

Introns and Templates Matter: Rethinking Linkage in GP-GOMEA

Johannes Koch
Centrum Wiskunde & Informatica
Amsterdam, The Netherlands
Delft University of Technology
Delft, The Netherlands
johannes.koch@cw.nl

Tanja Alderliesten
Leiden University Medical Center
Leiden, The Netherlands
T.Alderliesten@lumc.nl

Peter A.N. Bosman
Centrum Wiskunde & Informatica
Amsterdam, The Netherlands
Delft University of Technology
Delft, The Netherlands
peter.bosman@cw.nl

Abstract

GP-GOMEA is among the state-of-the-art for symbolic regression, especially when it comes to finding small and potentially interpretable solutions. A key mechanism employed in any GOMEA variant is the exploitation of linkage, the dependencies between variables, to ensure efficient evolution. In GP-GOMEA, mutual information between node positions in GP trees has so far been used to learn linkage. For this, a fixed expression template is used. This, however, leads to introns for expressions smaller than the full template. As introns have no impact on fitness, their occurrences are not directly linked to selection. Consequently, introns can adversely affect the extent to which mutual information captures dependencies between tree nodes. To overcome this, we propose two new measures for linkage learning, one that explicitly considers introns in mutual information estimates, and one that revisits linkage learning in GP-GOMEA from a grey-box perspective, yielding a measure that needs not to be learned from the population but is derived directly from the template. Across five standard symbolic regression problems, GP-GOMEA achieves substantial improvements using both measures. We also find that the newly learned linkage structure closely reflects the template linkage structure, and that explicitly using the template structure yields the best performance overall.

CCS Concepts

• Theory of computation → Evolutionary algorithms; Genetic programming.

Keywords

Genetic programming, symbolic regression, linkage learning, GP-GOMEA

ACM Reference Format:

Johannes Koch, Tanja Alderliesten, and Peter A.N. Bosman. 2026. Introns and Templates Matter: Rethinking Linkage in GP-GOMEA. In *Genetic and Evolutionary Computation Conference (GECCO '26)*, July 13–17, 2026, San Jose, Costa Rica. ACM, New York, NY, USA, 17 pages. <https://doi.org/10.1145/3795095.3805139>



This work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License.

GECCO '26, San Jose, Costa Rica

© 2026 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-2487-9/2026/07

<https://doi.org/10.1145/3795095.3805139>

1 Introduction

Especially in domains where decisions can have severe consequences on lives and livelihoods, such as healthcare, accountability and responsibility are paramount. In recognition of this, the fields of eXplainable AI (XAI) and interpretable machine learning aim to either explain model outputs or learn models that are inherently transparent [40, 55]. Symbolic regression (SR), the task of finding mathematical expressions that accurately model relationships in data, is one form of such inherently interpretable models [28]. Small mathematical expressions describing one target variable consisting of atomic functions (e.g., $\{+, -, \times, \div, \sin\}$), numerical parameters, and input features potentially are not only readable for human experts, but also have the potential to be interpretable and can help uncover yet unknown knowledge.

GP-GOMEA, a Genetic Programming (GP) version of the Gene-pool Optimal Mixing Evolutionary Algorithm (GOMEA), is one of the leading algorithms when it comes to finding small, yet accurate symbolic expressions [3, 29]. GP algorithms typically start with an initially random population of solutions that is refined through repeated selection and variation. One key aspect to GP-GOMEA is the explicit modeling and exploitation of key dependencies between decision variables, also called linkage [53]. Linkage is well-known to be a crucial factor in how effectively an algorithm can tackle problems in general, where problems can be exponentially harder to solve without the correct linkage structure [46]. If no linkage information is known a priori, it can be detected and learned during optimization. Oftentimes, this happens based on statistics taken from the current (and potentially previous) population(s). Specifically, a measure is needed that expresses how strongly variables are related. For this, it was previously proposed to use mutual information (MI) between decision variables in GP-GOMEA, with a modification to avoid the initial detection of false linkage due to non-uniform initialization methods [53].

Similarly to other approaches that explicitly maintain distributions over the expressions in the search space [22, 34, 42], a fixed template is used to ensure all solutions have the same number of variables and that variable indices correspond to identical expression positions across solutions. Oftentimes, without further knowledge of the type of expression being searched for, the template corresponds to a full binary tree with a fixed depth. For expressions that do not use the entire template, some decision variables are conditionally inactive, often called introns [30].

However, from the perspective of linkage learning, introns can be problematic, as the absence of selection pressure on introns causes them to act as a source of noise to the statistics used to learn linkage. Inactive variables even have been recognized as a source of difficulty beyond GP [19, 20, 35, 36].

The goal of this paper is to study whether explicitly taking the presence of introns into account during linkage learning can improve the detection of important dependencies and consequently lead to improved performance by GP-GOMEA. Moreover, we wish to study closer what dependencies are discovered once the noise caused by the presence of introns is removed.

The remainder of this paper is organized as follows: Related work is covered in Section 2. GP-GOMEA and new ways of detecting linkage are discussed in Section 3. Experiments and results thereof are covered in Section 4. Finally, we discuss our results in Section 5 and conclude in Section 6.

2 Related Work

The importance of linkage in evolutionary optimization is well-known [13] and it has been shown that the computational effort needed to solve optimization problems can exponentially increase without the right linkage model [46]. There are various ways to handle linkage to more effectively propagate building blocks with less disruption: The solution representation can be optimized towards highly linked low-order schemata [24], variables can be re-ordered during evolution to improve linkage [17, 30], domain specific crossovers can be used [38], distribution-based algorithms can encode linkage in probabilistic models [8, 33], or linkage can be modeled and learned explicitly. GOMEAs fall into the last category of model-based evolutionary algorithms, where linkage information is either known or learned [4, 15, 35].

