

Computational Complexity Analysis of Simple Genetic Programming On Two Problems Modeling Isolated Program Semantics

Greg Durrett
MIT CSAIL
32 Vassar Street
Cambridge, MA 02139
gdurrett@mit.edu

Frank Neumann
Max-Planck-Institut für
Informatik
Campus E 14, Room 317
66123 Saarbrücken, Germany
fne@mpi-inf.mpg.de

Una-May O'Reilly
MIT CSAIL
32 Vassar Street
Cambridge, MA 02139
unamay@csail.mit.edu

ABSTRACT

Analyzing the computational complexity of evolutionary algorithms (EAs) for binary search spaces has significantly informed our understanding of EAs in general. With this paper, we start the computational complexity analysis of genetic programming (GP). We set up several simplified GP algorithms and analyze them on two separable model problems, ORDER and MAJORITY, each of which captures a relevant facet of typical GP problems. Both analyses give first rigorous insights into aspects of GP design, highlighting in particular the impact of accepting or rejecting neutral moves and the importance of a local mutation operator.

General Terms

Genetic Programming Theory, Computational Complexity, Hill Climbing

1. INTRODUCTION

Because of the complexity of genetic programming (GP) variants and the challenging nature of the problems they address, it is arguably impossible in most cases to make formal guarantees about the number of fitness evaluations needed for an algorithm to find an optimal solution. Current theoretical approaches investigate foundational aspects of GP tangential to this goal, such as schema theories, search spaces, bloat and problem difficulty [13]. However, in this work, we instead choose to follow the path taken for evolutionary algorithms working on fixed-length binary strings. Initial work on pseudo-Boolean functions illustrated the working principles of simple evolutionary algorithms (see e. g. [7, 15, 2]); subsequently, results have been derived for a wide range of classical combinatorial optimization problems such as shortest paths, maximum matchings or minimum spanning trees (see e. g. [9]). These studies have contributed substantially to our theoretical understanding of evolutionary

algorithms for binary representations. Poli et al. [13] state, “we expect to see computational complexity techniques being used to model simpler GP systems, perhaps GP systems based on mutation and stochastic hill-climbing.” This contribution is one fulfillment of this prediction: its goal is to show a GP variant that identifies optimal solutions in provably low numbers of fitness function evaluations for two much simplified, but still relevant, problems that exhibit a few simple aspects of program structure.

The simple parameterized GP algorithm we analyze can succinctly be described as both a hill climber and a randomized algorithm. It has four parametric instantiations we call (1+1) GP-single, (1+1) GP-multi, (1+1) GP*-single, and (1+1) GP*-multi that differ in the acceptance criterion and the size of the mutation proposed. Initially, a solution is chosen at random. We produce by random mutation exactly one offspring of the current solution, and replace the current solution by this proposal as specified by the acceptance criterion. The algorithm iterates until it finds an optimal solution. This simple form of GP algorithm has historical precedent in very early comparisons between Koza-style genetic programming and GP stochastic iterated hill climbing [11, 10, 12], though it does not include a finite bound on fitness evaluations, random restarts or a limit on how many times mutation will be applied to the current solution. Another simplification of the algorithm is that it uses a genetic operator that is as similar to bit-wise mutation as possible. A single bit-wise mutation is the smallest step possible in an binary EA’s search space. Our mutation operator makes the smallest alteration possible to the GP tree while respecting the key properties of the GP tree search space: variable length and hierarchical structure.

The two model problems we select for our analysis are ORDER and MAJORITY, defined exactly as in [3]. We have chosen ORDER and MAJORITY because they make complexity analysis tractable. They allow fitness function evaluation without explicitly executing the program defined by the GP tree. They are minimally sufficient to capture several key properties of GP, including the existence of multiple optimal solutions but they are not real world application problems. Neither are they ad-hoc toy problems intended to demonstrate GP’s strength (such as Boolean multiplexer for classical GP [5] or lawnmower for GP with automatically

defined functions [6]). Each problem has a simple relation to more realistic GP problems: ORDER requires correct ordering as in conditional programs and MAJORITY requires the correct set of solution components.

We proceed as follows: in Section 2, we formally describe the GP variants and the two problems. This requires that we first describe program initialization from a primitive set (2.1) and our mutation operator which is called HVL-Prime (2.2). We then proceed in Sections 3 and 4 with our analyses of ORDER and MAJORITY in terms of the expected number of fitness evaluations until our algorithms have produced a globally optimal solution for the first time. This is called the expected optimization time of the algorithm. Our results are followed by a discussion in Section 5 and conclusions and future work in Section 5.5.

2. DEFINITIONS

2.1 Program Initialization

To use tree-based genetic programming, one must first choose a set of primitives A , which contains a set F of functions and a set L of terminals. Each primitive has explicitly defined semantics; for example, a primitive might represent a Boolean condition, a branching statement such as an IF-THEN-ELSE conditional, the value bound to an input variable, or an arithmetic operation. Functions are parameterized. Terminals are either functions with no parameters, i.e. arity equal to zero, or input variables to the program.

In our derivations, we assume that a GP program is initialized by its parse tree construction. In general, we start with a root node randomly drawn from A and recursively populate the parameters of each function in the tree with subsequent random samples from A , until the leaves of the tree are all terminals. Functions constitute the internal nodes of the parse tree, and terminals occupy the leaf nodes. The exact properties of the tree generated by this procedure will not figure into the analysis of the algorithm, so we do not discuss them in depth.

2.2 HVL-Prime

The HVL-Prime operator is an update of O’Reilly’s HVL mutation operator ([10, 11]) and motivated by minimality rather than inspired from a tree-edit distance metric. HVL first selects a node at random in a copy of the current parse tree. Let us term this the **currentNode**. It then, with equiprobability, applies one of three sub-operations: insertion, substitution, or deletion. Insertion takes place above **currentNode**: a randomly drawn function from F becomes the parent of **currentNode** and its additional parameters are set by drawing randomly from L . Substitution changes **currentNode** to a randomly drawn function of F with the same arity. Deletion replaces **currentNode** with its largest child subtree, which often admits large deletion sub-operations.

The operator we consider here, HVL-Prime, functions slightly differently, since we restrict it to operate on trees where all functions take two parameters. Rather than choosing a node followed by an operation, we first choose one of the three sub-operations to perform. The operations then proceed as shown in Figure 1. Insertion and substitution are exactly as in HVL; however, deletion only deletes a leaf and its parent

to avoid the potentially macroscopic deletion change of HVL that is not in the spirit of bit-flip mutation. This change makes the algorithm more amenable to complexity analysis and specifies an operator that is only as general as our simplified problems require, contrasting with the generality of HVL, where all sub-operations handle primitives of any arity. Nevertheless, both operators respect the nature of GP’s search among variable-length candidate solutions because each generates another candidate of potentially different size, structure, and composition.

In our analysis on these particular problems, we make one further simplification of HVL-Prime: substitution only takes place at the leaves. This is because our two problems only have one generic “join” function specified, so performing a substitution anywhere above the leaves is a vacuous mutation. Such operations only constitute one-sixth of all operations, so this change has no impact on any of the runtime bounds we derive.

