

Genetic Programming & Bloat

Jiří Kubalík
Department of Cybernetics, CTU Prague

Substantial part of this material is based on

Sean Luke and Liviu Panait: A Comparison of Bloat Control Methods for Genetic Programming,
See <http://portal.acm.org/citation.cfm?id=1182892.1182897>

Sara Silva: Handling Bloat in GP,
See <http://doi.acm.org/10.1145/2001858.2002146>



<http://cw.felk.cvut.cz/doku.php/courses/a0m33eoa/start>

Theories of Code Bloat Based on Introns

Hitchhiking – based on genetic algorithms, where unfit building blocks propagate in the population because they join highly fit building blocks.

- There is no real need to get rid of hitchhikers that do not damage fitness of the program.

Introns are hitchhikers in GP.

- The theory only suggests a propagation method.

It does not explain why it is more likely that the introns become attached in the first place than to be removed eventually.

Theories of Code Bloat Based on Introns

Removal Bias – to maintain fitness, the removed subtree must be contained within the inviable region. Since the inserted subtree can have any size, offspring are bigger than average while retaining the fitness of their parents.

- Inactive code in a GP tree tends to be low in the tree, residing therefore in smaller-than-average subtrees.
- Inserted subtree is bigger than the excised ones.

Non-Intron Theories of Code Bloat

Fitness Causes Bloat – there are many more longer ways than shorter ways to represent the same program, so a natural drift occurs to longer programs.

- Beyond a certain program length, the distribution of fitness does not vary with size.
- Since there are more longer programs, the number of long programs of a given fitness is greater than the number of short programs of the same fitness.
- Over time, GP samples longer and longer programs simply because there are more of them.

Non-Intron Theories of Code Bloat

Fitness Causes Bloat – there are many more longer ways than shorter ways to represent the same program, so a natural drift occurs to longer programs.

- Beyond a certain program length, the distribution of fitness does not vary with size.
- Since there are more longer programs, the number of long programs of a given fitness is greater than the number of short programs of the same fitness.
- Over time, GP samples longer and longer programs simply because there are more of them.

Modification Point Depth

- When a genetic operator modifies an individual, the deeper the modification point the smaller the change in fitness.
- Small changes are less likely to be disruptive, so there is a preference for deeper modification points, and consequently a preference for larger trees (removal bias).

Non-Intron Theories of Code Bloat

Fitness Causes Bloat – there are many more longer ways than shorter ways to represent the same program, so a natural drift occurs to longer programs.

- Beyond a certain program length, the distribution of fitness does not vary with size.
- Since there are more longer programs, the number of long programs of a given fitness is greater than the number of short programs of the same fitness.
- Over time, GP samples longer and longer programs simply because there are more of them.

Modification Point Depth

- When a genetic operator modifies an individual, the deeper the modification point the smaller the change in fitness.
- Small changes are less likely to be disruptive, so there is a preference for deeper modification points, and consequently a preference for larger trees (removal bias).

Crossover Bias

- Subtree crossover operators do not add to or remove from the population any amount of genetic code, they simply swap it between individuals.
- So the average program length in the population is not changed by the crossover.
- There is a bias of the crossover operators to create many small, and unfit, individuals.
- When these small unfit individuals compete for breeding, they are always discarded by selection in favor of the larger ones.

Lexicographic Parsimony Pressure Method: Direct Bucketing

Realization: The number of buckets, b , is specified beforehand, and each is assigned a quality rank from 1 to b (the bucket with rank 1 contains the worst-fit individuals).

1. The population of size p is sorted by fitness.
2. The bottom $\lceil p/b \rceil$ individuals are placed in the worst bucket.

All individuals remaining in the population with the same fitness as the best individual in the bucket are placed in the bucket as well.

This is to guarantee that all individuals of the same fitness fall into the same bucket (they have the same rank).

3. The same procedure is used to fill in the second worst bucket, the third one etc.
This continues until there are no individuals in the population.
4. The fitness of each individual is set to the rank assigned to the bucket holding it.

Lexicographic Parsimony Pressure Method: Direct Bucketing

Realization: The number of buckets, b , is specified beforehand, and each is assigned a quality rank from 1 to b (the bucket with rank 1 contains the worst-fit individuals).