In the field of GP, too, decision variables are typically arranged such that related variables are grouped and varied together, for example by representing solutions as trees in standard GP. All the above-mentioned approaches arguably have been explored to various extents, albeit not necessarily with an explicit intent towards linkage in GP: Tree-based GP keeps related variables close together [28], linear and cartesian GP allow re-ordering subroutines by decoupling semantic from spatial adjacency [58], various context aware crossover operators [31] and distribution based GP algorithms [34, 42, 45] have been proposed. How linkage between variables is modeled and exploited is generally highly specific to the solution representation and algorithm used.

When it comes to explicitly modeling and learning linkage as done in GOMEA, various topics have been explored. From different ways of modeling and learning linkage, to similarity measures, to filtering extraneous linkage, to conditional variation operators that consider the context of already varied linked variables [4, 6, 7, 9, 15, 35, 37]. In particular, fitness-based linkage learning, i.e., performing targeted fitness evaluations to uncover definitive linkage instead of using potentially untruthful statistical measures, is gaining ground [4, 35, 37]. Due to the use of a fixed template in GP-GOMEA and the knowledge that it is possible for any two variables to interact, this form of linkage learning does not apply to symbolic regression (SR). The linkage measure used typically heavily depends on the domain [5, 10, 41].

To the best of our knowledge, additional information available in the GP setting such as the user-defined template structure or knowledge about inactive variables has not yet been considered when it comes to linkage. Nonetheless, an adjustment to the estimated linkage to prevent the detection of spurious linkage after

initialization and a binning approach to handle numerical constants were previously proposed [53]. This adjusted linkage measure is re-visited and compared to in this work.

3 Method

In this section, GP-GOMEA is introduced, before the proposed changes to how linkage is learned, are presented.

3.1 GP-GOMEA

As in other population-based algorithms, GP-GOMEA works by iteratively improving an initial, often random, set of candidate solutions through variation and selection. Compared to most other algorithms, a linkage learning step is performed every generation in GP-GOMEA, followed by variation using the gene-pool optimal mixing (GOM) operator. GOM combines both variation and selection on a per-solution basis as shown in Algorithm 1.

Each solution consists of a fixed-length string of decision variables, where a common tree-based template structure is used to map decision variables to tree nodes, as shown in Figure 1. Typically, this template is a single n -ary tree, where n is the maximum arity in the function set used [53]. However, the use of multiple output trees or modular subfunctions is also possible [21, 44]. By using a fixed structural template compared to a representation with a dynamic length, the same variable position corresponds to the same tree node across all individuals. Hence, the linkage between variables corresponds to the linkage of the respective tree nodes. However, this inherently introduces conditionally inactive variables or *introns*, as for smaller expressions not all decision variables are used in the corresponding semantic expression.




Figure 1: A fixed-length string of decision variables is mapped to a fixed tree template, which defines the corresponding semantic expression. Shaded variables are introns, and do not affect the semantic meaning of the expression.

In each generation, first a linkage model is learned, aimed at capturing the dependencies between variables. To represent this linkage information, a set containing multiple subsets of all decision variables, called a Family of Subsets (FOS, denoted $\mathcal{F}$), is commonly used. Each subset of variables in the FOS corresponds to a set of linked variables that are jointly optimized by using the subset as a crossover mask during variation. In GP-GOMEA, this FOS is learned during optimization based on a pairwise similarity measure between decision variables, such as MI. After estimating the similarity measure, hierarchical clustering with UPGMA [18] is performed to repeatedly merge subsets of variables starting from all single variable subsets. The last merge, i.e., the subset containing all variables, is discarded as the corresponding crossover mask would perform replacement instead of variation. This hierarchical FOS structure is commonly referred to as a Linkage Tree (LT) [15, 47].

Variation is performed separately for each solution, each of which is first cloned into an offspring solution. Next, for each offspring solution, each subset in the FOS is considered, and the offspring solution inherits all variables in the subset from a donor solution chosen uniformly randomly from the population. Specific to the GP domain, an evaluation is only necessary if this leads to a modification of any actively used part of the expression. If this is the case, the fitness of the resulting solution is compared with a backup and regressions in fitness are reverted.

Algorithm 1: GP-GOMEA

```

1  $\mathcal{P} \leftarrow \text{INITIALIZEANDEVALUATEPOPULATION}()$ 
2 while  $\neg \text{TERMINATIONCRITERIASATISFIED}$  do
    // 1. Linkage learning
3    $S \leftarrow \text{ESTIMATESIMILARITY}(\mathcal{P})$ 
4    $\mathcal{F} \leftarrow \text{BUILDLINKAGETREE}(S)$ 
    // 2. GOM
5    $O \leftarrow \text{CLONE}(\mathcal{P})$ 
6   for  $i \in \{1, \dots, |\mathcal{O}|\}$  do
7     for  $s \in \text{SHUFFLE}(\mathcal{F})$  do
8        $\text{donor} \leftarrow \text{RANDOM}(\{1, \dots, |\mathcal{P}|\} \setminus \{i\})$ 
9        $\text{backup} \leftarrow \text{CLONE}(O_i)$ 
10       $O_i[s] \leftarrow \mathcal{P}_{\text{donor}}[s]$ 
11      if  $O_i[\text{active}] \neq \text{backup}[\text{active}]$  then
12         $\text{EVALUATE}(O_i)$ 
13        if  $f(O_i)$  worse than  $f(\text{backup})$  then
14           $O_i \leftarrow \text{backup}$ 
15    $\mathcal{P} \leftarrow O$ 

```

To avoid determining an appropriate population size for a problem, (GP-)GOMEA is often paired with the interleaved multistart scheme (IMS), in which populations of exponentially increasing sizes perform generations in a scheduled manner, with larger populations performing generations less frequently than smaller populations, and under-performing populations are terminated when converged or once larger populations perform better [53].