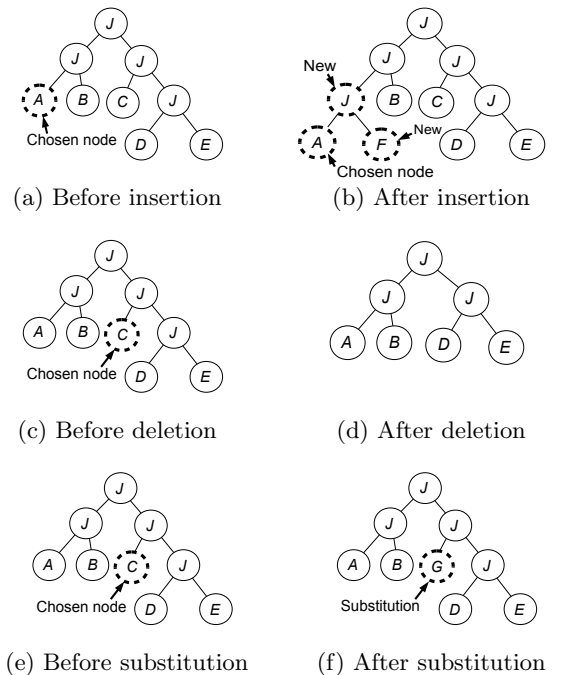


Figure 1: Example of the operators from HVL-Prime.

2.3 Algorithms

We define two genetic programming variants called (1+1) GP and (1+1) GP*. Both algorithms work with a population of size one and produce in each iteration one single offspring. (1+1) GP is defined in Algorithm 1 and accepts an offspring if it is as least as fit as its parent.

ALGORITHM 1 ((1+1) GP).

1. Choose an initial solution X .
2. Set $X' := X$.
3. Mutate X' by applying HVL-Prime k times. For each application, randomly choose to either substitute, insert, or delete.

- If substitute, replace a randomly chosen leaf of X' with a new leaf $u \in L$ selected uniformly at random.
- If insert, randomly choose a node v in X' and select $u \in L$ uniformly at random. Replace v with a join node whose children are u and v , with the order of the children chosen randomly.
- If delete, randomly choose a leaf node v of X' , with parent p and sibling u . Replace p with u and delete p and v .

4. If $f(X') \geq f(X)$, set $X := X'$.

5. Go to 2.

(1+1) GP* differs from (1+1) GP by accepting only solution that are strict improvements (see Algorithm 2).

ALGORITHM 2 (ACCEPTANCE FOR (1+1) GP*).

4'. If $f(X') > f(X)$, set $X := X'$.

For each of (1+1) GP and (1+1) GP* we consider two further variants which differ in using one application of HVL-Prime (“single”) or in using more than one (“multi”). For (1+1) GP-single and (1+1) GP*-single, we set $k = 1$, so that we perform one mutation at a time according to the HVL-Prime framework. For (1+1) GP-multi and (1+1) GP*-multi, we choose $k = 1 + \text{Pois}(1)$, so that the number of mutations at a time varies randomly according to the Poisson distribution.

We will analyze these four algorithms in terms of the expected number of fitness evaluations to produce an optimal solution for the first time. This is called the expected optimization time of the algorithm.

2.4 The ORDER problem

We consider two separable problems called ORDER and MAJORITY that have an independent, additive fitness structure. They both admit multiple solutions on their objective function, which we feel is a key property of a model GP problem because it holds generally for all real GP problems. They also both use the same primitive set:

- $F := \{J\}$, J has arity 2.
- $L := \{x_1, \bar{x}_1, \dots, x_n, \bar{x}_n\}$

x_i is the complement of \bar{x}_i .

ORDER represents problems where the primitive sets include conditional functions, which gives rise to conditional execution paths. GP classification problems, for example, often employ a numerical comparison function (e.g. greater than X , less than X , or equal to X). This sort of function has two arguments (subtrees), one which will be executed only when the comparison returns true, the other only when

it returns false [5]. Thus, a conditional function results in a conditional execution path, so the GP algorithm must identify and appropriately position the conditional functions to achieve the correct conditional execution behavior for all inputs.

ORDER is an abstracted simplification of this challenge in that it determines the conditional path execution of a program by tree inspection rather than execution. Instead of evaluating a condition test and then executing the appropriate condition body explicitly, an ORDER program’s conditional execution path is determined by simply inspecting whether a primitive or its complement occurs first in an in-order leaf parse. Correct programs for the ORDER problem must express each positive primitive x_i before its corresponding complement \bar{x}_i . This correctness requirement is intended to reflect a property commonly found in the GP solutions to problems where conditional functions are used: there exist multiple solutions, each with different conditional execution paths.

ALGORITHM 3 ($f(X)$ FOR ORDER).

1. Derive conditional execution path P of X :

Init: l an empty leaf list, P an empty conditional execution path

1.1 Parse X in order and insert each leaf at the rear of l as it is visited.

1.2 Generate P by parsing l front to rear and adding (“expressing”) a leaf to P only if it or its complement are not yet in P (i.e. have not yet been expressed).

2. $f(X) = |\{x_i \in P\}|$.

For example, for a tree X , with (after the in order parse) $l = (x_1, \bar{x}_4, x_2, \bar{x}_1, x_3, \bar{x}_6)$, $P = (x_1, \bar{x}_4, x_2, x_3, \bar{x}_6)$ and $f(X) = 3$ because $x_1, x_2, x_3 \in P$.

2.5 The MAJORITY problem

MAJORITY is a GP equivalent of the GA OneMax problem [3]. MAJORITY reflects a general (and thus weak) property required of GP solutions: a solution must have correct functionality and no incorrect functionality. Like ORDER, MAJORITY is a simplification that uses tree inspection rather than program execution. A correct program in MAJORITY must exhibit at least as many occurrences of a primitive as of its complement and it must exhibit all the positive primitives of its terminal (leaf) set. Both the independent sub-solution fitness structure and inspection property of MAJORITY are necessary to make our analysis tractable.

ALGORITHM 4 ($f(X)$ FOR MAJORITY).

1. Derive the combined execution statements S of X :

Init: l an empty leaf list, S is an empty statement list.

1.1 Parse X in order and insert each leaf at the rear of l as it is visited.

1.2 For $i \leq n$: if $\text{count}(x_i \in l) \geq \text{count}(\bar{x}_i \in l)$ and $\text{count}(x_i \in l) \geq 1$, add x_i to S

2. $f(X) = |S|$.

For example, for a tree X , with (after the inorder parse) $l = (x_1, \bar{x}_4, x_2, \bar{x}_1, \bar{x}_3, \bar{x}_6, x_1, x_4)$, $S = (x_1, x_2, x_4)$ and $f(X) = 3$.

3. ANALYSIS FOR ORDER

Here we present bounds for ORDER on the number of run-time evaluations needed in the execution of (1+1) GP and (1+1) GP*.

We will analyze this GP problem using fitness-based partitions [2]. This requires us to compute the probability of improving the fitness from k to $k + 1$ for each value of k between 0 and $n - 1$, inclusive. Although our HVL-Prime operator is complex, we can obtain a lower bound on the probability of making an improvement by considering fitness improvements that arise from insertions. This is described in the following lemma.

LEMMA 1. Define p_k to be the probability that we perform an insertion that improves the fitness value of the GP tree from k to $k + 1$. For the single- and multi-operation variants of (1+1) GP and (1+1) GP* applied to the ORDER problem,

$$p_k = \Omega\left(\frac{(n-k)^2}{n \max\{T, n\}}\right)$$

where n is the number of variables and T is the number of leaves in the GP tree at the particular iteration.