1. The population of size p is sorted by fitness.
2. The bottom $\lceil p/b \rceil$ individuals are placed in the worst bucket.

All individuals remaining in the population with the same fitness as the best individual in the bucket are placed in the bucket as well.

This is to guarantee that all individuals of the same fitness fall into the same bucket (they have the same rank).

3. The same procedure is used to fill in the second worst bucket, the third one etc.

This continues until there are no individuals in the population.

4. The fitness of each individual is set to the rank assigned to the bucket holding it.

Characteristics:

- It has the effect of trading off fitness differences for size.
- The larger the bucket, the stronger the emphasis on size as a secondary objective.
- The topmost bucket with the best-fit individuals can hold fewer than $\lceil p/b \rceil$ individuals.

Lexicographic Parsimony Pressure Method: Ratio Bucketing

Realization: The buckets are proportioned, so that low-fitness individuals are placed into larger buckets than high-fitness individuals. A parameter of the method is the bucket ratio $1/r$.

1. The population of size p is sorted by fitness.
2. The bottom $\lceil 1/r \rceil$ fraction of individuals are placed into the worst bucket.
All individuals remaining in the population with the same fitness as the best individual in the bucket are placed in the bucket as well.
3. The same procedure is used to fill in the second worst bucket with the bottom $\lceil 1/r \rceil$ fraction of the remaining population, etc.
This continues until every individual of the population has been placed in a bucket.
4. The fitness of each individual is set to the rank assigned to the bucket holding it.

Lexicographic Parsimony Pressure Method: Ratio Bucketing

Realization: The buckets are proportioned, so that low-fitness individuals are placed into larger buckets than high-fitness individuals. A parameter of the method is the bucket ratio $1/r$.

1. The population of size p is sorted by fitness.
2. The bottom $\lceil 1/r \rceil$ fraction of individuals are placed into the worst bucket.
All individuals remaining in the population with the same fitness as the best individual in the bucket are placed in the bucket as well.
3. The same procedure is used to fill in the second worst bucket with the bottom $\lceil 1/r \rceil$ fraction of the remaining population, etc.
This continues until every individual of the population has been placed in a bucket.
4. The fitness of each individual is set to the rank assigned to the bucket holding it.

Characteristics:

- As the remaining population decreases, the $\lceil 1/r \rceil$ fraction decreases as well.
- Higher-ranked buckets hold fewer individuals than lower-ranked buckets.
Thus, the tree-size comparisons are more frequently applied to low-fitness individuals than high-fitness individuals.
- Both bucketing schemes require user-specified bucket parameters b or r that determines how strong an effect of parsimony can have on the selection procedure.

Lexicographic Parsimony Pressure Method: Performance

Plain Lexicographic Parsimony Pressure

- Successful on all problems but the symbolic regression.

The symbolic regression is unusual in that occurrence of individuals of exactly the same fitness in the population is rare since small changes in fitness can be achieved by adding code fragments to the bottom of trees.

Direct bucketing

- Successful on all four problem, but no single setting of the parameter b that would be consistently good across all problems was found.

$b \in \{25, 50, 100\}$ is good for the symbolic regression.

$b = 250$ is the common setting for other problems.

Ratio bucketing

- Nearly uniformly superior in any setting, considering $r = \frac{1}{2}, \frac{1}{3}, \frac{1}{4}, \frac{1}{5}, \frac{1}{6}, \frac{1}{7}, \frac{1}{8}, \frac{1}{9}, \frac{1}{10}$.

Linear Parametric Parsimony Method

Idea: Parsimony pressure methods consider **size as part of the selection process** – a fitness of the program is a function of its quality and size. A fitness of a program is decreased by an amount that depends on its size. The intensity with which bloat is controlled is determined by a parameter called *parsimony coefficient*.

- If it is too small then the control of bloat is not effective.
- If it is too large then the minimization of tree size will become a primary target and fitness will be ignored.