3.2 Linkage Measures in GP-GOMEA

The FOS that dictates how variation in GP-GOMEA adheres to the learned variable dependencies, is constructed using a pairwise similarity measure. This measure is computed by interpreting the population as a dataset without regard for the mapping to the tree template. In this work, two new approaches are proposed to compute more fitting similarity measures by exploiting so far unused domain knowledge: the knowledge about inactive variables and the fixed template structure.

3.2.1 Existing measures. For GP, the previously explored measures are MI, an adjusted version thereof, and random noise [52, 53]. For the random measure, pairwise similarity is sampled from $\mathcal{U}(0, 1)$. For two random variables X and Y , the MI is defined as:

$$MI(X, Y) = H(X) + H(Y) - H(X, Y) \quad (1)$$

where $H(X) = \sum_x -P(X = x) \log_2(P(X = x))$ defines entropy.

In GP, expressions are typically initialized with domain-specific methods such as Half-and-Half [53] instead of uniformly randomly. While this often achieves more desirable diversity in some aspects,




Figure 2: Top left: population of 4 solutions. Bottom left: tree template. Matrices in blue show entropy (above the black line) and MI (below the black line). The corresponding linkage tree for both the normal and proposed masked MI are shown as well. The masked version treats all introns (shaded) as the same separate "masked" symbol.

such as the initial expression depths, biases can appear in the initial MI estimates [53]. To avoid this, normalization factors were proposed to adjust the MI such that no linkage is detected directly after initialization, leading to increased performance [53].

3.2.2 Exploiting knowledge about introns. The use of MI in GOMEA can often lead to successfully detecting variable dependencies. However, inactive variables in the form of introns are inherent to GP-GOMEA. The impact that this may have on dependency detection, has so far not been considered. For each solution, it is known which variables are currently active or inactive. The values of inactive variables do not influence the semantics and thus are freely interchangeable and not subject to selection pressure. From the perspective of entropy estimation, the values of inactive variables can therefore be considered to be noise, hiding the true signal of active variables. To avoid this undesirable effect, we propose to mask the inactive variables using a special intron label during the entropy calculation, i.e., the alphabet of any variable now includes the "masked" label. An example is shown in Figure 2, where masking introns leads to a different linkage tree compared to using normal MI. Masking makes the estimated MI more accurate, by ensuring that the estimate is fully based on active variables under selection pressure. Note that this approach is not compatible with the adjustments for biased initialization methods proposed in [53]. This is because initially, variables that are inactive in the entire population would lead to divisions by zero.

3.2.3 Exploiting the known template. In discrete and real-valued evolutionary optimization literature, it is common to model dependencies with a variable interaction graph (VIG) [4, 51, 57]. In such a graph, each variable is a vertex, and an edge exists between two vertices if there exists a dependency between the two variables. Commonly, a VIG is associated with grey-box optimization scenarios in which it is assumed that the problem being optimized is sufficiently known. By studying the problem formulation, dependencies between variables may then be directly gauged by identifying which variables are jointly part of a subfunction. As an extension, or more informed, variant of the VIG, the edges in the graph can be weighted. The resulting wVIG then expresses not

only whether or not two variables are jointly dependent, but also how strong that dependence is, relative to other dependencies. In GOMEA, such a wVIG can be used directly to build a linkage tree (or other types of linkage structures such as conditional linkage sets [4]) by creating a similarity matrix S that contains the weight of an edge between vertices i and j at S_{ij} and S_{ji} .

The function represented by a GP tree template in GP-GOMEA, is a nested function. It can be defined using two functions f and g in which the first argument of f determines the operation represented in the corresponding node of the GP tree, which may include returning the value of a constant or input variable and ignoring the other arguments. Function g is used only in leaf nodes of the GP tree template, returning the value of a constant or an input variable. Given a data vector $\mathbf{d}$, the function $GP(\mathbf{x})$ that a tree template represents, can be written as a nested application of function f with function g at the deepest level. I.e., for depth 2, we have:

$$GP(\mathbf{x}, \mathbf{d}) = f(x_0, f(x_1, g(x_3, \mathbf{d}), g(x_4, \mathbf{d}), \mathbf{d}), f(x_2, g(x_5, \mathbf{d}), g(x_6, \mathbf{d}), \mathbf{d})) \quad (2)$$

This function is the core of what imposes dependencies between the variables that we optimize over in GP-GOMEA. Consider the SSE optimization function for symbolic regression over all data vectors $\mathbf{d}^i$ and targets t^i :

$$\min_{\mathbf{x}} \left\{ \sum_i (GP(\mathbf{x}, \mathbf{d}^i) - t^i)^2 \right\} \quad (3)$$

Essentially this is a large sum of independent terms. Focusing on a single term, we find $GP(\mathbf{x}, \mathbf{d}^i)^2 - 2GP(\mathbf{x}, \mathbf{d}^i)t^i + (t^i)^2$. Since the variables in GP are already fully coupled through nonlinear interactions due to the nested function composition, squaring the function does not alter the dependency structure, but amplifies existing interactions. Hence, we can construct the wVIG based on the definition of $GP(\mathbf{x}, \mathbf{d})$.

A straightforward and common way of doing so is counting how many times a variable is contained in the same subfunction as another variable. For the depth 2 case, we get the similarity matrix defined in the middle of Figure 3.

