runs. The percentage of unsuccessful runs is given in Figure 2. These results are similar to those in [Gathercole and Ross, 1996] except [Gathercole and Ross, 1996, Figures 3 and 4] do not contain data for depths of 7 or 8.

Table 1: MAX Problem

Objective:	Find a program that returns the largest value
Primitives:	$+, \times, 0.5$
Max depth	3 ... 8 (NB root node is depth 0)
Init depth	5 or Max depth
Fitness:	Value of tree
Selection:	Tournament group size of 2 to 8, generational plus elitism.
Parameters	Pop = 200, G = 500, 99.5% crossover, no mutation. Crossover points selected uniformly between nodes.




Figure 1: Mean number of generations need by the successful runs. Error bars indicate standard error.

4.1 Impact of Depth Restriction on Crossover

The size of programs within the population grows quickly until almost all individuals are full trees. Due to the depth restriction crossover cannot make the trees any bigger and so crossover fragments are either the same size as the code they are replacing or smaller. In the case of full trees this means crossover can move code at the same level in the trees or higher but not lower.

In the first run with $D = 8$ a third of crossovers are rejected because the subtree to be inserted would cause the offspring to violate the maximum depth. In these cases the roles of the two parents are reversed and a shorter subtree is inserted instead. Of the remaining two thirds half (i.e. $\frac{1}{3}$ of the total) result in replacing a subtree with one of the same height, so in total two thirds of crossovers replace a subtree with a shorter one. Only 0.74% of crossovers resulted in a taller subtree replacing a shorter one. These occur throughout most of the run, half occurring by generation 85 and the last in generation 441 (the mean generation is 105).




Figure 2: Percentage of runs ending in failure

4.2 Trapping by Suboptimal Solutions

[Gathercole and Ross, 1996, page 295] suggest the reason why GP finds the MAX problem hard is that the population quickly finds suboptimal solutions which contain + nodes near the root of the tree. To improve on such solutions crossover must replace them with × nodes. They suggest that “once the trees have reached the depth limit, the only way the higher levels are affected is through the promotion of lower level subtrees, which contain no × nodes, and the movement of subtrees within the same level”. However the lack of × nodes in lower levels was not observed in our populations. For example at the end of all but one of the 50 runs with $D = 8$ the population contains thousands of × nodes at levels 3 and 4, even in runs where levels 0, 1 or 2 contain large numbers of + nodes.

Crossover readily moves such nodes to higher levels. To explain why the suboptimal + nodes are able to remain in the population we need to consider the fitness of the resulting offspring. Suppose both parents are near optimal (i.e. they consist mainly of +, × and 0.5 at the right level of the tree). The value of each subtree with a × as its root node is $4^{2^{D_x-3}}$ (where D_x is its depth) while if rooted with a + node it is $2 \times 4^{2^{D_+-4}}$. If an otherwise optimal subtree with a + as its root node is replaced with an optimal subtree from a lower level (i.e. $D_+ > D_x$) then the value of this part of the offspring is reduced. For example if an otherwise optimal subtree of depth 5 rooted with a + node (value 32) is replaced by an optimal subtree which is one level shorter (value 16). Then given the other subtrees at the same level are also near optimal (i.e. have values ≥ 1), the value (and hence fitness) of the offspring as a whole will be less than that of its parent by a factor of at least $\frac{1}{2}$. Therefore the offspring is unlikely to have children in the next generation. In contrast the $\frac{1}{3}$ of crossovers that exchange

subtrees at the same level may produce offspring with same fitness as their parents and thus are more likely to have children themselves. NB programs with a few + nodes near the root can readily reproduce copies of themselves but their children where a + is replaced with a × are unlikely to survive and so the GP can remain trapped for long periods.

4.3 Modelling the Rate of Improvement

[Gathercole and Ross, 1996, page 295] says “if there are no × nodes in any tree at a particular high level, it [is] now impossible for crossover to introduce a × node to this level; the population has converged to being duplicates of a sub-optimal tree, and no further improvement is possible.” However (taking the example with $D = 8$ again) in only 7 cases does the population at the end of the run contain zero × nodes at any level from 0 to 3, i.e. in 22 of the runs that failed the population did contain × nodes in all of the higher levels. (The impossibility of the 7 runs succeeding was confirmed experimentally by running them again but this time to 5,000 generations. None succeeded, however in one case the population did improve finding better suboptimal solutions by replacing + nodes with × nodes where these were available at the same level).