PROOF. When the fitness value is k , it must be the case that k different x_i appear before their corresponding \bar{x}_i . To improve the fitness, we must insert one of the $n - k$ unexpressed x_i as a leaf that will be visited before a leaf containing the corresponding \bar{x}_i . Assume for notational ease that these unexpressed x_i are indexed by $\{x_1, \dots, x_{n-k}\}$. Define A_i to be the event that we insert x_i into the tree with our mutation operation, and define B_i to be the event that x_i is inserted before the corresponding \bar{x}_i . Given this, we can write out p_k as follows.

$$p_k = \sum_{i=1}^{n-k} \Pr(A_i) \Pr(B_i | A_i)$$

With a single operation, the probability of choosing to insert a particular x_i is $\frac{1}{6n}$, since we choose to insert with probability $\frac{1}{3}$ and select the variable uniformly at random from the set of $2n$ possible terminals. We can cover the multi-operation case with this analysis as well because the number of operations is sampled according to $1 + \text{Pois}(1)$, so the probability of performing exactly one operation is $\frac{1}{e}$. The probability of A_i is therefore at least $\frac{1}{6en}$, so in both the single- and multi-operation cases, we have

$$p_k \geq \frac{1}{6en} \sum_{i=1}^{n-k} \Pr(B_i | A_i)$$

We need to analyze two cases in computing this sum. Preliminarily, we define S to be the total number of nodes in the GP tree. Note that $S = 2T - 1$, so $S = \Theta(T)$.

Case 1: $T \geq n - k$. We first note that the probability of inserting x_i such that it is visited between the $j - 1$ st leaf and the j th leaf in the traversal is at least $\frac{1}{2S}$, since we choose to insert at the j th leaf with probability $\frac{1}{S}$ and then add x_i as a left child of the new join node with probability $\frac{1}{2}$.

Inserting any of the $n - k$ unexpressed x_i before the first leaf in the tree clearly improves the fitness. If we insert at the second position instead, there must still be at least $n - k - 1$ choices of x_i that yield an improvement: there is only one node that will be traversed before this position in the tree, so there is at most one \bar{x}_i expressed before this position. We can iterate this argument to see that at the i th position, there are still $n - k - i + 1$ x_i that can be inserted for an improvement to the fitness. By reindexing the x_i , we then have that x_i can be inserted in at least the first i positions in the tree. Using the fact that the number of leaves T is at least $n - k$, we have that

$$\begin{aligned} p_k &\geq \frac{1}{6en} \sum_{i=1}^{n-k} \Pr(B_i | A_i) \\ &\geq \frac{1}{6en} \sum_{i=1}^{n-k} \frac{n-k-i+1}{2S} = \frac{1}{6en} \sum_{i=1}^{n-k} \frac{i}{2S} \\ &= \frac{1}{6en} \frac{(n-k)(n-k+1)}{4S} \geq \frac{1}{6en} \frac{(n-k)^2}{\max\{4S, n\}} \end{aligned}$$

Noting that $S = \Theta(T)$, the asymptotic result follows.

Case 2: $T < n - k$: We can apply the argument of Case 1 up to the T th position. After this, we have that for $n - k - T + 1$ of the unexpressed x_i , the corresponding \bar{x}_i appears nowhere in the tree, so the probability of an insertion improving the fitness is 1. We also note that $S < 2n$ in this case, allowing us to simplify our expression for p_k as follows.

$$\begin{aligned} p_k &\geq \frac{1}{6en} \sum_{i=1}^{n-k} \Pr(B_i | A_i) \\ &\geq \frac{1}{6en} \left[\sum_{i=1}^T \frac{i}{2S} \right] + \frac{1}{6en} \sum_{i=T+1}^{n-k} 1 \\ &\geq \frac{1}{6en} \left[\sum_{i=1}^T \frac{i}{4n} \right] + \frac{n-k-T}{6en} \\ &= \frac{T(T+1)}{24en^2} + \frac{(n-k-T)(n-k)}{6en(n-k)} \end{aligned}$$

If $T = \Omega(n - k)$, then we lower-bound p_k using only the first term, which behaves asymptotically in this case as $\Omega\left(\frac{(n-k)^2}{n^2}\right)$. Otherwise, if $T = o(n - k)$, then we use the second term, which then grows according to $\Omega\left(\frac{(n-k)^2}{n^2}\right)$. In either case, because T is less than n , we have the desired asymptotic behavior. \square

With this lemma, we can now state the general theorem about the number of fitness evaluations needed for our (1+1) GP variants.

THEOREM 1. The expected optimization time of the single- and multi-operation cases of (1+1) GP and (1+1) GP* on

ORDER is $O(nT_{\max})$ in the worst case, where n is the number of x_i and T_{\max} denotes the maximal tree size at any stage during the evolution of the algorithm.

PROOF. We can apply Lemma 1 to these algorithms, which implies an asymptotic lower bound on p_k , the probability of improving the fitness from k to $k + 1$ via an insertion. This also serves as an asymptotic lower bound on the probability of improving the fitness at all, and therefore provides an expected time necessary to improve the fitness, regardless of whether or not we accept neutral moves. In order to determine the total number of evaluations, we must sum the expected number of fitness function evaluations over all intermediate fitness values, from $k = 0$ to $k = n - 1$.

The expected optimization time is therefore upper bounded by

$$\begin{aligned} \sum_{k=0}^{n-1} \frac{1}{p_k} &= \sum_{k=0}^{n-1} O\left(\frac{n \max\{T_k, n\}}{(n-k)^2}\right) \\ &= nT_{\max} \sum_{k=0}^{n-1} O\left(\frac{1}{(n-k)^2}\right) \\ &= O\left(nT_{\max} \sum_{j=1}^{\infty} \frac{1}{j^2}\right) \\ &= O(nT_{\max}) \end{aligned}$$

where the second equality follows from the fact that $T_{\max} \geq T_n \geq n$, and the last equality follows from the fact that $\sum_{j=1}^{\infty} \frac{1}{j^2} \leq 2$. \square

Note that most GP algorithms explicitly limit the maximum tree size that can be used in an algorithm. Choosing a linear maximum tree size that would still allow us to generate an optimal solution, i.e. a tree with at least n leaves, gives an algorithm that solves the ORDER problem in expected time $O(n^2)$. However, it is also sometimes possible to show that the tree does not get too big during the optimization process. We examine this for (1+1) GP*-single and present an upper bound on the expected optimization time.

COROLLARY 1. *The expected optimization time of (1+1) GP*-single on ORDER is $O(n^2)$ if the tree is initialized with $O(n)$ terminals.*

PROOF. We note that the maximum value of the fitness is n , and the fitness is integer-valued, so if it is strictly increasing with each operation that is accepted, there must be no more than n operations accepted. In the single-operation framework, each operation adds at most two nodes to the tree (if it is an insertion), which means that $T_{\max} \leq O(n) + 2n = O(n)$ holds during the run of the algorithm. \square

The case of (1+1) GP*-multi is more difficult to analyze because the expected length of accepted moves may be very different from the expected length of proposed moves, as conditioning on accepting the move will skew the distribution. We conjecture that the bound from Corollary 1 holds in this case as well, but do not present a proof of this.

We also note that because of how our fitness-based partition argument is structured, invoking the average case does not enable us to find any better bounds. Although k will initially be somewhat greater than zero, we will generally still need to improve the fitness $\Theta(n)$ times, so we will have the same asymptotic result.

4. ANALYSIS FOR MAJORITY

We next consider the MAJORITY problem. We start with some preliminary definitions.

DEFINITION 1. *For a given GP tree, let $c(x_i)$ be the number of x_i variables and $c(\bar{x}_i)$ be the number of negated x_i variables present in the tree. For a GP tree representing a solution to the MAJORITY problem, we define the deficit in the i th variable by*

$$D_i = c(\bar{x}_i) - c(x_i).$$

DEFINITION 2. *In a GP tree for MAJORITY, we say that x_i is expressed when $D_i \leq 0$ and $c(x_i) > 0$.*

The fitness of a tree T is simply the number of variables that are expressed.