However, for multiple reasons, the dependence between the arguments of a parent node in the template tree is likely almost always smaller than the dependence between an argument and its parent node. One of these reasons is that some functions may impose no dependencies between its arguments. Another reason is that generally, the higher a node is in the template tree (and therefore the less deeply nested in the functional description), the larger its impact on changing the output of the tree and thus on the regression error. Moreover, internal nodes may be terminals. Thus, the second and third arguments of f may not always be used. In that case, it does not matter what values the variables take in the deeper nested functions (as they are then introns), and thus, in those cases, those variables are independent. We can convey this information in a systematic manner by defining a distance function between any two nodes in the GP template that is the smallest distance between these nodes to travel along the template. This makes the distance between two arguments of the same parent node 1 larger than the distance between each of the arguments and the parent node (i.e., 2 and 1, respectively).




Figure 3: An example template with highlighted subfunctions and the corresponding node proximity and subfunction based similarity measures.

This distance measure can be turned into a normalized similarity measure as follows:

$$S_{i,j} = \begin{cases} 1 - \frac{d(i,j)}{1 + \max_{k,l \in \mathcal{I}} \{d(k,l)\}} & \text{if } i \text{ and } j \text{ are connected} \\ 0 & \text{otherwise} \end{cases} \quad (4)$$

where the distance $d(i, j)$ is the total number of edges between nodes i and j and their closest ancestor in the template, and $\mathcal{I}$ is the set of all node indices.

As shown on the right of Figure 3, while nodes are still equally similar by definition using this measure of structural relatedness, there is more granularity than in the direct subfunction encapsulation approach. During the hierarchical clustering, ties are broken randomly if there are multiple merges with the same similarity. By rebuilding the linkage tree every generation, different LT FOS structures representing the similarity matrix are created. Note that compared to MI estimates, this form of pre-defined similarity disregards semantic relations due to variable values entirely. However, there is no need for population sizes sufficiently large to get statistically reliable estimates and literature has shown that appropriate offline similarity measures can lead to much better results [11, 12, 50].

4 Experiments and Results

In this section, two experiments are performed to evaluate the proposed measures. First, the proposed measures are compared with previously used measures. The second experiment aims to uncover what linkage structures are learned.

4.1 Experimental Setup

To isolate the effect of linkage learning from other algorithm settings such as the fixed structural template and whether linear scaling [26] (hereafter LS) is enabled, all combinations are tested. LS is a much-used post-processing technique that provides scale and translation invariance by performing a linear regression on the output of the GP expression. The other settings are detailed in Table 2. The operator set was chosen based on [32]. The coefficient of determination (R^2 Score) is used as objective:

$$R^2(y_{pred}, y_{target}) = 1 - \frac{MSE(y_{pred}, y_{target})}{var(y_{target})} \quad (5)$$

$$MSE(y_{pred}, y_{target}) = \frac{1}{N} \sum_{i=1}^N (y_{pred_i} - y_{target_i})^2 \quad (6)$$

where y_{pred} and y_{target} are the predicted output (as in Equation (2)) and target variable respectively, and N is the dataset size. Note that with LS enabled, the predicted values are scaled and transformed before computing the prediction error.

Table 1: The problems used in the experiments.

Dataset	#Instances	#Features
Airfoil [49]	1503	5
Bike Sharing (Daily) [16]	731	12
Concrete Compressive Strength [25]	1030	8
Dow Chemical [1]	1066	57
Tower [56]	4999	25

The problems we used in our experiments (listed in Table 1) were selected to contain varying numbers of input features. For each problem, 25% of the instances form a held-out test set. The remaining data is split into 5 cross-validation folds, using 80% for training and 20% for validation. Each fold is repeated 6 times with different seeds, yielding 30 independent runs. Seeds are shared across algorithms so that all algorithms have identical initial populations.

Numerical constants in the form of ephemeral random constants (ERCs) are used for which random values are sampled during initialization that remain constant throughout evolution. During linkage learning, ERCs are handled as in [43], i.e., for the MI-based measures, constant values are split into 25 discrete bins (note: measures not considering population statistics are unaffected).

In the first experiment, a computational budget of 10^7 evaluations was used. As in [43], the IMS starting population size is set to 64 with a subgeneration factor of 10 generations before larger populations are started. All experiments follow the setup described in Table 2 and were run on a separate physical core of a machine with two Intel Xeon E5-2699v4 processors, ensuring that runtimes are comparable.

In the second experiment, instead of IMS, a single population with a fixed size of 1024 and a generational budget of 20, a template height of 5 (31 Nodes), and LS enabled, is used. This population size aligns with earlier work on GP-GOMEA [27, 53, 54].

We consider the following linkage measures: random linkage, MI, bias adjusted MI as in [53], the proposed approach of masking inactive variables when computing the MI, and the similarity based on node proximity in the structural template with and without repeated building of the linkage tree every generation, hereafter abbreviated as Random, MI, $MI_{adjusted}$, MI_{masked} , Node, and Node (static) respectively. Both building of the linkage tree only once and rebuilding it every generation are considered for the node proximity measure to assess the effect of the aforementioned random tie-breaking during hierarchical clustering. Results for another baseline assuming full variable independence and the subfunction based measure are presented and discussed in the supplementary material, where the results confirm that variables are not independent and the Node measure outperforms the subfunction based measure.

Table 2: Experiment Setup

Setting	Experiment I	Experiment II
Objective	R^2 Score	
Initialization	Half-and-Half [53]	
ERC Initialization	$\mathcal{U}(\min(y_{train}), \max(y_{train}))$	
Operator set	{+, −, ×, ÷, sin}	
Termination	10^7 evaluations	20 generations
Linear scaling	Yes / No	Yes
Template Height	5 / 7 (31 / 127 Nodes)	5 (31 Nodes)

To analyze the results, we follow the recommendations and implementation provided by [2], where interval estimates via stratified bootstrap confidence intervals are preferred over point estimates and statistical significance tests to take uncertainty and the effect size into account. We mainly consider training performance, as that is directly optimized and we are primarily interested in the effectiveness of the linkage models instead of the generalization performance which can be addressed separately, for example through model selection [39]. Statistical analysis is also done as recommended in [14] with the implementation provided by [23].