The reason why the 22 runs failed (and why the successful runs took so long) despite having × nodes available is in part due to the low level of crossover activity near the root of the trees. That is crossover is able to improve suboptimal trees but it takes a long time. We can estimate how long with the following model.

Assume the population has converged to a suboptimal tree containing n_{+d} + and $n_{\times d}$ × nodes at level d ($d < D - 3$) (so $n_{+d} + n_{\times d} = 2^d$). Each individual will be a full (or nearly full) binary tree of height D and so contain $2^{D+1} - 1$ nodes. Therefore the chance of selecting one of the + nodes at level d to be the crossover point is

$$\frac{n_{+d}}{2^{D+1} - 1}$$

and the chance of replacing it with a × node from the same level is

$$\frac{n_{\times d}}{2^{D+1} - 1} \quad (1)$$

Firstly we note that the chance of improvement is much higher with large d . I.e. the nearer to the root the + node is the harder it will be for crossover to shift (in the case of the root node, it is impossible as there can be no × nodes also at the root). Secondly improvement is easiest when the number of + nodes is equal to the number of × nodes and the last + node at each level is the most difficult to remove. Also note that if the number of × nodes is small it will be difficult to increase it.

As replacing a + node at any of the higher levels will produce an improvement, by only considering crossovers which find improved solutions by moving subtrees at the same level as the + nodes we can form a lower bound on the overall chance of the offspring being an improvement:

$$\sum_{d=0}^{D-4} \frac{n_{+d}n_{\times d}}{(2^{D+1}-1)^2} = \sum_{d=0}^{D-4} \frac{n_{+d}(2^d - n_{+d})}{(2^{D+1}-1)^2} \quad (2)$$

Crossover closer to the root can also find improved solutions by replacing trees containing + nodes with others where they are replaced by \times nodes, including where the inserted subtree contains + nodes further from the root. In general where there are + nodes at different levels deciding how many crossover points will yield an improved solution becomes complex. However if there are few incorrectly positioned + nodes we need only consider crossovers where the crossover points in both parents are at the same level. For the special case of a single + node at level d then crossover at any node connecting it to the root with a different node at the same level will yield the optimal solution. There are

$$\sum_{i=1}^d 2^i - 1 = 2^{d+1} - d - 2 \quad (3)$$

such crossover point pairs. Note in this case the exact figure is about twice the approximation of $n_{\times d}$ used in (1) and (2).

As each crossover is independent, the number required to replace a + node with a \times node has an exponential distribution. We can estimate the mean and standard deviation using (2) as

$$\text{crossovers to improvement} \leq \frac{(2^{D+1}-1)^2}{\sum_{d=0}^{D-4} n_{+d}(2^d - n_{+d})}$$

Thus the expected number of generations until the next improved solution is found in a population which has converged towards a good but sub-optimal solution is

$$\text{gens to improvement} \leq \frac{(2^{D+1}-1)^2}{p p_c \sum_{d=0}^{D-4} n_{+d}(2^d - n_{+d})} \quad (4)$$

where p is the population size and p_c is the crossover probability (NB no mutation).

The predictions of (3) were tested by measuring the number of generations required to replace the last misplaced + node with $D = 8$ (this required extending 14 runs up to 5,000 generations). Figure 3 shows good agreement between prediction and measurement. At every depth the mean lies within 1.4 standard errors of the predicted value.

To test further the model we looked at multiple + nodes. Since this is more complex we looked at only




Figure 3: Number of generations to displace last misplaced + ($D = 8$). Error bars indicate standard error.

Table 2: Generations at which improved solutions were found in an extended run ($D = 8$). The estimated number of crossover pairs which give improved solutions and the corresponding expected number of generations between improved solutions are given, for both the simple $\sum n_{+}(2^d - n_{+})$ model and for the detailed counting model.

depth	Gen	ΔGen	$\sum n_{+} \dots \text{Pred}$	Count	Pred
4	523	23	55	73	24
4	542	19	46	65	29
4	569	27	35	51	37
3	584	15	21	34	62
4	611	27	15	26	87

one run. We selected the first unsuccessful run for $D = 8$ after the run used in Section 4.1. In this run the population had converged towards a tree with one + node at level 3 and four at level 4. It was run on past generation 500 and the generations where improved solutions were found was noted. The results are shown in Table 2. In the last but one column of Table 2 the exact count of crossovers which would generate an improved solution is given, the right most column gives the corresponding expected number of generations to find an improved solution. There is reasonable agreement with the actual number of generations, recorded in column 3, whereas the simpler model (4), columns 4 and 5, consistently overestimates the time required to find the next solution.