We note a property of HVL-Prime for this particular problem that we will make use of later.

DEFINITION 3. *The substitution decomposability property (SDP) for MAJORITY states that a substitution is exactly equivalent to a deletion followed by an insertion, which are accepted or rejected as a unit.*

This property follows from the fact that the order of the terminals has no bearing on the fitness of a solution for MAJORITY. The variable to be replaced by substitution is selected uniformly at random from the set of leaves of the tree. This is identical to how the variable to be deleted is chosen when using the deletion operator. Substitution then inserts a variable selected uniformly at random from the set of possible terminals, just as the insertion operator does.

We begin our analysis, in 4.1, with worst case bounds for (1+1) GP-single, (1+1) GP*-single, and (1+1) GP*-multi. (1+1) GP-single solves the problem quite efficiently, yielding polynomial-time worst-case complexity. However, not accepting neutral moves, as in (1+1) GP*, results in poor performance: (1+1) GP*-single fails to terminate in the worst case, and (1+1) GP*-multi requires a number of fitness evaluations exponential in the size of the initial tree.

In 4.2 we derive average case bounds that assume the initial solution tree has $2n$ terminals each selected uniformly at random from L . This random tree initialization allows us to bound the maximum deficit in any variable. We show that (1+1) GP-single runs in time $O(nT_{\max} \log \log(n))$ in the average case. By contrast, (1+1) GP*-single has a constant probability of failing to terminate, and so the expected runtime is infinite.

4.1 Worst Case Bounds

4.1.1 (1+1) GP-single

We will show here some properties of (1+1) GP-single on MAJORITY and give a polynomial-time worst-case bound on the performance. Our analysis considers the evolution of the deficits D_i over the course of the algorithm as n parallel random walks. We will show that each positive D_i reaches zero at least as quickly as a balanced random walk, which is the condition for the corresponding x_i to be expressed; this, then, gives us the expected number of operations that we are required to perform on a particular variable before it is expressed. Because these arguments do not easily extend to (1+1) GP-multi, we omit from this section any treatment of that case.

We begin by establishing the validity of modeling the temporal sequence of each of the D_i as a random walk.

LEMMA 2. *For (1+1) GP-single on MAJORITY:*

- a) *The probability of proposing an operation that changes either the number of x_i or the number of \bar{x}_i is $\Omega(\frac{1}{n})$.*
- b) *If some x_i has a deficit $D_i = d > 0$, we require in expectation $O(dT_{\max})$ proposed operations involving that variable before it is successfully expressed, where T_{\max} is the maximum number of nodes in our GP tree at any timestep of the algorithm.*

PROOF. a) To see that a particular operation involves x_i or \bar{x}_i with probability $\Omega(\frac{1}{n})$, we simply note that the probability of inserting one of the two variables is $\frac{1}{3} \times \frac{2}{2n} = \Omega(\frac{1}{n})$.

b) We address each of the three types of operations in turn and show that each is at least as favorable as a balanced random walk in terms of reducing D_i to zero.

Insertion: The probability of inserting x_i into the tree is $\frac{1}{6n}$, which is the same as the probability of inserting \bar{x}_i . Therefore, given that we change D_i with an insertion, we increase it or decrease it in a balanced manner, with probability $\frac{1}{2}$.

Deletion: The probability of a deletion changing D_i is

$$\frac{c(x_i) + c(\bar{x}_i)}{T}$$

where T is the size of the GP tree. Given that we do such a deletion, we increase D_i with probability $\frac{c(x_i)}{c(x_i) + c(\bar{x}_i)}$ and decrease it with probability $\frac{c(\bar{x}_i)}{c(x_i) + c(\bar{x}_i)}$, since we pick the variable to delete uniformly at random. However, note that because $D_i > 0$, we have that $c(x_i) < c(\bar{x}_i)$, so the probability of decreasing D_i is greater than the probability of increasing it, so this is actually slightly better than a balanced random walk for the purpose of reducing D_i .

Substitution: We now make use of the substitution decomposability property (SDP) defined previously to observe that substitution consists of a deletion followed by an insertion. Therefore, a substitution is simply equivalent to taking one or two steps that tend to reduce D_i with probability at least $\frac{1}{2}$ if D_i is greater than 0.

Consider the 1-dimensional random walk on the integers $0, 1, \dots, n$, with n being a reflecting barrier and 0 being an absorbing barrier. The expected time to reach 0 when starting at k is $O(kn)$, following the analysis for random walks on undirected graphs carried out in [1]. This is precisely the setting we have for our random walk on the D_i if we set $k = d$ and $n = T_{\max}$, so we have that the random walk performed by the D_i reaches zero after at most $O(dT_{\max})$ accepted operations.

We now must address the question of how many operations on the variable must be proposed in order to accept $O(dT_{\max})$ of them. Note that if x_i is unexpressed, any insertion or deletion affecting D_i will be accepted, since it cannot possibly decrease the fitness value. The probability of a substitution affecting D_i is, by the SDP, less than or equal to the probability that an insertion affects D_i plus the probability that a deletion affects D_i . Therefore, even if every substitution is rejected, we still accept a constant fraction of proposed operations that affect D_i , so we only require $O(dT_{\max})$ proposed operations involving x_i and \bar{x}_i in order to have $O(dT_{\max})$ accepted operations.

Once D_i reaches zero, we are done and x_i is expressed unless $c(x_i) = c(\bar{x}_i) = 0$. In this case, we clearly cannot do any more deletes, but will either add in an x_i or \bar{x}_i via an insertion or a substitution. Through either operation, we add each variable with probability $\frac{1}{2}$, and therefore successfully express x_i with probability $\frac{1}{2}$. In the case where we insert an \bar{x}_i and increase D_i to one, we again apply our one-dimensional random walk result with $k = 1$ and $n = T_{\max}$ to see that we will return to zero again after only $O(T_{\max})$ additional moves, whereupon either x_i is present in the tree and we are done or we can once again attempt to add it. Because we expect to do this procedure only twice before succeeding, it only adds $O(T_{\max})$ steps, and therefore does not change our bound of $O(dT_{\max})$. \square

This lemma allows us to establish an upper bound on the number of evaluations for (1+1) GP on MAJORITY given a bound on the largest deficit.

THEOREM 2. *Let $D = \max_i D_i$ for an instance of MAJORITY initialized with T terminals drawn from a set of size $2n$ (i.e. terminals $x_1, \dots, x_n, \bar{x}_1, \dots, \bar{x}_n$). Then the expected optimization time of (1+1) GP-single is*

$$O(n \log n + DT_{\max} n \log \log n)$$

in the worst case.

PROOF. We draw upon a result from Myers and Wilf [8] about a generalized form of the coupon collector problem. If we have n coupons and wish to acquire at least k of each coupon, we need to draw, in expectation, $O(n \log n + kn \log \log n)$ coupons. When k is at least $\log n$, this is a slight improvement over the naive bound of $O(kn \log n)$ from simply iterating the basic coupon collector problem k times.

Lemma 2 tells us two things. Firstly, we have that a proposed operation involves x_i or \bar{x}_i with probability $\Omega(\frac{1}{n})$, so we have a coupon collector problem with slightly perturbed

coupon probabilities. Secondly, we find that we need to propose $O(DT_{\max})$ operations involving each terminal in order to express all of our variables. Plugging these into the bound described above yields an asymptotic requirement of $O(n \log n + DT_{\max} n \log \log n)$ fitness function evaluations, as desired.