Linear Parametric Parsimony Method

Idea: Parsimony pressure methods consider **size as part of the selection process** – a fitness of the program is a function of its quality and size. A fitness of a program is decreased by an amount that depends on its size. The intensity with which bloat is controlled is determined by a parameter called *parsimony coefficient*.

- If it is too small then the control of bloat is not effective.
- If it is too large then the minimization of tree size will become a primary target and fitness will be ignored.

Realization:

- Linear Parametric Parsimony Method treats the individual's size as a linear factor in fitness

$$g = x \cdot f + y \cdot s$$

where the parameters x and y weight contributions of raw fitness f and the size s to the final fitness g , that is to be minimized.

- Linear Parametric Parsimony Method with a limit applies the size component only if s greater than some specified limit z . Then

$$g = xf, \text{ if } s \leq z$$
$$g = x \cdot f + y \cdot (s - z), \text{ otherwise.}$$

Linear Parametric Parsimony Method

Characteristics:

- A user must set up the *parsimony coefficient* so that it optimally defines f as being worth so many units of s .
 - This can be difficult when the fitness assessment procedure is nonlinear.
Assume a situation where a difference between 0.9 and 0.91 in raw fitness is much more dramatic than a difference between 0.7 and 0.9. Then the size can be given an advantage over the raw fitness when the difference in raw fitness is only 0.01 as opposed to 0.2.
 - Proper setting of the *parsimony coefficient* can be hard when the raw fitness values are converging late in the evolution procedure.
- As $x \rightarrow \infty$, the method approaches Lexicographic Parsimony Pressure.
- In the experiments, parameter x was varied from 1/16 to 65536, doubling x each time.
There are several settings ($x=32, 64, 128, 512, 1024$) for which the method was effective on all four problems.

Double Tournament Method

Characteristics:

- D should be smaller than 2, otherwise it puts too much pressure on parsimony.

Double Tournament Method

Characteristics:

- D should be smaller than 2, otherwise it puts too much pressure on parsimony.
In order to permit D to hold real values between 1.0 and 2.0 the following rule was implemented: Given two individuals, the smaller one wins the tournament with probability $D/2$, else the larger one wins.
Thus, $D = 1$ is random selection, while $D = 2$ is the same as a plain parsimony-based tournament of size 2.
- An individual passes the double tournament if it is generally
low in size and high in fitness.
- $D = 1.4$ was consistently superior on all four problems.

Proportional Tournament Method

Idea: A proportion, P , of tournaments is based on tree size; remaining $1 - P$ tournaments are based on fitness.

Realization:

- Each tournament selection flips a coin to determine which objective to use.

Characteristics:

- Higher values of P imply less of an emphasis on fitness, and vice versa.

$P = 0.0$: All tournaments will select based on fitness.

$P = 0.5$: Tournaments will select based on fitness or size with equal probability.

- An individual passes the proportional tournament if it is generally

low in size or high in fitness.

- $P = 0.2$ was consistently superior on all four problems.

Comparison of Methods: Experiment Setup

General conclusion: The combination of a method with depth limiting was nearly universally superior to either the method alone or depth limiting alone.

Double Tournament and Biased Multi-objective appeared the best across all problem domains when tuned to their optimal per-problem values.

Question: Which method with its single problem-independent setting is the best across all four problems?

Dynamic Operator Equalisation: Calculating the Target Distribution

Target number of individuals in bin b is proportional to the average fitness of individuals within the bin, calculated as

$$bin_capacity_b = round(n \times (\bar{f}_b / \sum_i \bar{f}_i))$$

where

- \bar{f}_i is the average fitness in the bin i ,
- \bar{f}_b is the average fitness of the individuals in bin b , and
- n is the number of individuals in the population.

The target is updated every generation.

Bins with better average fitness will have higher capacity – allowing the population to sample regions where the search proved to be more successful.

Reading

- [Poli08] Poli, R., Langdon, W., McPhee, N.F.: *A Field Guide to Genetic Programming*, 2008.
- [Luke06] Luke, S. and Panait, L.: A Comparison of Bloat Control Methods for Genetic Programming. *Evolutionary Computation*, Volume 14 Issue 3, 2006.
<http://portal.acm.org/citation.cfm?id=1182892.1182897>

