Fitness Causes Bloat: Mutation

W. B. Langdon and R. Poli

School of Computer Science, University of Birmingham, Birmingham B15 2TT, UK
{W.B.Langdon,R.Poli}@cs.bham.ac.uk http://www.cs.bham.ac.uk/~wbl, ~rmp
Tel: +44 (0) 121 414 4791, Fax: +44 (0) 121 414 4281

Abstract. The problem of evolving, using mutation, an artificial ant to follow the Santa Fe trail is used to study the well known genetic programming feature of growth in solution length. Known variously as "bloat", "fluff" and increasing "structural complexity", this is often described in terms of increasing "redundancy" in the code caused by "introns". Comparison between runs with and without fitness selection pressure, backed by Price's Theorem, shows the tendency for solutions to grow in size is caused by fitness based selection. We argue that such growth is inherent in using a fixed evaluation function with a discrete but variable length representation. With simple static evaluation search converges to mainly finding trial solutions with the same fitness as existing trial solutions. In general variable length allows many more long representations of a given solution than short ones. Thus in search (without a length bias) we expect longer representations to occur more often and so representation length to tend to increase. I.e. fitness based selection leads to bloat.

1 Introduction

The tendency for programs in genetic programming (GP) populations to grow in length has been widely reported [Tac93; Tac94; Ang94; Tac95; Lan95; NB95; SFD96]. This tendency has gone under various names such as "bloat", "fluff" and increasing "structural complexity". The principal explanation advanced for bloat has been the growth of "introns" or "redundancy", i.e. code which has no effect on the operation of the program which contains it. ([WL96] contains a survey of recent research in biology on "introns"). Such introns are said to protect the program containing them from crossover [BT94; Bli96; NFB95; NFB96]. [MM95] presents an analysis of some simple GP problems designed to investigate bloat. This shows that, with some function sets, longer programs can "replicate" more "accurately" when using crossover. I.e. offspring produced by crossover between longer programs are more likely to behave as their parents than children of shorter programs. [RB96a] provides a detailed analysis of bloat using tree schemata specifically for GP.

We advance a more general explanation which should apply generally to any discrete variable length representation and generally to any progressive search technique. That is bloat is not specific to genetic programming applied to trees using tree based crossover but should also be found with other genetic operators and non-population based stochastic search techniques such as simulated annealing and stochastic iterated hill climbing ([Lan97] contains examples where such bloat does indeed occur).

The next section summarises our argument that bloat is inherent in variable length representations such as GP [LP97b]. In Sects. 3 and 4 we expand our previous analysis of a typical GP demonstration problem to include solution by mutation in place of crossover, again showing that it suffers from bloat and also showing that bloat is not present in the absence of fitness based selection. (This improves earlier experiments [LP97c] by removing the arbitrary limit on program size). Section 5 describes the results we have achieved and this is followed in Sect. 6 by a discussion of the potential advantages and disadvantages of bloat and possible responses to it. Finally Sect. 7 summarises our conclusions.

2 Bloat in Variable Length Representations

In general with variable length discrete representations there are multiple ways of representing a given behaviour. If the evaluation function is static and concerned only with the quality of each trial solution and not with its representation then all these representations have equal worth. If the search strategy were unbiased, each of these would be equally likely to be found. In general there are many more long ways to represent a specific behaviour than short representations of the same behaviour. Thus we would expect a predominance of long representations.

Practical search techniques are biased. There are two common forms of bias when using variable length representations. Firstly search techniques often commence with simple (i.e. short) representations, i.e. they have an in built bias in favour of short representations. Secondly they have a bias in favour of continuing the search from previously discovered high fitness representations and retaining them as points for future search. I.e. there is a bias in favour of representations that do at least as well as their initiating point(s).

On problems of interest, finding improved solutions is relatively easy initially but becomes increasingly more difficult. In these circumstances, especially with a discrete fitness function, there is little chance of finding a representation that does better than the representation(s) from which it was created (cf. "death of crossover" [Lan96a, page 222]). So the selection bias favours representations which have the same fitness as those from which they were created.

In general the easiest way to create one representation from another and retain the same fitness is for the new representation to represent identical behaviour. Thus, in the absence of improved solutions, the search may become a random search for new representations of the best solution found so far. As we said above, there are many more long representations than short ones for the same solution, so such a random search (other things being equal) will find more long representations than short ones. In GP this has become known as bloat.

3 The Artificial Ant Problem

The artificial ant problem is described in [Koz92, pages 147–155]. It is a well studied problem and was chosen as it has a simple fitness function. Briefly the problem is to devise a program which can successfully navigate an artificial ant along a twisting trail on a square 32×32 toroidal grid. The program can use three operations, Move, Right and Left, to move the ant forward one square, turn to the right or turn to the left. Each of these operations takes one time unit. The sensing function IfFoodAhead looks into the square the ant is currently facing and then executes one of its two arguments depending upon whether that square contains food or is empty. Two other functions, Prog2 and Prog3, are provided. These take two and three arguments respectively which are executed in sequence.

The artificial ant must follow the "Santa Fe trail", which consists of 144 squares with 21 turns. There are 89 food units distributed non-uniformly along it. Each time the ant enters a square containing food the ant eats it. The amount of food eaten is used as the fitness measure of the control program.

The evolutionary system we use is identical to [LP97b] except the crossover operator is replaced by mutation. The details are given in Table 1, parameters not shown are as [Koz94, page 655]. On each version of the problem 50 independent runs were conducted.

Table 1. Ant Problem

Objective:	Find an ant that follows the "Santa Fe trail"
Terminal set:	Left, Right, Move
Functions set:	IfFoodAhead, Prog2, Prog3
Fitness cases:	The Santa Fe trail
Fitness:	Food eaten
Selection:	Tournament group size of 7, non-elitist, generational
Wrapper:	Program repeatedly executed for 600 time steps.
Population Size:	500
Max program size:	32,767
Initial population:	Created using "ramped half-and-half" with a max depth of 6
Parameters:	90% mutation, 10% reproduction
Termination:	Maximum number of generations G = 50

Note in these experiments we allow the evolved programs to be far bigger than required to solve the problem. For example the 100% correct solution given in [Koz92, page 154] takes about 543 time steps to traverse the Santa Fe trail but has a length of only 18 nodes and this is not the most compact solution possible.

4 Tree Mutation

For our purposes it is necessary that the mutation operator be able to change the size of the chromosomes it operates. Ideally this should be unbiased in the sense of producing, of itself, no net change in size. The following operator, in all cases examined, produces offspring which are on average almost the same size as their parent (cf. Sect. 5.6). (NB this may not be the case with different ratios of node branching factors in either the terminal/function set or in the evolving populations). An alternative tree mutation operator which also tries to avoid size bias by creating random trees of a randomly chosen size is proposed in [Lan97].