4.4 No. of Steps to Climb the Hill

The number of improved solutions found before the optimal solution is found is plotted in Figure 4. In success-




Figure 4: Mean number of improved solutions found before optimal solution in successful runs. Error bars show standard error for tournament size of 6.

ful runs, it takes about as many improvement steps as there are functions in the optimal solution (i.e. $2^D - 1$). This supports the notion that, in this problem, GP finds solutions one step at a time and is not benefiting from *implicit parallelism* expected where complete solutions can be assembled from building blocks. From (4) we see each step takes $O(2^{2D})$ generations, so we expect the total number of generations required, in successful runs, to grow as $O(2^{2D})$, i.e. parallel to the straight line in Figure 1. Figure 1 shows good agreement until the limit of 500 generations per run acts to cut short runs which would otherwise have eventually found a solution.

5 Variety

5.1 Variety in the Initial Population

The rightmost column of Table 3 contains the average number of different programs in the initial populations. It highlights another potential problem with GP when it is applied to the MAX problem: in contrast with normal GP practice the initial populations contains a large number of duplicates (see also Figures 5 and 6). This is inherent in the ramped half-and-half method when used in a problem with such small function and terminal sets. Recall that with the ramped half-and-half method random trees are created with a maximum depth evenly distributed between 1 (counting the root node as zero) and the maximum depth allowed in the initial population. Half are created as full trees of this maximum depth and half are created using a grow method which creates trees of different shapes. ([Iba, 1996] and [Bohm and Geyer-Schulz, 1996] discuss alternative means of generating random trees for use as the initial population in GP runs.)

With $D = 8$ about $\frac{1}{8}$ (i.e. $\frac{1}{D}$) of trees will be created

Table 3: Variety in initial populations, Predicted v. Mean of 50 runs (ramped half-and-half)

D	$\frac{1}{8} + \frac{7}{4D}$	\times popsize	Predicted	Mean
3	.7083	141.67	58.33	66.3
4	.5625	112.5	87.5	94.1
5	.475	95	105	111.2
6	.4167	83.33	116.67	120.6
7	.375	75	125	128.9
8	.3438	68.75	131.25	136.4
∞	.125	25	175	

with a maximum depth of 1 and therefore will contain a single function. I.e. about 25 trees will consist of a single $+$ node and two 0.5 terminals or a \times node and two 0.5 terminals, and therefore they are very likely to be identical to at least one other member of the population. $\frac{1}{D}$ of trees are created with a maximum depth of 2 and so will contain no more than two or three functions. As there are 16 possible trees with two or three functions it is likely most trees with a depth of two will not be unique. With larger trees the number of possible trees grows rapidly and so most will be unique. $\frac{1}{4}$ of trees generated by the grow method will not grow either branch from the root and so a further $\approx \frac{1}{2} \frac{D-2}{D} \frac{1}{4}$ (i.e. 18.75) trees will be very small and so unlikely to be unique.

In summary we expect at least $\frac{1}{D} + \frac{1}{D} + \frac{1}{8} \frac{D-2}{D} = \frac{1}{8} + \frac{7}{4D}$ duplicates in the initial population (unless the initial population is rather less than 200). Table 3 shows this close to the mean number of duplicate individuals found in 50 initial populations. In problems with larger function or terminal sets the chance of randomly generating duplicates will be much smaller, but a high proportion of small programs can still be expected.

5.2 Evolution of Variety

While the populations converge to the extent that most of the trees in the population are similar to each other, the population does not converge to the extent that all of the population are identical. Where the maximum depth is large, on average after 250 generations, the number of different programs in the populations is about two thirds of the total population size (variety ≈ 133 , cf. Figure 5). I.e. one third of the population is composed of programs which are identical to at least one other in the population. Figure 6 shows in the case of smaller trees population variety is lower still.