4.2 Experiment 1: Effectiveness of Various Linkage Measures

To assess the effectiveness of the various linkage measures, we consider both the final accuracies achieved (Figure 5) and intermediate results throughout optimization in Figures 4 and 6. The supplementary material provides per problem results (Section C), and a repeat of the first experiment showing that the conclusions made extend to varying operator set sizes (Section E).

The expected final R^2 scores on both the training and test splits are shown in Figure 5 with two different aggregate metrics, split up into the various combinations of template height and linear scaling that we considered. Two different aggregates, based on the mean performance without outliers, and the median per problem performance are provided, as both measures are differently affected by asymmetric accuracy distributions.

Across all settings, the Node measure performs best, followed by MI_{masked} , indicating that the proposed methods provide more effective linkage information. Interestingly, the static version of the Node measure, where the initial linkage tree is re-used within the same population for all subsequent generations, clearly performs worse than the randomized version. This aligns with [48], where it was found that there is not necessarily one best static linkage structure and randomization can be beneficial. This could be due




Figure 4: The interquartile mean (IQM) training R^2 and bootstrapped 95% confidence intervals (as per [2]) over evaluations and approximate runtime, across all runs performed. For the runtime, the last recorded value before each point in time was used, and dashed horizontal lines show the final values reached for the existing measures. The vertical line indicates the computational budget (not to scale w.r.t runtime) where the final values reached are marked with crosses.




Figure 5: Aggregate R^2 scores (higher is better) on the problems for each combination of template height and LS considered. The interquartile mean corresponds to the mean after discarding the bottom and top 25% of runs for each problem, and the mean of medians corresponds to the mean of the median performances on each problem. The colored bar corresponds to the 95% confidence interval estimated using a percentile bootstrap with stratified sampling as per [2] with the expected value in black.




Figure 6: Median training R^2 score (higher is better) over evaluations and approximate runtime across problems with template height 7 and without linear scaling. The filled area corresponds to the interquartile range, and the runtime where the first run finished is marked with a cross, after which the final value obtained is re-used for completed runs until all runs have finished.

to the static version introducing bias by breaking ties only one way, while the repeated building of linkage trees with ties being broken randomly, leads to an unbiased FOS representation of the estimated linkage structure over multiple generations. Another effect is that the randomization possibly promotes propagation of building blocks between overlapping subsets of subsequent FOS structures better than the static version.

In general, increasing the template height and thus the possible expression sizes, leads to better accuracies, and the same holds for enabling LS. While the template size and LS have a larger impact on accuracy, clearly, linkage learning too has a noticeable impact. However, considering only the already existing methods, the random baseline is not clearly outperformed for all settings, suggesting that previously proposed linkage measures are suboptimal for GP.




Figure 7: The different linkage measures over generations, averaged over 30 runs on the Bike Sharing dataset with a template height of 5 (31 Nodes), LS and a fixed population size of 1024. Note that different scales are used for each measure, as the focus lies on differences between variable pairs within measures instead of differences between linkage measures. The Node measure is shown only once per row as it does not change throughout evolution to ease comparisons.

When it comes to the relative performance of the linkage measures, increasing the number of decision variables (template height) widens the gap between the various methods. LS, as arguably could be expected, has the opposite effect and decreases the differences. Moreover, of course, there is a generalization gap when the training and test performance are compared, however, the relative performance of the various linkage measures mostly remains unaffected. Overall, MI tends to outperform the bias-adjusted version, while the setting using a template height of 5 and LS as used in [53] where statistical tests showed the opposite. Compared to the experiments performed in [53], the budget used here is an order of magnitude larger, and thus any beneficial effects of mitigating initial biases are likely less visible in the final results achieved after exhausting the computational budget. Figure 4, where the convergence aggregated over all combinations and problems considered is shown, also shows an initial advantage of $MI_{adjusted}$ over MI, aligning with the results obtained in [53]. This suggests that the bias adjustment itself does work, however, as optimization progresses, the initial bias becomes less and continued mitigation can be harmful as the MI measure outperforms its adjusted counterpart in various settings.

Intermediate results and the runtime performance across all runs are shown in Figure 4, and Figure 6 provides convergence results per problem for template height 7 without LS. Results for the other combinations are in Section C of the supplementary material. In both figures, the Node and MI_{masked} measures outperform the other

measures throughout evolution with respect to both evaluations performed and runtime. In terms of evaluations, use of the proposed measures outperforms use of the existing measures at only half the computational budget, and the same holds for comparing the IQM achieved using the Node measure after $5e5$ evaluations, and the existing methods after $1e6$ evaluations. Considering runtime, the varying number of instances between the problems (see Table 1) leads to varying runtimes until the full budget is exhausted, ranging from several minutes to hours. Nonetheless, the Node and MI_{masked} measures exceed the final accuracy reached by the existing methods after 10 and 20 minutes, respectively, showing a substantial speed-up in terms of runtime and highlighting the importance of exploiting linkage. The observed efficiency improvement possibly is proportional to the number of nodes in the template, as the advantage in Figure 6 is larger compared to Figure 4 where all settings are aggregated. Statistical differences are shown in Figure 8 across all settings considered, where a critical difference diagram is used to show the average measure ranks. It can be seen that all measures are significantly different from each other, except the normal and bias adjusted variants of mutual information. Clearly, the proposed linkage measures exploit the problem structure more effectively compared to the existing measures, with the template informed Node measure performing best followed by MI_{masked} .