The mutation operator selects uniformly at random a node (which may be a function or terminal) and replaces it and the subtree descending from it with a randomly created tree. The new tree is created using the same "half and half" method used to create the initial population [Koz92, Page 92–93] however its maximum height is chosen at random from one to the height of the tree it is to replace. NB a terminal will always be replaced by a randomly selected terminal (i.e. there is no change in size) but a function can be replaced by a larger (but not deeper) or smaller tree, so that the program's size may change.

5 Results

5.1 Standard Runs

In 50 independent runs 9 found "ants" that could eat all the food on the Santa Fe trail within 600 time steps. The evolution of maximum and mean fitness averaged across all 50 runs is given in Fig. 1. (In all cases the average minimum fitness is near zero). These curves show the fitness behaving as with crossover with both the maximum and average fitness rising rapidly initially but then rising more slowly later in the runs. However with mutation the rise in mean fitness is less rapid, e.g. it passes 50 at generation 36 compared to generation 17 with crossover. The population converges in the sense that the average fitness approaches the maximum fitness. However the spread of fitness values of the children produced in each generation remains large and children which eat either no food or only one food unit are still produced even in the last generation.

Figure 2 shows the evolution of maximum and mean program size averaged across all 50 runs. We see in the first four generations the average program length grows rapidly from an initial size of 23.5 to 70.0. Unlike with crossover the average program length then falls to about 40 but the maximum program size grows continually (excepting some noisy oscillation). The mean program size remains fairly constant until generation 19 when bloat sets in and it begins to rise progressively.

As with crossover the mean program length grows faster than the size of the "best" program within the population. Between generations 2 and 12 the mean length exceeds that of the "best" program (cf. Sect. 5.1). In the later generations the mean program length and length of the best program on average lie close to each other.

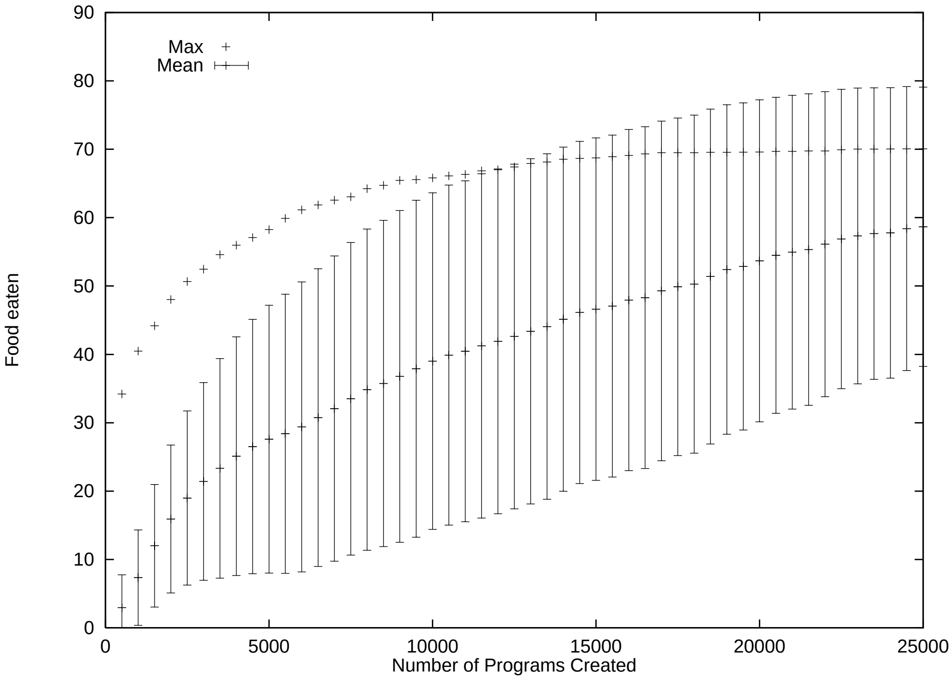

Fig. 1. Evolution of maximum and population mean of food eaten. Error bars indicate one standard deviation. Means of 50 runs.

Fitness is Necessary for Bloat – Price's Theorem Applied to Representation Size. Price's Covariance and Selection Theorem [Pri70] from population genetics relates the expected change in frequency of a gene Δq 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 z produced by individuals in that population:

$$\Delta q = \frac{\mathrm{Cov}(z,q)}{\overline{z}} \tag{1}$$

We have used it to help explain the evolution of the number of copies of functions and terminals in GP populations [Lan96a; LP97a]. In our experiments the size of the population does not change so $\overline{z} = 1$ and the expected number of children is given by the parent's rank so in large populations the expected change is approximately $\mathrm{Cov}(t(r/p)^{t-1}, q)$ as long as crossover is random. (t is the tournament size and r is each program's rank within the population of size p).

Where representation length is inherited, such as in GP and other search techniques, (1) should hold for representation length. More formally Price's theorem applies (provided length and genetic operators are uncorrelated) since representation length is a *measurement function* of the genotype [Alt95, page 28]. If it

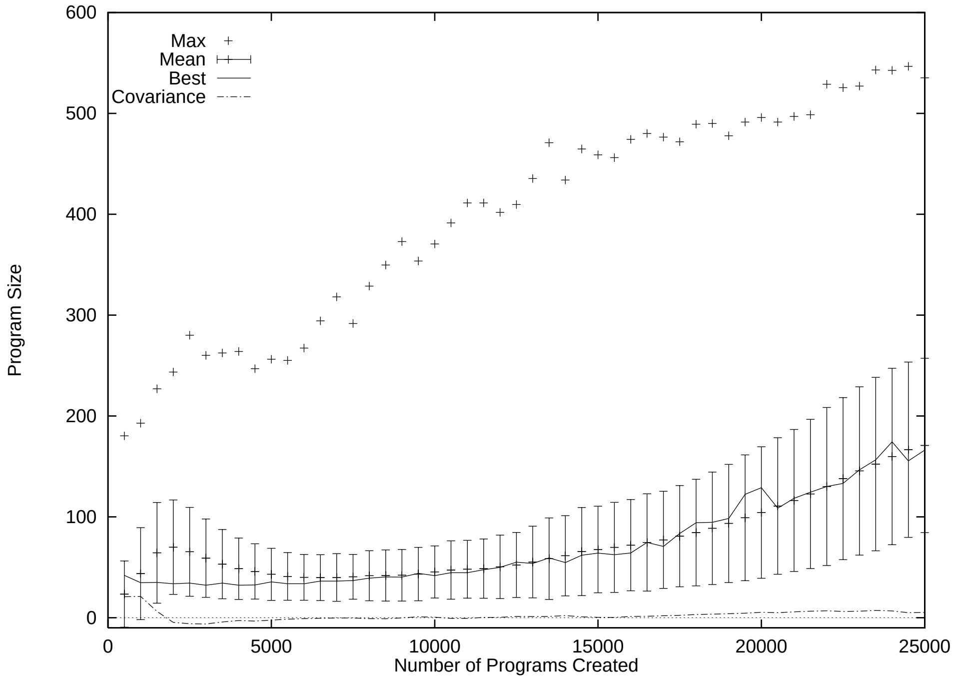

Fig. 2. Evolution of maximum and population mean program length. Error bars indicate one standard deviation. Solid line is the length of the "best" program in the population. Covariance of length and normalised rank based fitness shown dotted. Means of 50 runs.

held exactly a plot of covariance vs. change in mean length would be a straight line (assuming $\overline{z}$ is constant). The solid line plot in Fig. 3 shows good agreement between theory and measurement. (The slight discrepancy is due to a small bias in the mutation operator, which is plotted in Fig. 11).