5.3 Modelling Variety

Recall from Section 4.1 that about two thirds of children are of a different length to their host parent and




Figure 5: The number of different trees in the population (variety) and best fitness in the first run and mean of 50 runs ($D = 8$)




Figure 6: Number of different trees, mean of 50 runs

about one third are the same. With a relatively small population it seems reasonable to assume that children which are different from their parent are also likely to be different from every other child in the population. I.e. $\frac{2}{3}$ of children are likely to be unique. On the other hand, due to the convergence of the population, children produced by the $\frac{1}{3}$ of crossovers which swap subtrees at the same level are likely to be the same as one parent. Therefore they are likely to be the same as another child of the same parent or an identical individual in the previous generation. Therefore we expect the population to be composed of one third copies and two thirds unique individuals. This is approximately the case for $D = 8$.

The lower variety measured for $D < 8$ can be explained if one considers the breeding population to be composed mainly of copies of a full binary tree with only one node type at each level. The number of different trees that can be produced by crossover between

Table 4: Number of different trees that can be produced by crossing over two identical full binary trees, where each level contains only one type of node

D	3	4	5	6	7	8
number	12	27	58	121	248	503

such trees is limited (see Table 4). In the cases of $D = 4$ and 5 variety is on average close to the figure given Table 4 while for $D = 6$ and 7 the $\frac{2}{3}$ limit is also important.

6 Selection Pressure

In the MAX problem, with high selection pressure, we see a rapid convergence of the population towards solutions that are beneficial in the first few generations. However early solutions may not readily evolve to acceptable solutions. We suspect this is common to many GP problems. In the case of the MAX problem early high fitness individuals may use $+$ functions to form comparatively large values. Later generations can then yield still larger values by joining subtrees composed of $+$ nodes with \times functions. However, as [Gathercole and Ross, 1996] point out, the depth restriction and the mechanics of the crossover operator make the insertion of \times nodes near the root of large trees difficult. This section was motivated by the suggestion that the selection pressure was too high (i.e. the tournament size was too big) and this was responsible for driving the initial generations towards convergence too quickly.

The chance of an individual i being selected by a tournament is given by its rank r_i within the population of size p and the tournament size t by the formula [Blickle and Thiele, 1995; Langdon, 1996]

$$\frac{(r_i/p)^t - (r_{i-1}/p)^t}{r_i - r_{i-1}}$$

In a large and diverse population this can be approximated by

$$t(r_i/p)^{t-1} \quad (5)$$

The expected number of children produced by individuals of rank R or less is

$$p \sum_{i=1}^R \frac{(r_i/p)^t - (r_{i-1}/p)^t}{r_i - r_{i-1}} \approx p(R/p)^t$$

I.e. the worst $\approx \sqrt[t]{1/2}$ of the current generation produce half the children of the next generation, as do the best $1 - \sqrt[t]{1/2}$. With a tournament size of 6 this means on average the best 11% of the population (i.e. 22 individuals) produce half the next generation (i.e. 100 children). Given one third of these are likely to be identical to their




Figure 7: Mean number of generations need by the successful runs with tournament sizes from 2 to 8. Error bars indicate standard error.

parents, the elite members of the population can readily pass identical copies of themselves to the next and succeeding generations.

Figures 7 and 8 give the average number of generations in successful runs and proportion of runs that failed respectively with different tournament sizes. Each point represents the mean of 50 independent runs. For clarity only the data for the runs where the maximum height in the initial population is the same as that in later generations are displayed. We see performance is essentially independent of selection pressure except for tournament sizes of two and three.

With lower selection pressure ($t = 2$ or 3) GP perform worse on the MAX problem. This may be because the fitter members of the population have on average only two or three (i.e. t) children. Since only one third of these will be identical it becomes impossible for them to keep passing on identical copies of themselves to later generations. (The GP remains elitist and so passes on one copy of the best individual to the next generation without crossover). So lower selection pressure increases the diversity of the population. This permits greater exploration but at the expense of making the GP perform worse as a hill climber. The anticipated benefit of reduced premature convergence does not happen (cf. lower curves on Figure 8). Perhaps selection pressure below that given by a tournament of two is needed?