Figure 8: Statistical differences across all problem, template height and linear scaling combinations considered obtained by performing a Friedman test followed by the pairwise Nemenyi test using [23]. Lower ranks indicate better performance, and measures not connected by a bar exceed the critical difference (CD) and are significant (at $p = 0.05$).

4.3 Experiment 2: What linkage structures are learned?

The learned pairwise similarities pertaining to the various measures averaged over 30 runs on the Bike Sharing problem are shown in Figure 7. Results for the other problems are shown in Section D of the supplementary material. Note that both Node versions use the same similarity measure that does not change throughout evolution and Random is not shown as it is not informative. The similarities shown at generation 0 are recorded before any variation is performed, i.e. directly after initialization.

In the first row of Figure 7, the bias in the mutual information estimate due to non-uniform initialization (as addressed in [53]) is clearly visible in the first generation, and notably absent for $MI_{adjusted}$. As the number of generations performed increases, the bias in MI (and the bias adjustment of $MI_{adjusted}$) becomes less visible and both measures learn similar patterns of linked variables.

Compared to the existing measures, MI_{masked} does not seem to suffer from the same bias as MI, suggesting that not accounting for inactive variables is the cause of the bias. The linkage structure learned by MI_{masked} is highly similar to the node proximity measure introduced. However, with increasing generations, parts of the template structure learned by MI_{masked} degrade over time. The patterns learned in later generations by MI and $MI_{adjusted}$, as well as the linkage structure learned by MI_{masked} resemble the Node measure, suggesting that indeed the distance of variables in the fixed template is highly informative. Clearly, once introns are taken into account, this structure can be learned from population statistics. Section D of the supplementary material shows that this general trend of mutual information based linkage learning approximating the template structure holds across all problems considered.

5 Discussion

Learning and exploitation of linkage information can be crucial for effective optimization. We re-visited linkage learning in GP-GOMEA with the aim to more fully exploit the available problem knowledge, in particular knowledge about inactive variables and the user-defined structural template shared across all solutions. The experiments clearly show that the proposed measures outperform existing measures. This holds across all combinations of template heights and linear scaling considered, with clear speed-ups in terms of evaluations and runtime performance.

The proposed methods include no new algorithm settings, and either are a simple modification to how introns are handled during the calculation of the mutual information between variable pairs or

require no runtime cost in the case of the node proximity measure. Notably, masking introns to improve the signal-to-noise ratio of variables not under selection pressure is not specific to GP, and can be applied to other settings where it is known which variables are inactive, as in [19, 35, 36].

For the first time in GP-GOMEA, we also explored what linkage structures are learned by the various measures. Interestingly, all mutual information based variants seem to approximate the structural template, suggesting that at least for the problems considered, the solution representation is more important than variable values when it comes to the dependency strengths between variables. However, there could be problems where this does not hold.

Clearly, the best solutions differ between problems, suggesting that the importance of different parts in the template can vary as well. This could potentially still be captured by runtime measures, especially the masked measure that we proposed that takes into account the conditionally inactive variables that appear due to the use of a tree template. It would be interesting to see in future work if such hybrid linkage measures could improve performance further.

Another aspect to consider is the modeling of the learned linkage, i.e. why randomization of the node proximity performs better, and how much randomization of the FOS structure is best. Currently, the tie breaking in UPGMA only resorts to randomness in the face of equally strongly linked subsets of variables, which might not truly lead to FOS structures that over multiple generations form an unbiased representation of the linkage information.

Beyond improving the learning, modeling and exploitation of linkage, other factors potentially interacting with linkage learning in GP were not considered, but could be equally important. This includes, for example, representing numerical constants differently as done in [27] and the analysis of possible effects on interpretability from an explainability viewpoint, which is not considered here. Nonetheless, the fixed template size used in GP-GOMEA typically restricts the search space to small, and thus more likely humanly interpretable expressions.

6 Conclusion

Learning and exploiting linkage is a key mechanism of GP-GOMEA, an algorithm among the current state-of-the-art when it comes to SR. With this in mind, we revisited this core aspect of the algorithm, leading to two new linkage measures that exploit more of the available problem knowledge.

Our experiments, through which we evaluated the use of various linkage measures in the context of varying template sizes and the use of linear scaling, showed that by exploiting known problem information, clear accuracy and runtime performance improvements are possible across all combinations considered. Moreover, the linkage structures that are learned, align with the user-defined structural template available a priori, with the best performing linkage measure being the one that encapsulates exactly the structure of the template, requiring no additional on-line linkage learning from the population statistics. This establishes a clear bridge between GP-GOMEA and other linkage learning EAs, particularly GOMEA variants, highlighting the advantage of explicitly modeling compositional problem structure, a principle widely recognized in discrete and real-valued optimization.