Essentially (1) gives a quantitative measurement of the way genetic algorithms (GAs) search. If some aspect of the genetic material is positively correlated with fitness then, other things being equal, the next generation's population will on average contain more of it. If it is negative, then the GA will tend to reduce it in the next generation.

For a given distribution of representation lengths (1) says the change in mean representation length will be linearly related to the selection pressure, z. This provides some theoretical justification for the claim [Tac94, page 112] that "average growth in size ... is proportional to selection pressure". His claim is based upon experimental measurements with a small number of radically different selection and crossover operators on Tackett's "Royal Road" problem. (Solutions to the Royal Road problem are required to have prespecified syntactic properties which mean "that the optimal solution has a fixed size and does not admit extraneous code selections". Such solutions are not executable programs).

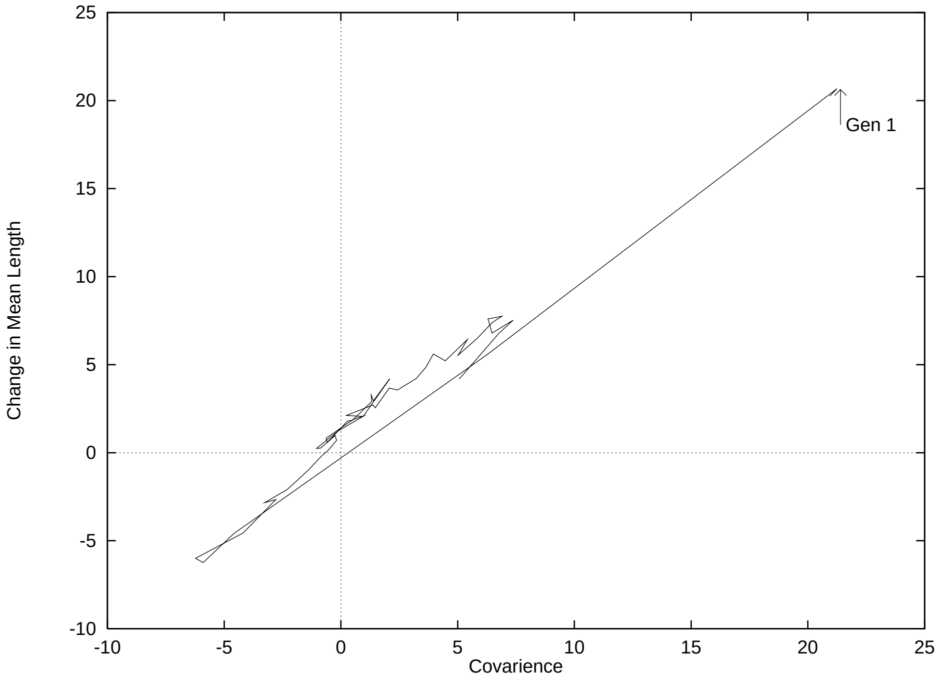

Fig. 3. Covariance of program length and normalised rank based fitness vs. change in mean length in next generation. Means of 50 runs.

Referring back to Fig. 2 and considering the length of the best solution and the covariance in the first four generations, we see the covariance correctly predicts the mean length of programs will increase. In contrast if we assumed the GA would converge towards the individual with the best fitness in the population, the fact that the mean length is higher than the mean length of the best individual would lead us to predict a fall in mean program length. That is Price's Theorem is the better predictor. This is because it considers the whole population rather than just one (albeit the best in the population).

Where fitness selection is not used (as in the following sections), each individual in the population has an equal chance of producing children and so the covariance is always zero. Therefore Price's Theorem predicts on average there will be no change in length.

5.2 No Selection

A further 50 runs were conducted using the same initial populations and no fitness selection. As with crossover no run found a solution and the maximum, mean and other fitness statistics fluctuate a little but are essentially unchanged (cf. Fig. 4). In contrast to the case with crossover only, our mutation operator has a slight bias. In the first 50 generation this causes the population to gradually increase in size at the rate of $\frac{1}{3}$ of a node per generation (cf. Fig. 5). Given mutation produces random changes in length we expect the spread of lengths to gradually increase. Unlike with crossover where the mean maximum program size fell slowly, the average maximum size grows continually and the mean standard deviation grows steadily from 32.8 to 74.8 (In contrast with crossover the mean standard deviation fell from 32.8 in the initial population to 27.3 at the end of the runs).

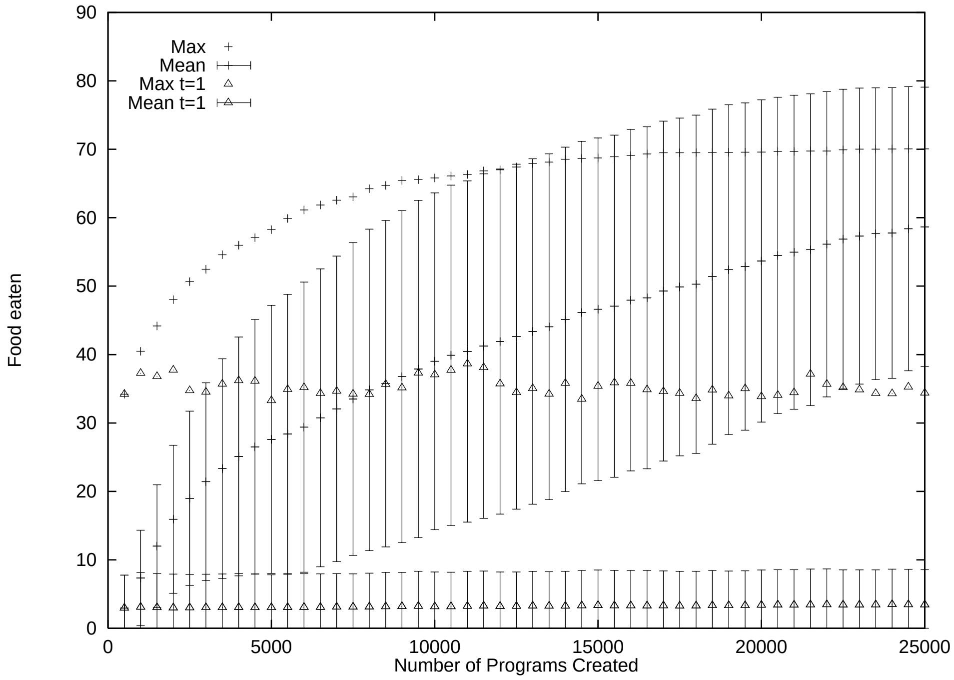

Fig. 4. Evolution of maximum and population mean of food eaten, in runs with and without fitness selection (i.e. tournament sizes of 7 and 1). Error bars indicate one standard deviation. Means of 50 runs.