7 Price's Covariance and Selection Theorem

Price's Covariance and Selection Theorem [Price, 1970] from population genetics relates the change in frequency of a gene in a population from one generation to the next, to the covariance of the gene's frequency in the original population with the number of offspring pro-




Figure 8: Percentage of runs ending in failure with tournament sizes from 2 to 8. Lower curves shows percentage of failed runs with no \times in a higher level.

duced by individuals in that population.

$$\Delta q = \frac{\text{covariance}(z, q)}{\bar{z}} \quad (6)$$

- q = Frequency of gene
- z = Number of offspring
- \bar{z} = Mean number of children

The theorem holds “for a single gene or for any linear combination of genes at any number of loci, holds for any sort of dominance or epistasis (non-linear interaction between genes), for sexual or asexual reproduction, for random or non-random mating, for diploid, haploid or polyploid species, and even for imaginary species with more than two sexes” [Price, 1970]. While [Altenberg, 1994] says it applies to genetic algorithms (GAs) this (and [Langdon, 1996]) are the first experimental tests of its applicability to GP.

In our GP the size of the population does not change so $\bar{z} = 1$ and the expected number of children an individual has is given by its rank. So we expect the change in frequency of a gene to be given by $\text{covariance}(q, t(r_i/p)^{t-1})$, cf. (5), as long as crossover is random. Figure 9 shows good agreement between theory and measurement for $+$ nodes near the root but at the maximum allowed depth for $+$ nodes, the crossover depth restriction acts to depress the change in frequency. After three generations the number of $+$ nodes no longer increases dramatically and instead clusters near zero even though the covariance is positive. Similar effects are seen with \times nodes and the opposite with 0.5 nodes where the change in frequency tends to be larger than predicted later in the run when most trees have reached near the maximum size.

Figure 10 considers 3 runs, the first one (which failed but does find a solution if run for long enough), a suc-




Figure 9: Change in Frequency v. Covariance for + in the first $D = 8$ run (+ at depth 7 rescaled by dividing by 16).

successful run and one which can never find the optimal solution. The three are principally separated by the number of \times nodes at the second level of the tree (i.e. $d = 1$). In the successful run this rises rapidly so by generation four there are on average nearly 2.0 such nodes per individual in the population. This remains true throughout the rest of the run.

In the first run, the number of \times nodes at level 1 rises in generations 1 and 2 as it does in the successful run but then its covariance with fitness drops to near zero and the its frequency converges towards 1.0 where it remains until the end of the run. The unsuccessful run starts like the other two runs but then in generation three the covariance becomes negative and frequency falls to near zero in the next generation. The population eventually converges to zero i.e. there are no \times nodes at level 1 anywhere in the population and (as explained above) it cannot escape from this trap. NB the eventual outcome of these three runs 500 generations later can be predicted from the covariance of gene frequency with (rank based) fitness in generations two and three.

Figure 11 considers the fate of \times nodes in successful and unsuccessful runs with $D = 5$. We see from Figure 11 that the two sets of runs start from approximately the same point as their initial population but diverge radically after the first generation created by crossover. The mean covariance in successful runs is positive and remains positive, with the frequency rising. In contrast in unsuccessful runs the covariance remains low and the frequency of \times nodes fails to rise, eventually being totally displaced by + nodes.

8 Conclusions

Analysis shows on the larger MAX problems GP has two serious problems:




Figure 10: Covariance with fitness and frequency for \times in the second level of the tree in the first run, a successful run and a prematurely converged run ($D = 8$).

1. As described in [Gathercole and Ross, 1996] GP populations have a significant risk of losing vital components of the solutions at the very beginning of the run and these components can not be recovered later in the run. We have used Price's Theorem to analyse why this happens and which runs will be effected. (Similar effects are reported in [Langdon, 1996]).
2. Where solution is possible, the later stages of GP runs are effectively performing randomised hill climbing and so solution time grows exponentially with depth of the solution.

We have extended the analysis of the difficulties crossover experiences presented in [Gathercole and Ross, 1996] to include a quantitative model of the later evolution of MAX problem populations. Comparison with experimental results indicate the model gives a reasonable indication of the likely rate of improvement.

While [Gathercole and Ross, 1996] suggest that MAX problem populations converge, measurements indicate after many generations up to two thirds of the population are unique depending upon maximum depth and selection pressure. A model of this based on the interaction between crossover and the depth restriction has been presented. Additionally a model of the number of duplicate individuals in the initial populations has been presented which highlights a potential problem with the standard technique for generating the initial random population.