References

- [1] 2012. Symbolic Regression Competition. <https://web.archive.org/web/20120628140646/http%3A/casnew.iti.upv.es/index.php/evocompetitions/105-symregcompetition>.
- [2] Rishabh Agarwal, Max Schwarzer, Pablo Samuel Castro, Aaron Courville, and Marc G Bellemare. [n.d.]. Deep Reinforcement Learning at the Edge of the Statistical Precipice. ([n.d.]).
- [3] Guilherme S. Imai Aldeia, Hengzhe Zhang, Geoffrey Bomarito, Miles Cranmer, Alcides Fonseca, Bogdan Burlacu, William G. La Cava, and Fabricio Olivetti de França. 2025. Call for Action: Towards the next Generation of Symbolic Regression Benchmark. In *Proceedings of the Genetic and Evolutionary Computation Conference Companion*. 2529–2538. arXiv:2505.03977 [cs] doi:10.1145/3712255.3734309
- [4] Georgios Andreadis, Tanja Alderliesten, and Peter A. N. Bosman. 2024. Fitness-Based Linkage Learning and Maximum-Clique Conditional Linkage Modelling for Gray-box Optimization with RV-GOMEA. arXiv:2402.10757 [cs] doi:10.48550/arXiv.2402.10757
- [5] Peter A.N. Bosman, Ngoc Hoang Luong, and Dirk Thierens. 2016. Expanding from Discrete Cartesian to Permutation Gene-pool Optimal Mixing Evolutionary Algorithms. In *Proceedings of the Genetic and Evolutionary Computation Conference 2016 (GECCO '16)*. Association for Computing Machinery, New York, NY, USA, 637–644. doi:10.1145/2908812.2908917
- [6] Peter A.N. Bosman and Dirk Thierens. 2012. Linkage Neighbors, Optimal Mixing and Forced Improvements in Genetic Algorithms. In *Proceedings of the 14th Annual Conference on Genetic and Evolutionary Computation (GECCO '12)*. Association for Computing Machinery, New York, NY, USA, 585–592. doi:10.1145/2330163.2330247
- [7] Peter A.N. Bosman and Dirk Thierens. 2013. More Concise and Robust Linkage Learning by Filtering and Combining Linkage Hierarchies. In *Proceedings of the 15th Annual Conference on Genetic and Evolutionary Computation (GECCO '13)*. Association for Computing Machinery, New York, NY, USA, 359–366. doi:10.1145/2463372.2463420
- [8] Peter A. N. Bosman and Dirk Thierens. 1999. Linkage Information Processing in Distribution Estimation Algorithms. In *Proceedings of the 1st Annual Conference on Genetic and Evolutionary Computation - Volume 1 (GECCO '99)*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 60–67.
- [9] Peter A. N. Bosman and Dirk Thierens. 2012. On Measures to Build Linkage Trees in LTGA. In *Parallel Problem Solving from Nature - PPSN XII*, David Hutchison, Takeo Kanade, Josef Kittler, Jon M. Kleinberg, Friedemann Mattern, John C. Mitchell, Moni Naor, Oscar Nierstrasz, C. Pandu Rangan, Bernhard Steffen, Madhu Sudan, Demetri Terzopoulos, Doug Tygar, Moshe Y. Vardi, Gerhard Weikum, Carlos A. Coello Coello, Vincenzo Cutello, Kalyanmoy Deb, Stephanie Forrest, Giuseppe Nicosia, and Mario Pavone (Eds.). Vol. 7491. Springer Berlin Heidelberg, Berlin, Heidelberg, 276–285. doi:10.1007/978-3-642-32937-1_28
- [10] Anton Bouter, Tanja Alderliesten, Cees Witteveen, and Peter A. N. Bosman. 2017. Exploiting Linkage Information in Real-Valued Optimization with the Real-Valued Gene-Pool Optimal Mixing Evolutionary Algorithm. In *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO '17)*. Association for Computing Machinery, New York, NY, USA, 705–712. doi:10.1145/3071178.3071272
- [11] Anton Bouter, Stefanus C. Maree, Tanja Alderliesten, and Peter A. N. Bosman. 2020. Leveraging Conditional Linkage Models in Gray-Box Optimization with the Real-Valued Gene-Pool Optimal Mixing Evolutionary Algorithm. In *Proceedings of the 2020 Genetic and Evolutionary Computation Conference*. ACM, Cancun Mexico, 603–611. doi:10.1145/3377930.3390225
- [12] Anton Bouter, Dirk Thierens, and Peter A. N. Bosman. 2025. The Pitfalls and Potentials of Adding Gene-invariance to Optimal Mixing. In *Proceedings of the Genetic and Evolutionary Computation Conference Companion (GECCO '25 Companion)*. Association for Computing Machinery, New York, NY, USA, 1918–1926. doi:10.1145/3712255.3734293
- [13] Ying-ping Chen, Tian-Li Yu, K. Sastry, and D. Goldberg. 2007. A Survey of Linkage Learning Techniques in Genetic and Evolutionary Algorithms.
- [14] Janez Demšar. 2006. Statistical Comparisons of Classifiers over Multiple Data Sets. *J. Mach. Learn. Res.* 7 (Dec. 2006), 1–30.
- [15] Arkadiy Dushatskiy, Marco Virgolin, Anton Bouter, Dirk Thierens, and Peter A. N. Bosman. 2021. Parameterless Gene-pool Optimal Mixing Evolutionary Algorithms. arXiv:2109.05259 [cs] doi:10.48550/arXiv.2109.05259
- [16] Hadi Fanaee-T. 2013. Bike Sharing. doi:10.24432/C5W894
- [17] David E. Goldberg, Bradley Korb, and Kalyanmoy Deb. [n.d.]. Messy Genetic Algorithms: Motivation, Analysis, and First Results. *Complex Systems* 3, 5 ([n.d.]).
- [18] Ilan Gronau and Shlomo Moran. 2007. Optimal Implementations of UPGMA and Other Common Clustering Algorithms. *Inform. Process. Lett.* 104, 6 (Dec. 2007), 205–210. doi:10.1016/j.ipl.2007.07.002