5.3 Removing Selection

A final 50 runs were conducted in which fitness selection was removed after generation 25 (i.e. these runs are identical to those in Sect. 5.1 up to generation 25). The evolution of maximum and mean fitness averaged across all 50 runs is given in Fig. 6. As expected Fig. 6 shows in the absence of fitness selection the fitness of the population quickly falls.

Of the ten runs which found ants that could eat all the food on the Santa Fe trail within 600 movements, three first found solutions after generation 26 and

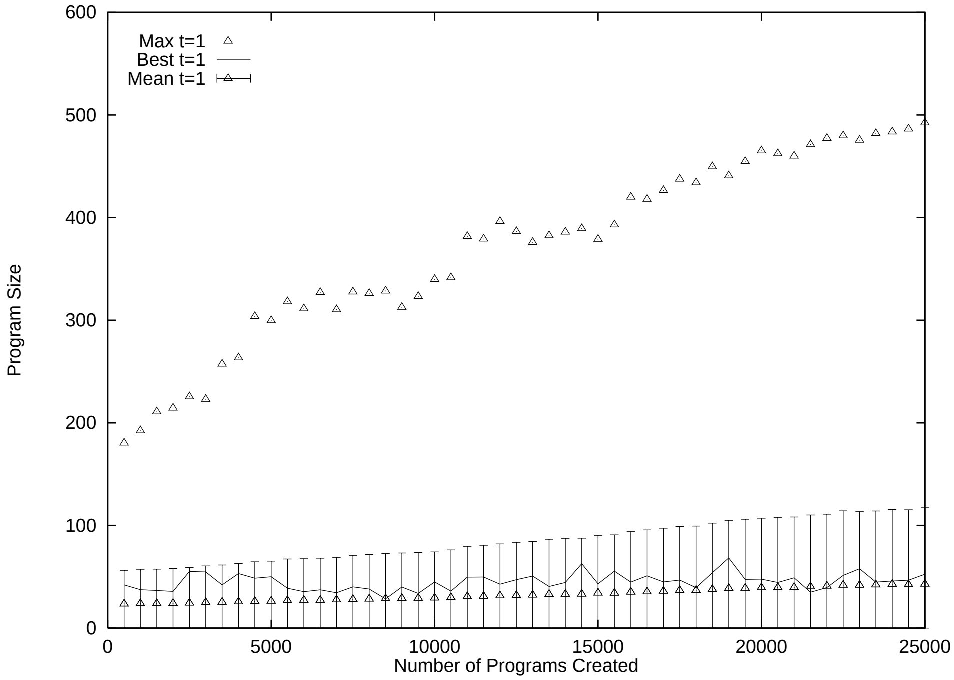

Fig. 5. Evolution of maximum and population mean program length when fitness selection is absent. Error bars indicate one standard deviation. Solid line is the length of the "best" program in the population. Means of 50 runs.

so failed to find solutions when fitness selection was removed, however two other runs did find solutions when it was removed. In one run the population managed to retain individuals with the maximum score until the end of the run. In all other cases mutation eliminated solutions from the population after fitness selection was removed, typically within 8 generations.

After fitness selection is removed, the length of programs behaves much as when there is no selection from the start of the run. I.e. bloat is replaced by the slow growth associated with the mutation operator and the spread of program lengths gradually increases (cf. Fig. 7)

5.4 Correlation of Fitness and Program Size

Figure 8 plots the correlation coefficient of program size and amount of food eaten. (Correlation coefficients are equal to the covariance after it has been normalised to lie in the range $-1 \ldots +1$. By considering food eaten we avoid the intermediate stage of converting program score to expected number of children required when applying Price's Theorem). Figure 8 shows in all cases there is a positive correlation between program score (rather than fitness) and length of programs. In the initial population, long random programs do better than short

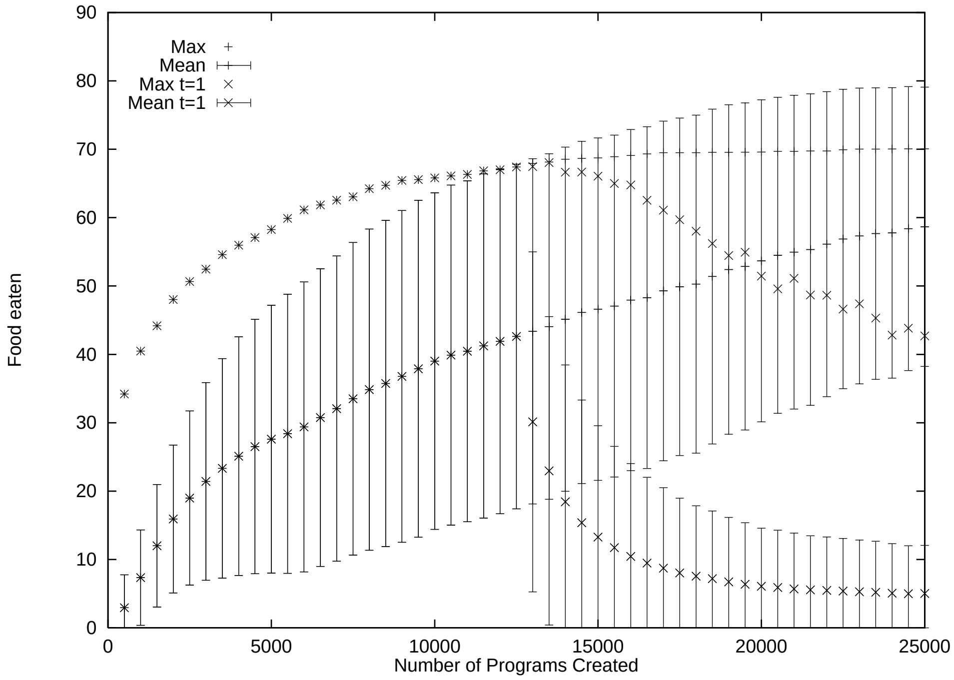

Fig. 6. Evolution of maximum and population mean of food eaten. t=1 curves show effect of removing fitness selection after generation 25. Error bars indicate one standard deviation. Means of 50 runs.

ones. This may be because they are more likely to contain useful primitives (such as Move) than short programs. Also short ones make fewer moves before they are reinitialised and re-executed. This may increase the chance of them falling into unproductive short cyclic behaviour that longer, more random, programs can avoid. Selection quickly drives the population towards these better individuals, so reducing the correlation. With our mutation operator it falls to a much lower level than with crossover (when it was typically in the region of +0.3). This might indicate that mutation is less disruptive than crossover. In the absence of selection (or after selection has been removed halfway through a run) mutation tends to randomise the population so that the correlation remains (or tends towards) that in the initial population.