The only wrinkle in this picture is that the coupon collector assumes that a variable is “complete” after a set number of coupons have been collected. While we do not accept moves that reduce the fitness value, an expressed variable x_i could become unexpressed if, during the course of a substitution operation, another variable x_j were simultaneously expressed. However, in this case, we must have had $D_i = 0$ and $D_j = 1$, and we have merely reversed the two, which amounts to a relabeling of the x_i and x_j . Because the D_i are the only state variables that we care about in this case, this move effectively does nothing except cause us to make a vacuous move. Because substitutions only make up $\frac{1}{3}$ of all of the proposed moves, such wasted moves can only make up a constant fraction of the total number of moves, and therefore do not change the asymptotics. \square

As a corollary of 2 we can bound the initial D by considering tree initialization.

COROLLARY 2. *When MAJORITY is initialized with $m = O(n)$ terminals drawn from a set of size $2n$, the expected optimization time of (1+1) GP-single is*

$$O(n^2 T_{\max} \log \log n)$$

PROOF. This follows from Theorem 2 with $D = m$, since the deficit cannot be greater than the number of terminals in the tree. \square

We can consider the outcome of the worst case tree initialization both intuitively and experimentally. We have $D = m$ when all of the leaves consist of instances of one bar variable, say \bar{x}_1 . Since the \bar{x}_1 occupy such a large fraction of the tree, they will frequently be substituted out or deleted. This suggests that the balanced random walk argument is quite pessimistic given this circumstance. We thus expect that, in practice, this initial condition will be quickly erased. If we put $T_{\max} = O(n)$, we know, from the coupon collector problem, that after an initial phase of $O(n \log n)$ steps, we will have proposed a deletion on every leaf that was initialized in the GP tree. Because deletions are always accepted on negated variables, we will have deleted all of the initial \bar{x}_1 variables by the end of this “erasure” phase, and only expect to introduce at most $O(\log n)$ of any particular bar variable through insertions and substitutions. This implies that, after this relatively short phase, $D = O(\log n)$, giving an optimization time of $O(n^2 \log n \log \log n)$, a bound very close to the average-case optimization time we present in 4.2.1.

We experimented with this initialization to confirm our intuition. Figure 2 shows the results of solving MAJORITY using (1+1) GP-single with increasing problem size and trees initialized with $2n$ leaves, each occupied by \bar{x}_1 . We tracked

the number of fitness evaluations required and, even though we imposed no bound on the tree size, the order of growth relative to n appears to be just barely superlinear. This empirical evidence supports the intuition that the worst-case performance is much closer to the average-case than Corollary 2 would suggest.




Figure 2: Plot of the average optimization time given a “bad” initialization of size $2n$ with $D_1 = 2n$ for x_1 . Fifty trials were used to compute each point. Circles indicate the mean number of fitness function evaluations for each value of n , and error bars show the standard deviation of the 50 trials.

4.1.2 (1+1) GP*

Unlike in the case of the ORDER problem, where accepting or not accepting neutral moves makes no difference in the performance of the algorithm, for MAJORITY, such a distinction matters tremendously. Intuitively, this behavior arises because there is a notion of “working towards” a solution here that is absent from the ORDER problem. In ORDER, our analysis relied on our ability to express an x_i by simply inserting it as a terminal early enough in the tree, which required only one step. However, in MAJORITY, if there are k x_i and ℓ \bar{x}_i present in the GP tree, at least $\lceil \frac{\ell-k}{2} \rceil$ mutation operations will be required to make up this deficit, all but the last of which will be neutral moves.

Because of the importance of neutral moves, we find that (1+1) GP*-single and (1+1) GP*-multi perform quite badly. Even when we initialize with a tree with size linear in n , the number of terminal symbols, we can demonstrate an initialization where (1+1) GP*-single never terminates and (1+1) GP*-multi takes an exponential amount of time to do so. Consider the tree T_{bad} which has as leaves the variables

$$x_1, x_2, x_3, \dots, x_{n-1}, \underbrace{\bar{x}_n, \bar{x}_n, \dots, \bar{x}_n}_{n+1 \text{ of these}}$$

THEOREM 3. *Let T_{bad} be the initial solution to MAJORITY. Then the expected optimization time of (1+1) GP*-single is infinite.*

PROOF. It is clear that, with one move, the deficit in x_n can only be changed by at most two. There is a deficit of $n+1$ to make up, which is impossible, therefore (1+1) GP*-single will never find its way out of this local optimum. \square

THEOREM 4. *Let T_{bad} be the current solution to MAJORITY. Then the expected optimization time of (1+1) GP*-multi is at least exponential in n .*

PROOF. The fitness value of T_{bad} is $n-1$, with x_1 through x_{n-1} expressed, so the only way to improve the fitness is to make a move that expresses x_n . Therefore, the moves that achieve this are the only moves that will be accepted. We compute the probability of making such a move in this configuration in order to determine the expected time to make such a move.

Note that any mutation operation that successfully improves the fitness must make up for a deficit of $n+1$, which requires at least $\lceil \frac{n+1}{2} \rceil$ operations, assuming that we, in each case, substitute an \bar{x}_n with an x_n . The number of moves per mutation is distributed as $1+\text{Pois}(1)$, so the Poisson random variable must take a value of at least $\lceil \frac{n-1}{2} \rceil$. The probability of this when $\lambda = 1$ is given by

$$\sum_{i=\lceil \frac{n-1}{2} \rceil}^{\infty} \frac{e^{-1}}{i!} \leq \frac{1}{\left(\frac{n-1}{2}\right)!} = O\left(\left(\frac{n}{2e}\right)^{-\frac{n}{2}}\right)$$

by Stirling's formula.

We can take this probability as a (very weak) upper bound on the probability of improving the fitness. Inverting it, we see that the expected number of moves required is $\Omega\left(\left(\frac{n}{2e}\right)^{\frac{n}{2}}\right)$.

4.2 Average Case Bounds

To provide average case bounds we consider a GP tree which is initialized with what we term "unity expectation": it has $2n$ terminals (leaves) each selected uniformly at random from the set of possible terminals.

4.2.1 (1+1) GP

The average case bound follows more or less directly from Theorem 2 once a result from the literature is applied to given an expected bound on the maximum initial deficit.

COROLLARY 3. For MAJORITY with a terminal set of size $2n$ under unity expectation initialization, the expected optimization time of (1+1) GP-single is $O(nT_{\max} \log n)$.

PROOF. A result from Raab and Steger [14] tells us that, with probability at least $1 - O\left(\frac{1}{n^k}\right)$ for any integer k , no \bar{x}_i appears more than $O\left(\frac{\log n}{\log \log n}\right)$ times in the GP tree, so $D = O\left(\frac{\log n}{\log \log n}\right)$. Set $k = 2$, so that the probability of having a larger deviation is $O\left(\frac{1}{n^2}\right)$. The worst-case bound of $O(n^2 T_{\max} \log \log n)$ from Corollary 2 ensures that these uncommon cases contribute only an $O(T_{\max} \log \log n)$ term to the expectation. Substituting $D = O\left(\frac{\log n}{\log \log n}\right)$ into the expression in Theorem 2 gives us the desired bound for the common case, which is also the overall runtime bound. \square

4.2.2 (1+1) GP*

Assuming unity expectation initialization, we can improve on our result from 4.1.2 and show (1+1) GP*-single has a constant probability of failing to terminate. Our general strategy will be to prove that there is constant probability that, when starting with a deficit of size three in x_1 , this deficit will be preserved until the fitness is $n-1$. At this

point, when all the other variables are expressed, there will remain a gap that cannot be closed in a single step. Such a deficit could disappear over the course of the algorithm because substitution has the ability to shrink the deficit (by removing \bar{x}_1 and replacing it with a x_i in order to express that variable), but this proof shows that there is nonetheless a constant probability of the deficit being preserved.