In the last section we used the MAX problem to demonstrate the applicability of Price's Covariance and Selection Theorem of gene frequencies to GP populations, but noted GP's depth restriction influences crossover and the consequent implications on the changes in gene frequencies. Gene covariance was anal-




Figure 11: Covariance with fitness and frequency for \times in the second level of the tree. Means of successful and unsuccessful runs. Error bars indicate standard error ($D = 5$).

used to help explain why GP populations get locked into suboptimal solutions in the first few generations.

Acknowledgements

This research is supported by the Defence Research Agency in Malvern and the British Council. Our thanks to Andy Singleton for the original version of GP-QUICK on which our code is based.

References

- [Altenberg, 1994] Lee Altenberg. The evolution of evolvability in genetic programming. In Kenneth E. Kinnear, Jr., editor, *Advances in Genetic Programming*, chapter 3, pages 47–74. MIT Press, 1994.
- [Altenberg, 1995] Lee Altenberg. The Schema Theorem and Price’s Theorem. In L. Darrell Whitley and Michael D. Vose, editors, *Foundations of Genetic Algorithms 3*, pages 23–49. Morgan Kaufmann.
- [Blickle and Thiele, 1995] Tobias Blickle and Lothar Thiele. A mathematical analysis of tournament selection. In Larry J. Eshelman, editor, *Genetic Algorithms: Proceedings of the Sixth International Conference*, pages 9–16. Morgan Kaufmann.
- [Bohm and Geyer-Schulz, 1996] Walter Bohm and Andreas Geyer-Schulz. Exact uniform initialization for genetic programming. In Richard K. Belew and Michael Vose, editors, *Foundations of Genetic Algorithms IV*, University of San Diego, CA, USA.
- [Gathercole and Ross, 1996] Chris Gathercole and Peter Ross. An adverse interaction between crossover and restricted tree depth in genetic programming.

In John R. Koza, David E. Goldberg, David B. Fogel, and Rick L. Riolo, editors, *Genetic Programming 1996: Proceedings of the First Annual Conference*.

- [Gathercole, 1997] 16 Jan 97, Private communication.
- [Holland, 1992] John H. Holland. *Adaptation in Natural and Artificial Systems: An Introductory Analysis with Applications to Biology, Control and Artificial Intelligence*. MIT Press, 1992.
- [Iba, 1996] Hitoshi Iba. Random tree generation of genetic programming. In Hans-Michael Voigt, Werner Ebeling, Ingo Rechenberg, and Hans-Paul Schwefel, editors, *Parallel Problem Solving from Nature IV, Proceedings of the International Conference on Evolutionary Computation*. Springer Verlag.
- [Koza, 1992] John R. Koza. *Genetic Programming: On the Programming of Computers by Natural Selection*.
- [Koza, 1994] John R. Koza. *Genetic Programming II: Automatic Discovery of Reusable Programs*.
- [Langdon, 1996] W. B. Langdon. *Data Structures and Genetic Programming*. PhD thesis, University College, London, 27 September 1996.
- [O’Reilly and Oppacher, 1995] Una-May O’Reilly and Franz Oppacher. The troubling aspects of a building block hypothesis for genetic programming. In L. Darrell Whitley and Michael D. Vose, editors, *Foundations of Genetic Algorithms 3*, pages 73–88.
- [Poli and Langdon, 1997] Riccardo Poli and W. B. Langdon. A new schema theory for genetic programming with one-point crossover and point mutation. Technical Report CSRP-97-3, School of Computer Science, The University of Birmingham, B15 2TT, UK, January 1997. To be presented at GP-97.
- [Price, 1970] George R. Price. Selection and covariance. *Nature*, 227, August 1:520–521, 1970.
- [Punch *et al.*, 1996] William F. Punch, Douglas Zongker, and Erik D. Goodman. The royal tree problem, a benchmark for single and multiple population genetic programming. In Peter J. Angeline and K. E. Kinnear, Jr., editors, *Advances in Genetic Programming 2*, chapter 15, pages 299–316.
- [Tackett, 1995] Walter Alden Tackett. Greedy recombination and genetic search on the space of computer programs. In L. Darrell Whitley and Michael D. Vose, editors, *Foundations of Genetic Algorithms 3*.
- [Whigham, 1995] P. A. Whigham. A schema theorem for context-free grammars. In *1995 IEEE Conference on Evolutionary Computation*, pages 178–181.