5.5 Effect of Mutation on Fitness

In the initial generations mutation is disruptive (cf. Fig. 9) with 64.6% producing a child with a different score from its parent. However mutation evolves, like crossover, to become less disruptive. By the end of the run 74.6% of mutants have the same score as their parent.

The range of change of fitness is highly asymmetric; many more children are

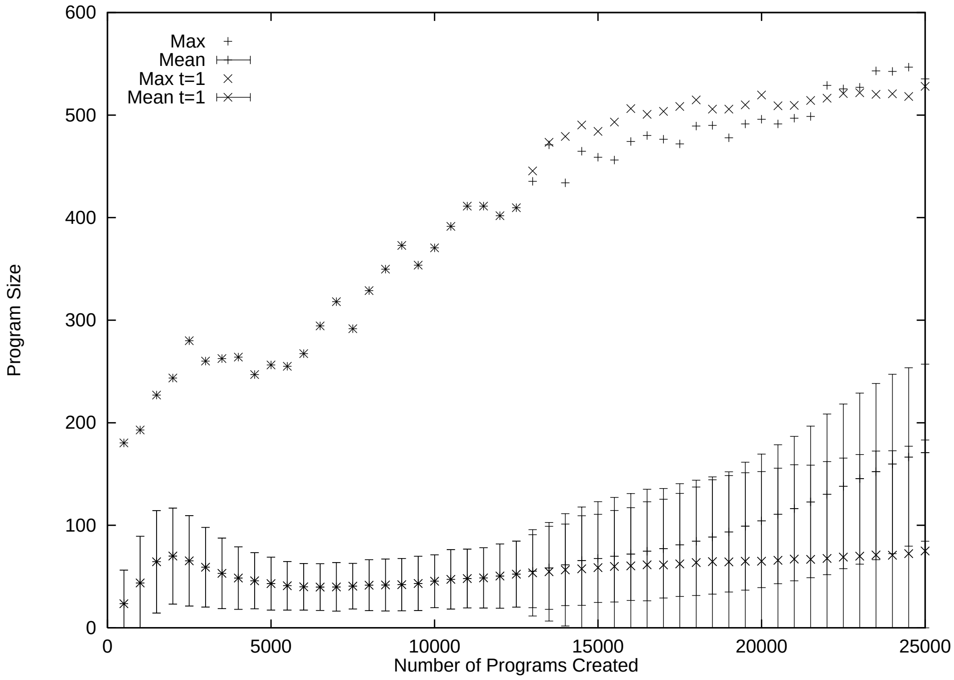

Fig. 7. Evolution of maximum and population mean program length. t=1 curves show effect of removing fitness selection after generation 25. Error bars indicate one standard deviation. Solid line is the length of the "best" program in the population. Means of 50 runs.

produced which are worse than their parent than those that are better. By the end of the run, only 0.02% of the population are fitter than their parent. Similar behaviour has been reported using crossover on other problems [NFB96] [RB96b, page 183] [Lan96a, Chapter 7].

5.6 Non-Disruptive Mutation and Program Length

In Sect. 2 we argued that there are more long programs with a given performance than short ones and so a random search for programs with a given level of performance is more likely to find long programs. We would expect generally (and it has been reported on a number of problems) that crossover produces progressively fewer improved solutions as evolution proceeds and instead in many problems selection drives it to concentrate upon finding solutions with the same fitness. Apart from the minimum size restriction mutation is nearly random so we would expect to see it finding more long solutions than short ones. The distribution of changes in program size produced by our mutation operator has a large central spike of programs with the same length and rapidly non-monotonic falling tails either side of zero. Figure 10 plots the change in length of only

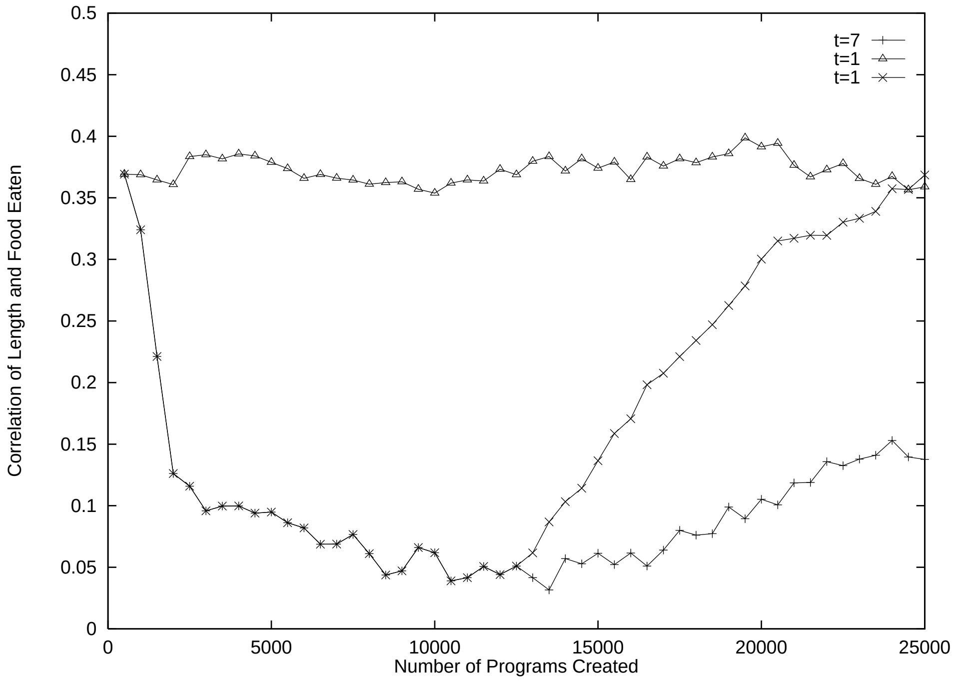

Fig. 8. Correlation of program length and fitness: normal runs, runs without selection and runs with selection removed halfway through. Means of 50 runs.