First, we establish a lemma about the prevalence of constant-size deficits arising based on our initialization.

LEMMA 3. Suppose we have a $2n$ -length instance of the MAJORITY problem with unity expectation initialization. Let A_k denote the event that \bar{x}_1 appears exactly k times without x_1 appearing at all, where k is any constant. Then $\Pr(A_k) = \Omega(1)$.

PROOF. To compute $\Pr(A_k)$, we count the number of $2n$ -length sequences of terminals for which this is true and divide by the total number of possible sequences. Under A_k , we must have k instances of \bar{x}_1 and zero instances of x_1 , so there are $\binom{2n}{k}$ positions that can be occupied by the \bar{x}_1 and \square the remaining $2n-k$ positions should each be occupied by one of the $2n-2$ elements that are not x_1 or \bar{x}_1 . In total, there are $(2n)^{2n-k}$ possible $2n$ -length sequences of terminals. Combining these facts yields

$$\begin{aligned} \Pr(A_k) &= \frac{\binom{2n}{k} (2n-2)^{2n-k}}{(2n)^{2n}} \\ &= \frac{1}{k!} \frac{(2n)!}{(2n-k)!(2n-2)^k} \left(\frac{2n-2}{2n}\right)^{2n} \\ &\geq \frac{1}{k!} \left(\frac{2n-k}{2n-2}\right)^k \left(1 - \frac{1}{n}\right)^{2n} \\ &= \frac{1}{k!} \times \Omega(1) \times \Omega(1) \\ &= \Omega(1) \end{aligned}$$

assuming that k is a constant. \square

Next, we lower-bound the size of the GP tree when running (1+1) GP*-single on MAJORITY. The tree must be large enough so that we are not too likely to substitute out the \bar{x}_1 over the course of the algorithm.

LEMMA 4. Using (1+1) GP-single on MAJORITY with any initialization of size $2n$, the size of the GP tree is always greater than $\frac{7n}{6}$ with probability one.

PROOF. A deletion can only improve the fitness if we delete some \bar{x}_i when $c(x_i) = n$ and $c(\bar{x}_i) = n+1$, with n positive. Such a configuration requires at least three occurrences of x_i and \bar{x}_i in the GP tree, so at most $\frac{2n}{3}$ variables can be present in this fashion initially. Of the at least $\frac{n}{3}$ variables that remain, at most half can be expressed by a deletion, because they must be first put into this configuration during the course of a substitution that expresses some other variable x_j . Therefore, we are forced to accept at least $\frac{n}{6}$ insertions or substitutions over the course of the algorithm, giving us an upper bound of $\frac{5n}{6}$ on the number

of deletions accepted. This in turn guarantees that our tree always remains larger than $2n - \frac{5n}{6} = \frac{7n}{6}$. \square

Finally, we can prove the claim directly.

THEOREM 5. *With probability $\Omega(1)$, the optimization time of (1+1) GP*-single with unity expectation initialization is infinite on MAJORITY.*

PROOF. Lemma 3 tells us that, with a constant probability, we initialize one of the variables, say x_1 , with $c(x_1) = 0$ and $c(\bar{x}_1) = 3$. We now show that, also with a constant probability, such a deficit is preserved during the course of the expression of at most $n - 1$ of the other variables.

We make such an argument by induction. Define the j th step of the algorithm as the period after j variables have been expressed, at the end of which we propose the move that expresses the $j + 1$ st move. Suppose that at the j th step, it is true that $c(x_1) = 0$ and $c(\bar{x}_1) = 3$. The $j + 1$ st variable expressed cannot possibly be x_1 , since there is no way to make up a deficit of three with a single move. If the move we accept to express the $j + 1$ st variable is an insertion or a deletion, we preserve our deficit of three and do not change the state of the variable x_1 at all, since we must either insert some variable in the set $\{x_2, \dots, x_n\}$ or delete some variable in the set $\{\bar{x}_2, \dots, \bar{x}_n\}$.

If we express the $j + 1$ st variable with a substitution, however, it is possible that we might insert an x_1 or delete one of the \bar{x}_1 . An accepted substitution must either replace some variable with a variable in the set $\{x_2, \dots, x_n\}$ or substitute out some $\{\bar{x}_2, \dots, \bar{x}_n\}$. However, a substitution also involves an “extraneous” insertion or deletion, by the SDP. If this operation impacts a variable different than the $j + 1$ st variable we are expressing, it must be an operation that, on its own, would keep the fitness constant. For an extraneous insertion, we note that it is always admissible to insert any of the n symbols in the set $\{x_1, x_2, \dots, x_n\}$ without decreasing the fitness. Therefore, the probability of inserting neither x_1 nor \bar{x}_1 in a neutral or better move is at least $1 - \frac{2}{n}$.

If the extraneous operation is a deletion, we note that it is always possible to delete at least $\frac{T-n}{2}$ terminals, where T is the current tree size. Any variable expressed with $c(x_i) = 1$ and $c(\bar{x}_i) = 0$ cannot be removed, so there might be as many as n terminals forbidden for this reason. For any variable not in this configuration, we have one of two cases. If the variable is unexpressed or is expressed with a deficit less than or equal to -1, any occurrence of x_i or \bar{x}_i can safely be deleted without decreasing the fitness. If the variable is expressed with a deficit of zero, there must be at least as many \bar{x}_i as there are x_i , and any of these \bar{x}_i can be safely deleted. Therefore, we set aside at most n “singleton” symbols that cannot be deleted, and of those remaining, it must always be acceptable to delete at least half, yielding $\frac{T-n}{2}$.

We therefore preserve our three \bar{x}_1 variables with probability at least

$$\frac{\frac{T-n}{2} - 3}{\frac{T-n}{2}} = 1 - \frac{6}{T-n}$$

We now invoke the result from Lemma 4. Because the size of the tree is at least $\frac{7n}{6}$ at all times, we can lower-bound the probability of preserving the \bar{x}_1 as $1 - \frac{36}{n}$.

These situations (extraneous inserts and extraneous deletes) are mutually exclusive, and of the two, the deletes are the more probable to interfere with our x_1 setup. Nevertheless, the probability of preserving our deficit of three in x_1 from the j th step to the $j + 1$ st step is at least $1 - \frac{O(1)}{n}$. Because there are at most $n - 1$ such steps of the algorithm, our overall probability of preserving the deficit is

$$\left(1 - \frac{O(1)}{n}\right)^{n-1} = \Omega(1)$$

We have constant probability of initializing with such a deficit, and a constant probability of preserving the deficit, in which case the algorithm never terminates. Therefore, with constant probability, (1+1) GP*-single never terminates on MAJORITY. \square

COROLLARY 4. *Using unity expectation initialization, the expected optimization time of (1+1) GP*-single on MAJORITY is infinite.*

PROOF. This follows directly from Theorem 5 \square

While Theorem 5 does demonstrate that the probability of getting stuck in a local optimum is at least a constant, the actual constant yielded by the proof is rather small. However, our proof technique made several very conservative assumptions for simplicity. To investigate further, we tried to solve MAJORITY with (1+1) GP*-single experimentally, observing when we would get stuck in a local minimum. Experimentally, the actual probability of (1+1) GP*-single failing to converge to the optimum is actually quite high, as demonstrated by Figure 3.




Figure 3: Plot showing the probability of (1+1) GP*-single failing to terminate on MAJORITY when the initial solution tree has $2n$ terminals each selected uniformly at random, i.e. with unity expectation initialization. Each probability was determined empirically over the course of 100 simulations for each value of n .