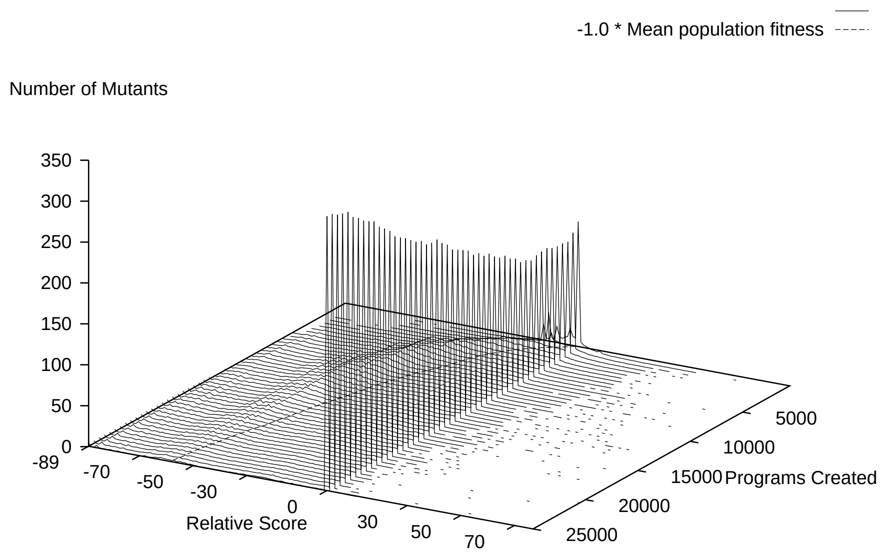

Fig. 9. Fitness relative to parent: normal runs. Means of 50 runs.

those programs produced which have worse fitness than their parent. As when considering all mutations, Fig. 10 has a large central spike of programs with the same length and rapidly falling tails either side of zero. I.e. the distribution of change of size produced by mutations which are worse than their parents is similar to the distribution of all mutations. The proportion of children produced by mutation which are identical to their parent remains on average fairly constant at about 13% throughout the runs.

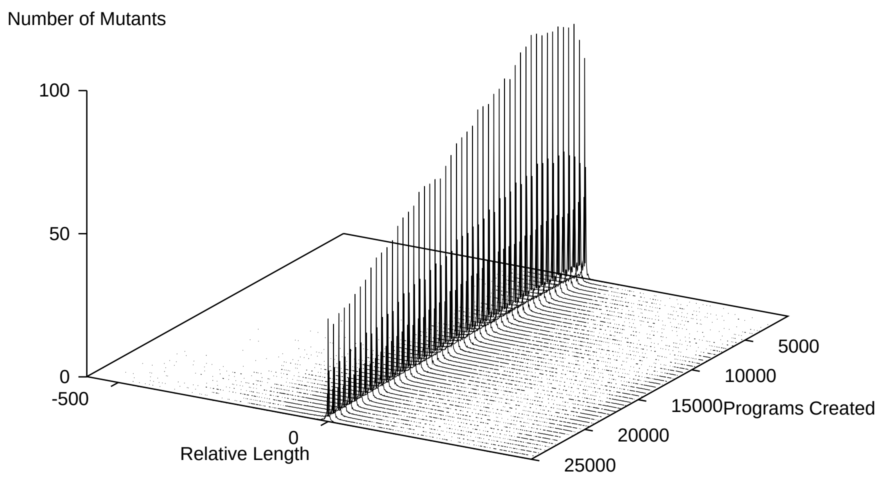

Fig. 10. Change in length of offspring relative to parent where score is worse: normal runs. Means of 50 runs.

As with crossover, mutation becomes increasingly less disruptive (after 50 generations on average only 25% of mutations produce a change in fitness). We see from Fig. 11 on average mutation makes little change to the length of programs, however, after bloat becomes established, our mutation operator produces children which are slightly shorter than their parents on average. I.e. once the change in the fitness of the population slows, program size bloats despite length changes introduced by mutation. The positive covariance of fitness and length shows this bloat is driven by fitness selection.

5.7 Fraction of Bloated Programs Used

Every program terminal uses one energy unit each time it is executed. Thus the amount of energy used when executing a program gives the actual number of terminals it has used. From Fig. 12 we see (like with crossover) typically only

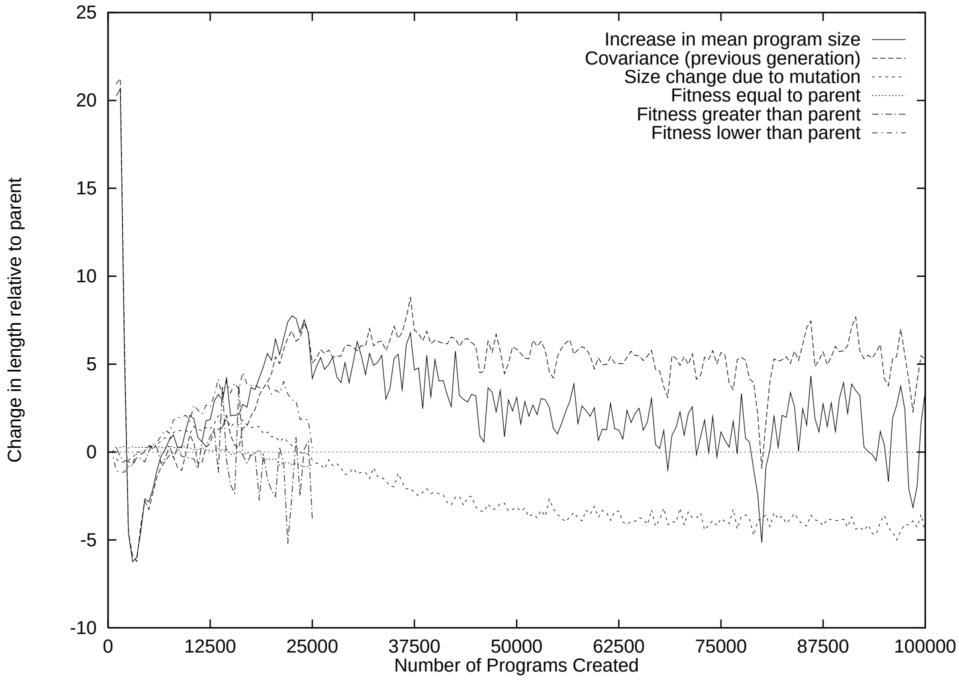

Fig. 11. Mean change in length of offspring relative to parent: normal runs. Means of 50 runs.

5–10 terminals are used per execution (a move and four turns allows the ant to move forward whilst searching either side of it). While different parts of the program could be executed each time it is called, we expect the large programs evolved contain only small functional parts. In one bloated solution only 20 of 180 nodes in the program could be executed and of these five more were executed but had no net effect.

6 Discussion

6.1 Do we Want to Prevent Bloat?

From a practical point of view the machine resources consumed by any system which suffers from bloat will prevent extended operation of that system. However in practice we may not wish to operate the system continually. For example it may quickly find a satisfactory solution or better performance may be achieved by cutting short its operation and running it repeatedly with different starting configurations [Koz92, page 758].

In some data fitting problems growth in solution size may be indicative of "over fitting", i.e. better matching on the test data but at the expense of general performance. For example [Tac93, page 309] suggests "parsimony may be an

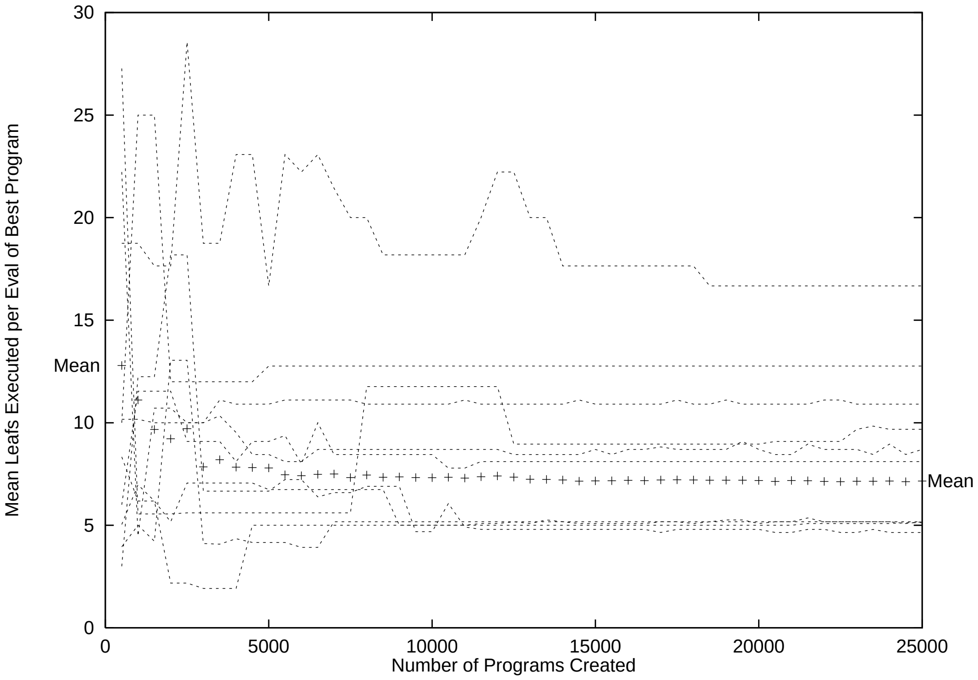

Fig. 12. Energy used per call of "best" program. 10 of 50 normal runs.

important factor not for 'aesthetic' reasons or ease of analysis, but because of a more direct relationship to fitness: there is a bound on the 'appropriate size' of solution tree for a given problem".

By providing a "defence against crossover" [NFB96, page 118] bloat causes the production of many programs of identical performance. These can consume the bulk of the available machine resources and by "clogging up" the population may prevent GP from effectively searching for better programs.

On the other hand [Ang94, page 84] quotes results from fixed length GAs in favour of representations which include introns, to argue we should "not ... impede this emergent property [i.e. introns] as it may be crucial to the successful development of genetic programs". Introns may be important as "hiding places" where genetic material can be protected from the current effects of selection and so retained in the population. This may be especially important where fitness criteria are dynamic. A change in circumstance may make it advantageous to execute genetic material which had previously been hidden in an intron. [Hay96] shows an example where a difficult GP representation is improved by deliberately inserting duplicates of evolved code.

In complex problems it may not be possible to test every solution on every aspect of the problem and some form of dynamic selection of test cases may be required [GR94]. For example in some cases co-evolution has been claimed to be beneficial to GP. If the fitness function is sufficiently dynamic, will there still be an advantage for a child in performing identically to its parents? If not, will we still see such explosive bloat?

6.2 Three Ways to Control Bloat

Three methods of controlling bloat have been suggested. Firstly, and most widely used (e.g. in these experiments), is to place a universal upper bound either on tree depth [Koz92] or program length. ([GR96; LP97a] discuss unexpected problems with this approach).

The second (also commonly used) is to incorporate program size directly into the fitness measure (often called parsimony pressure) [Koz92; ZM93; IdS94]. [RB96a] gives an analysis of the effect of parsimony pressure which varies linearly with program length. Multi-objective fitness measures where one objective is compact or fast programs have also been used [Lan96b].

The third method is to tailor the genetic operations. [Sim93, page 469] uses several mutation operators but adjusts their frequencies so a "decrease in complexity is slightly more probable than an increase". [Bli96] suggests targeting genetic operations at redundant code. This is seldom used, perhaps due to the complexity of identifying redundant code. [SFD96] showed bloat continuing despite their targeted genetic operations. Possibly this was because of the difficulty of reliably detecting introns. I.e. there was a route whereby the GP could evolve junk code which masqueraded as being useful and thereby protected itself from removal. While [RB96a] propose a method where the likelihood of potentially disruptive genetic operations increases with parent size.

7 Conclusions

We have generalised existing explanations for the widely observed growth in GP program size with successive generations (*bloat*) to give a simple statistical argument which should be generally applicable both to GP and other systems using discrete variable length representations and static evaluation functions. Briefly, in general simple static evaluation functions quickly drive search to converge, in the sense of concentrating the search on trial solutions with the same fitness as previously found trial solutions. In general variable length allows many more long representations of a given solution than short ones of the same solution. Thus (in the absence of a parsimony bias) we expect longer representations to occur more often and so representation length to tend to increase. I.e. fitness based selection leads to bloat.

In earlier work [LP97b] we took a typical GP problem and demonstrated with fitness selection it suffers from bloat whereas without selection it does not. In Sects. 3, 4 and 5 we repeated these experiments replacing crossover with mutation and showed fitness selection can still cause bloat. (In these experiments we remove the maximum size limit of 500 nodes used in [LP97c]. The results in Sect. 5 are similar to those in [LP97c] confirming our earlier work and showing a 500 node limit had little effect in those experiments). We have demonstrated that if fitness selection is removed, bloat is stopped and program size changes little on average. As expected in the absence of selection, mutation is free to change program size at random and the range of sizes increases, as does the mean size (albeit very slowly). In contrast crossover of small programs typical of the initial population does not seem so free and random growth was not observed in our earlier experiments. Detailed measurement of mutation confirms after an extended period of evolution, most are not disruptive (i.e. most children have the same fitness as their parents).

In Sect. 5.1 we applied Price's Theorem to program lengths within evolving populations. We confirmed experimentally that it fits unless bias in the genetic operators have significant impact. We used Price's Theorem to prove fitness selection is required for a change in average representation length. In Sect. 6 we discussed the circumstances in which we need to control bloat and current mechanisms which do control it but suggest a way forward may be to consider more complex dynamic fitness functions.

Acknowledgements

This research was funded by the Defence Research Agency in Malvern.

References

[Alt95] 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, Estes Park, Colorado, USA, 31 July–2 August 1994 1995. Morgan Kaufmann.

[Ang94] Peter John Angeline. Genetic programming and emergent intelligence. In Kenneth E. Kinnear, Jr., editor, *Advances in Genetic Programming*, chapter 4, pages 75–98. MIT Press, 1994.

[Bli96] Tobias Blickle. *Theory of Evolutionary Algorithms and Application to System Synthesis*. PhD thesis, Swiss Federal Institute of Technology, Zurich, November 1996.