However, we do not know how to show a similar result for (1+1) GP*-multi. We note that difficult MAJORITY instances, such as T_{bad} presented in 4.1.2, are exponentially unlikely to occur when the initial solution tree has $2n$ terminals each selected uniformly at random. From Raab and

Steger [14], we know that deficits larger than logarithmic occur with exponentially small probability, and in any case large deficits should tend to equalize over the course of the algorithm execution, even if we only accept a linear number of moves. However, if the last unexpressed variable has a deficit of size k , we will require at least $\Omega\left(n^{-\frac{k}{2}}\right)$ steps to correctly substitute out enough instances of \bar{x}_i for x_i even in the best case, so unless k can be bounded at a constant, we will have an expected runtime that is superpolynomial.

5. SUMMARY AND DISCUSSION

Table 1 aggregates our expected optimization time results for all algorithm variants and each problem.

ORDER		
	(1+1) GP	(1+1) GP*
single	$O(nT_{\max})$ w.c. †	$O(n^2)$ w.c.
multi	$O(nT_{\max})$ w.c. †	$O(nT_{\max})$ w.c. †
MAJORITY		
	(1+1) GP	(1+1) GP*
single	$O(n^2T_{\max} \log n)$ w.c. † $O(nT_{\max} \log n)$ a.c.	$\Omega(\infty)$ a.c.
multi	?	$\Omega\left(\left(\frac{n}{2e}\right)^{\frac{n}{2}}\right)$ w.c.

Table 1: Results of the computational complexity analysis for our sample problem. We use w.c. to denote a worst-case bound and a.c. to denote an average-case bound. The daggers indicate where we conjecture that better bounds exist.

From the perspective of a GP practitioner, the insights provided by this rigorous analysis may be more valuable than the complexity results themselves. In 5.1 we discuss how our treatment sheds light upon the important but subtle interactions between a problem, the acceptance criterion, and the genetic operator. In 5.2, we discuss the impact of the sub-operations on the mutation operator we considered. In 5.3, we address the implications of our design and analysis methodology for practical GP algorithm design. Section 5.4 covers some of our analysis techniques, and finally 5.5 presents future work avenues and concludes.

5.1 Accepting Neutral Moves in ORDER and MAJORITY

It might initially seem immaterial whether or not we accept neutral moves with our genetic operator for ORDER and MAJORITY. However, our analysis provides rigorous evidence that the differences in performance between (1+1) GP and (1+1) GP* are substantial for both of these problems. Similar results have already been obtained in the context of evolutionary algorithms for binary representations [4].

ORDER’s focus on condition semantics gives it the property that only the first occurrence of each terminal matters. A large tree makes the probability of improvements smaller because many of the mutations will change variables that have no effect on expression, being sequentially later than earlier occurrences of those same variables. Therefore, not accepting neutral moves helps prevent “bloat” and using (1+1) GP*

is significantly advantageous. (1+1) GP’s acceptance of neutral moves causes a feedback loop that stimulates growth of the tree: there is a slight bias towards accepting insertions as opposed to deletions, which makes the tree large, which increases the time to find an improvement and results in many neutral insertions, which increase the tree size even more. In general, to solve ORDER with runtime performance that respects the complexity analysis, the tree must not grow too large, and not accepting neutral moves assures this.

Solving MAJORITY, we see the opposite effect: (1+1) GP very handily beats (1+1) GP* in terms of expected optimization time. Neutral moves have the effect of balancing both the relative frequency of variables and the number of positive versus negative occurrences. This draws us toward a very favorable average case where every variable is either expressed or very close to being expressed. If neutral moves are not accepted, improvement can frequently stagnate, underscoring the fact that there are large flat regions in the search space. (1+1) GP is better equipped to escape these than (1+1) GP*, so one should, in fact, clearly choose (1+1) GP so that there is a guarantee of termination and to avoid the exponential-time worst-case associated with (1+1) GP*-multi.

Overall, these results highlight the fact that, in choosing whether or not to accept neutral moves, one should consider their general effect with respect to both the fitness landscape and growth in tree size. Tying this knowledge into expected optimization time also requires an understanding of the mechanisms by which the fitness increases. We recognize that for ORDER and MAJORITY, this is much easier to do rigorously than for more realistic problems. However, perhaps our exercise with ORDER and MAJORITY can provide intuitive insight to GP practitioners.

5.2 Mutation

Our results also tell us more about our HVL-Prime framework, and show that it has several interesting properties and behaves quite differently for the two problems.

Interestingly, the analysis for ORDER, which uses the fitness partition method, only relies on the use of insert. However, we could not run the algorithm with only insertion, because the tree would get very large and the expected time to termination would actually become infinite, in the absence of a strict bound on the tree size. Therefore, deletions are necessary to control the size of the tree, if nothing else. We could, however, envision designing an alternative operator without substitution, whose insertion and deletion probabilities are imbalanced. By choosing these probabilities appropriately, we could prevent the tree size from getting too large (without explicitly bounding it) while still doing as many insertions as possible, thus allowing the algorithm to reach the optimum with the lowest number of fitness evaluations.

For MAJORITY, the substitution decomposability property indicates that, for this particular problem, the substitution operation is more complex than insertion and deletion and is, in fact, a macro operator which is a combination of the two. A superficial glance at the operator does not necessarily reveal this; in fact, it is tempting to believe that substitu-

tion is generally the least complex of the operators, because it most closely resembles a bit-flip in a fixed-length representation and does not change the size or structure of the GP tree at all. However, it could be beneficial to dispense with substitution altogether; this would simplify some of the analysis and achieve the goal of making our mutation operator as “local” as possible for the given problem.

Locality is a property that depends on fitness landscape and operator. Here we see the interaction explicitly and reflect upon the influence of the fitness landscape, which itself depends upon the genotype to phenotype mapping. A substitution makes the same genotypic change because MAJORITY and ORDER share the same primitive set. But in MAJORITY, to a first order approximation, the amortized average change in fitness is larger than in ORDER because ORDER’s expression mechanism places emphasis on only the front of the parsed leaf list whereas MAJORITY’s depends on the entire set of leaves.

Both problems also reveal a fundamental asymmetry between deletions and insertions. Insertions select uniformly from the set of possible terminals, so each terminal is affected with the same probability, but deletions select uniformly from the set of leaves, so the probability of the operation changing a particular terminal depends on the concentration of that terminal in the tree. In the case of ORDER, this has ramifications for the evolution of the tree size over time, because insertions end up being less likely to decrease the fitness than deletions, so the tree grows over time. For MAJORITY, this phenomenon has a positive effect, rather than a negative one: if there are more occurrences of a negated variable than its corresponding positive variable, we will tend to remove those negations with higher probability, and simultaneously balance the relative concentration of each variable.

5.3 Informing GP Practice

This analysis also prompts one to revisit and review assumptions about the necessity of a population and a crossover operator in GP. It does not imply that they are unnecessary but it explicitly shows, at least, simple circumstances when they are not. This may advise GP practitioners to assure themselves empirically that a population and a crossover operator are needed (alone or together) when they start their algorithm development. These algorithms and operators are simple and easy to code, yielding a quick solution to the problem while at the same time supporting parallelism and featuring similar efficiency to conventional GP. Even though the problem at hand is not likely to have the simple problem structure observed here and even though it may require a more sophisticated operator, starting from a provably correct algorithm provides a platform for rationally exploring how to address the separate challenges presented by a harder problem or designing an operator specifically for the problem.