[BT94] Tobias Blickle and Lothar Thiele. Genetic programming and redundancy. In J. Hopf, editor, *Genetic Algorithms within the Framework of Evolutionary Computation (Workshop at KI-94, Saarbrücken)*, pages 33–38, Im Stadtwald, Building 44, D-66123 Saarbrücken, Germany, 1994. Max-Planck-Institut für Informatik (MPI-I-94-241).

[GR94] Chris Gathercole and Peter Ross. Dynamic training subset selection for supervised learning in genetic programming. In Yuval Davidor, Hans-Paul Schwefel, and Reinhard Männer, editors, *Parallel Problem Solving from Nature III*, pages 312–321, Jerusalem, 9-14 October 1994. Springer-Verlag.

[GR96] 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*, pages 291–296, Stanford University, CA, USA, 28–31 July 1996. MIT Press.

[Hay96] Thomas Haynes. Duplication of coding segments in genetic programming. In *Proceedings of the Thirteenth National Conference on Artificial Intelligence*, pages 344–349, Portland, OR, August 1996.

[IdS94] Hitoshi Iba, Hugo de Garis, and Taisuke Sato. Genetic programming using a minimum description length principle. In Kenneth E. Kinnear, Jr., editor, *Advances in Genetic Programming*, chapter 12, pages 265–284. MIT Press, 1994.

[Koz92] John R. Koza. *Genetic Programming: On the Programming of Computers by Natural Selection*. MIT Press, Cambridge, MA, USA, 1992.

[Koz94] John R. Koza. *Genetic Programming II: Automatic Discovery of Reusable Programs*. MIT Press, Cambridge Massachusetts, May 1994.

[Lan95] W. B. Langdon. Evolving data structures using genetic programming. In L. Eshelman, editor, *Genetic Algorithms: Proceedings of the Sixth International Conference (ICGA95)*, pages 295–302, Pittsburgh, PA, USA, 15-19 July 1995. Morgan Kaufmann.

[Lan96a] W. B. Langdon. *Data Structures and Genetic Programming*. PhD thesis, University College, London, 27 September 1996.

[Lan96b] William B. Langdon. Data structures and genetic programming. In Peter J. Angeline and K. E. Kinnear, Jr., editors, *Advances in Genetic Programming 2*, chapter 20, pages 395–414. MIT Press, Cambridge, MA, USA, 1996.

[Lan97] W. B. Langdon. Fitness causes bloat: Simulated annealing, hill climbing and populations. Technical Report CSRP-97-22, University of Birmingham, School of Computer Science, 2 September 1997.

[LP97a] W. B. Langdon and R. Poli. An analysis of the MAX problem in genetic programming. In John R. Koza, Kalyanmoy Deb, Marco Dorigo, David B. Fogel, Max Garzon, Hitoshi Iba, and Rick L. Riolo, editors, *Genetic Programming 1997: Proceedings of the Second Annual Conference*, pages 222–230, Stanford University, CA, USA, 13-16 July 1997. Morgan Kaufmann.

[LP97b] W. B. Langdon and R. Poli. Fitness causes bloat. In P. K. Chawdhry, R. Roy, and R. K. Pan, editors, *Second On-line World Conference on Soft Computing in Engineering Design and Manufacturing*. Springer-Verlag London, 23-27 June 1997.

[LP97c] W. B. Langdon and R. Poli. Fitness causes bloat: Mutation. In John Koza, editor, *Late Breaking Papers at the GP-97 Conference*, pages 132–140, Stanford, CA, USA, 13-16 July 1997. Stanford Bookstore.

[MM95] Nicholas Freitag McPhee and Justin Darwin Miller. Accurate replication in genetic programming. In L. Eshelman, editor, *Genetic Algorithms: Proceedings of the Sixth International Conference (ICGA95)*, pages 303–309, Pittsburgh, PA, USA, 15-19 July 1995. Morgan Kaufmann.

[NB95] Peter Nordin and Wolfgang Banzhaf. Complexity compression and evolution. In L. Eshelman, editor, *Genetic Algorithms: Proceedings of the Sixth International Conference (ICGA95)*, pages 310–317, Pittsburgh, PA, USA, 15-19 July 1995. Morgan Kaufmann.

[NFB95] Peter Nordin, Frank Francone, and Wolfgang Banzhaf. Explicitly defined introns and destructive crossover in genetic programming. In Justinian P. Rosca, editor, *Proceedings of the Workshop on Genetic Programming: From Theory to Real-World Applications*, pages 6–22, Tahoe City, California, USA, 9 July 1995.

[NFB96] Peter Nordin, Frank Francone, and Wolfgang Banzhaf. Explicitly defined introns and destructive crossover in genetic programming. In Peter J. Angeline and K. E. Kinnear, Jr., editors, *Advances in Genetic Programming 2*, chapter 6, pages 111–134. MIT Press, Cambridge, MA, USA, 1996.

[Pri70] George R. Price. Selection and covariance. *Nature*, 227, August 1:520–521, 1970.

[RB96a] Justinian P. Rosca and Dana H. Ballard. Complexity drift in evolutionary computation with tree representations. Technical Report NRL5, University of Rochester, Computer Science Department, Rochester, NY, USA, December 1996.

[RB96b] Justinian P. Rosca and Dana H. Ballard. Discovery of subroutines in genetic programming. In Peter J. Angeline and K. E. Kinnear, Jr., editors, *Advances in Genetic Programming 2*, chapter 9, pages 177–202. MIT Press, Cambridge, MA, USA, 1996.

[SFD96] Terence Soule, James A. Foster, and John Dickinson. Code growth 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*, pages 215–223, Stanford University, CA, USA, 28–31 July 1996. MIT Press.

[Sim93] K. Sims. Interactive evolution of equations for procedural models. *The Visual Computer*, 9:466–476, 1993.

[Tac93] Walter Alden Tackett. Genetic programming for feature discovery and image discrimination. In Stephanie Forrest, editor, *Proceedings of the 5th International Conference on Genetic Algorithms, ICGA-93*, pages 303–309, University of Illinois at Urbana-Champaign, 17-21 July 1993. Morgan Kaufmann.

[Tac94] Walter Alden Tackett. *Recombination, Selection, and the Genetic Construction of Computer Programs.* PhD thesis, University of Southern California, Department of Electrical Engineering Systems, 1994.

[Tac95] 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*, pages 271–297, Estes Park, Colorado, USA, 31 July–2 August 1994 1995. Morgan Kaufmann.

[WL96] Annie S. Wu and Robert K. Lindsay. A survey of intron research in genetics. 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*, volume 1141 of *LNCS*, pages 101–110, Berlin, Germany, 22-26 September 1996. Springer-Verlag.

[ZM93] Byoung-Tak Zhang and Heinz Mühlenbein. Evolving optimal neural networks using genetic algorithms with Occam's razor. *Complex Systems*, 7:199–220, 1993.