Additionally, this analysis contrasts with conventional GP design practice. Conventionally, GP design proceeds in a very practical manner, but one which is antithetical to theoretical algorithm designers. Rather than derive an algorithm that is provably correct and of efficient complexity, practitioners use biological inspiration, empirical insight and cur-

rent GP theory. This current theory tries to provide transparent explanations of how GP bloats, how it constructs solutions from schema and how it navigates a fitness landscape with its operators and selection. The resulting heuristic can be expected to generate initial mixed results and require subsequent trial and error to “perfect” its use on the problem of interest. There exist some best practices and rules of thumb for robust algorithms but little useful guidance on algorithm customization (via, e.g. genetic operators) for this subsequent design phase. The process finally yields a heuristic which, though a randomized algorithm, is intractable to analyze post-hoc for correctness or efficiency. One can offer to a user its computational expense, which is the product of population size, generations, and number of runs, as well as some empirical estimate of the likelihood of finding a sufficiently optimal solution on a future problem instance. An open question is whether the algorithm design methodology taken in this contribution, that of algorithmic theoreticians, could be blended with or complement the method of practice. Our methodology yields a fundamentally different new form of theoretical result for GP: a randomized algorithm of established computational efficiency (which is different from computational expense), that is guaranteed to find a solution. However, our analysis is tractable only because the fitness structure of the model problems is simple, and the algorithms use only a simple hierarchical variable length mutation operator. It is an open question as to whether the first pass application of the simple algorithms and operators on a realistic problem might prove useful for insight or well founded design choices. Forums such as the annual Genetic Programming: From Theory to Practice workshop which encourage explicit interactions among theoreticians and practitioners may encourage the investigation of this question and provide a means of collecting the experiences.

5.4 Analysis Techniques

We also comment briefly on our analysis techniques. The analysis of ORDER used the method of fitness partitions [2], a very general method that has found many applications in the complexity analysis of EAs for binary representation. The method was successful because, in ORDER, we are always only one move away from expressing a particular variable. However, MAJORITY required different analysis techniques because, unlike in ORDER, there is not an easily computable probability of improving the fitness given only the fitness value; there is also a crucial dependence on the neutral moves we make. The coupon collector and random walk method considered optimizing all of the variables jointly, and in doing so achieved a bound on performance very close to the theoretical lower bound. This indicates that perhaps more complex fitness functions that are not separable might be admissible to similar styles of analysis.

5.5 Future Work

We see three main directions for future work in the computational complexity analysis of genetic programming. Obviously the goal of bridging a gap from what exists to practice is daunting. However, modest steps forward may be revealing. The first extension is to increase the complexity of the genetic operators that are acting on these two problems. Our 1 + 1 operators are essentially just stochastic hill climbers, and while understanding of such an optimization technique is valuable in and of itself, real-world GP implementations

clearly involve more individuals. From GA theory, there is a precedent of taking $1 + 1$ analysis of a problem and extending it to $\mu + 1$ analysis, where μ is the size of the population (see e.g. [16]). This would admit tree-based crossover operators (if they can be shown to be necessary for an efficient optimization time).

The other extension is to consider harder problems. While ORDER and MAJORITY each capture a couple very simple properties of a program, both rely upon inspection and neither problem's fitness function takes into account the hierarchical nature of a GP tree, which is of crucial importance for all practical applications of GP. We could extend these problems in several ways to address the latter issue. One could keep the same terminal set and join nodes, but make the fitness function take subtrees into account somehow. Alternatively, one could introduce a new type of join operation and use the fitness function to impose a constraint that forces us to optimize the higher levels of the tree as well, perhaps by giving higher fitness to individuals when this "join prime" has regular joins as its children and vice versa. Either of these changes would increase the interest of the problem and make the results more relevant to the way that GP is used in practice. However, any new problem in this form may require new or modified mutation operators to be admissible to analysis. We could also try to extend the difficulty of MAJORITY by designing a new objective that requires the correct material (and no incorrect material) to be in the right order, in addition to being present. A model problem which abstracts the semantics of iteration might also provide insights, if tractable. Whether the currently used proof techniques—fitness partitioning, random walks and use of the coupon collector—are sufficient to address more challenging setups is an open question.

6. REFERENCES

- [1] R. Aleliunas, R. M. Karp, R. J. Lipton, L. Lovász, and C. Rackoff. Random walks, universal traversal sequences, and the complexity of maze problems. In *FOCS*, pages 218–223. IEEE, 1979.
- [2] S. Droste, T. Jansen, and I. Wegener. On the analysis of the $(1+1)$ evolutionary algorithm. *Theor. Comput. Sci.*, 276:51–81, 2002.
- [3] D. E. Goldberg and U.-M. O'Reilly. Where does the good stuff go, and why? how contextual semantics influence program structure in simple genetic programming. In W. Banzhaf, R. Poli, M. Schoenauer, and T. C. Fogarty, editors, *Proceedings of the First European Workshop on Genetic Programming*, volume 1391 of *LNCS*, pages 16–36, Paris, 14–15 Apr. 1998. Springer-Verlag.
- [4] T. Jansen and I. Wegener. Evolutionary algorithms - how to cope with plateaus of constant fitness and when to reject strings of the same fitness. *IEEE Trans. Evolutionary Computation*, 5(6):589–599, 2001.
- [5] J. R. Koza. *Genetic Programming: On the Programming of Computers by Means of Natural Selection*. MIT Press, Cambridge, MA, USA, 1992.
- [6] J. R. Koza. *Genetic Programming II: Automatic Discovery of Reusable Programs*. MIT Press, Cambridge Massachusetts, May 1994.
- [7] H. Mühlenbein. How genetic algorithms really work: mutation and hillclimbing. In *Proc. of PPSN '92*, pages 15–26. Elsevier, 1992.
- [8] A. N. Myers and H. S. Wilf. Some new aspects of the coupon collector's problem. *SIAM J. Discret. Math.*, 17(1):1–17, 2004.
- [9] F. Neumann and C. Witt. *Bioinspired Computation in Combinatorial Optimization – Algorithms and Their Computational Complexity*. Springer, 2010.
- [10] U.-M. O'Reilly. *An Analysis of Genetic Programming*. PhD thesis, Carleton University, Ottawa-Carleton Institute for Computer Science, Ottawa, Ontario, Canada, 22 Sept. 1995.
- [11] U.-M. O'Reilly and F. Oppacher. Program search with a hierarchical variable length representation: Genetic programming, simulated annealing and hill climbing. In Y. Davidor, H.-P. Schwefel, and R. Manner, editors, *Parallel Problem Solving from Nature – PPSN III*, number 866 in *Lecture Notes in Computer Science*, pages 397–406, Jerusalem, 9–14 Oct. 1994. Springer-Verlag.
- [12] U.-M. O'Reilly and F. Oppacher. A comparative analysis of GP. In P. J. Angeline and K. E. Kinneary, Jr., editors, *Advances in Genetic Programming 2*, chapter 2, pages 23–44. MIT Press, Cambridge, MA, USA, 1996.
- [13] R. Poli, L. Vanneschi, W. Langdon, and N. McPhee. Theoretical results in genetic programming: the next ten years? *Genetic Programming and Evolvable Machines*, 11:285–320, 2010. 10.1007/s10710-010-9110-5.
- [14] M. Raab and A. Steger. Balls into bins - a simple and tight analysis. In *RANDOM '98: Proceedings of the Second International Workshop on Randomization and Approximation Techniques in Computer Science*, pages 159–170, London, UK, 1998. Springer-Verlag.
- [15] G. Rudolph. *Convergence properties of evolutionary algorithms*. Hamburg: Kovaáč, 1997.
- [16] C. Witt. Runtime analysis of the $(\mu + 1)$ EA on simple pseudo-boolean functions. *Evolutionary Computation*, 14(1):65–86, 2006.