- [19] Arthur Gijit, Dirk Thierens, Tanja Alderliesten, and Peter A. N. Bosman. 2024. Exploring the Search Space of Neural Network Combinations Obtained with Efficient Model Stitching. In *Proceedings of the Genetic and Evolutionary Computation Conference Companion (GECCO '24 Companion)*. Association for Computing Machinery, New York, NY, USA, 1914–1923. doi:10.1145/3638530.3664131
- [20] Arthur Gijit, Dirk Thierens, Tanja Alderliesten, and Peter A. N. Bosman. 2024. Stitching for Neuroevolution: Recombining Deep Neural Networks without Breaking Them. In *Proceedings of the Genetic and Evolutionary Computation Conference Companion*. 1914–1923. arXiv:2403.14224 [cs] doi:10.1145/3638530.3664131
- [21] Joe Harrison, Peter A. N. Bosman, and Tanja Alderliesten. 2025. Thinking Outside the Template with Modular GP-GOMEA. arXiv:2505.01262 [cs] doi:10.48550/arXiv.2505.01262
- [22] Erik Hemberg, Kalyan Veeramachaneni, James McDermott, Constantin Berzan, and Una-May O'Reilly. 2012. An Investigation of Local Patterns for Estimation of Distribution Genetic Programming. In *Proceedings of the 14th Annual Conference on Genetic and Evolutionary Computation (GECCO '12)*. Association for Computing Machinery, New York, NY, USA, 767–774. doi:10.1145/2330163.2330270
- [23] Steffen Herbold. 2020. Autorank: A Python Package for Automated Ranking of Classifiers. (2020).
- [24] John H. Holland. 1992. *Adaptation in Natural and Artificial Systems: An Introductory Analysis with Applications to Biology, Control, and Artificial Intelligence*. The MIT Press. doi:10.7551/mitpress/1090.001.0001
- [25] I-Cheng Yeh. 1998. Concrete Compressive Strength. doi:10.24432/C5PK67
- [26] Maarten Keijzer. 2004. Scaled Symbolic Regression. *Genetic Programming and Evolvable Machines* 5, 3 (Sept. 2004), 259–269. doi:10.1023/B:GENP.0000030195.77571.f9
- [27] Johannes Koch, Tanja Alderliesten, and Peter A. N. Bosman. 2024. Simultaneous Model-Based Evolution of Constants and Expression Structure in GP-GOMEA for Symbolic Regression. In *Parallel Problem Solving from Nature – PPSN XVIII*, Michael Affenzeller, Stephan M. Winkler, Anna V. Kononova, Heike Trautmann, Tea Tušar, Penousal Machado, and Thomas Bäck (Eds.). Vol. 15148. Springer Nature Switzerland, Cham, 238–255. doi:10.1007/978-3-031-70055-2_15
- [28] John R. Koza. 1994. Genetic Programming as a Means for Programming Computers by Natural Selection. *Statistics and Computing* 4, 2 (June 1994), 87–112. doi:10.1007/BF00175355
- [29] William La Cava, Bogdan Burlacu, Marco Virgolin, Michael Kommenda, Patryk Orzechowski, Fabricio Olivetti de França, Ying Jin, and Jason H. Moore. 2021. Contemporary Symbolic Regression Methods and Their Relative Performance. *Advances in neural information processing systems* 2021, DB1 (Dec. 2021), 1–16.
- [30] Fernando G Lobo, Kalyanmoy Deb, David E Goldberg, Georges R Harik, and Liwei Wang. [n.d.]. Compressed Introns in a Linkage Learning Genetic Algorithm. ([n.d.]).
- [31] Hammad Majeed and Conor Ryan. 2006. A Less Destructive, Context-Aware Crossover Operator for GP. In *Genetic Programming*, Pierre Collet, Marco Tomassini, Marc Ebner, Steven Gustafson, and Anikó Ekárt (Eds.). Springer, Berlin, Heidelberg, 36–48. doi:10.1007/11729976_4
- [32] Miguel Nicolau and Alexandros Agapitos. 2021. Choosing Function Sets with Better Generalisation Performance for Symbolic Regression Models. *Genetic Programming and Evolvable Machines* 22, 1 (March 2021), 73–100. doi:10.1007/s10710-020-09391-4
- [33] Martin Pelikan, David E. Goldberg, and Erick Cantú-Paz. 1999. BOA: The Bayesian Optimization Algorithm. In *Proceedings of the 1st Annual Conference on Genetic and Evolutionary Computation - Volume 1 (GECCO '99)*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 525–532.
- [34] Riccardo Poli and Nicholas Freitag McPhee. 2008. A Linear Estimation-of-Distribution GP System. In *Genetic Programming*, Michael O'Neill, Leonardo Vanneschi, Steven Gustafson, Anna Isabel Esparcia Alcázar, Ivanoe De Falco, Antonio Della Cioppa, and Ernesto Tarantino (Eds.). Springer, Berlin, Heidelberg, 206–217. doi:10.1007/978-3-540-78671-9_18
- [35] Michal W. Przewozniczek, Peter A.N. Bosman, Anton Bouter, Arthur Gijit, Marcin M. Komarnicki, and Dirk Thierens. 2025. Conditional Direct Empirical Linkage Discovery for Solving Multi-Structured Problems. In *Proceedings of the 18th ACM/SIGEVO Conference on Foundations of Genetic Algorithms (FOGA '25)*. Association for Computing Machinery, New York, NY, USA, 238–249. doi:10.1145/3729878.3746624
- [36] Michal W. Przewozniczek, Bartosz Frej, and Marcin M. Komarnicki. 2024. The Hop-like Problem Nature – Unveiling and Modelling New Features of Real-World Problems. arXiv:2406.01215 [cs] doi:10.48550/arXiv.2406.01215
- [37] Michal W. Przewozniczek, Marcin M. Komarnicki, and Bartosz Frej. 2021. Direct Linkage Discovery with Empirical Linkage Learning. In *Proceedings of the Genetic and Evolutionary Computation Conference*. ACM, Lille France, 609–617. doi:10.1145/3449639.3459333
- [38] Ozeas Quevedo De Carvalho and Darrell Whitley. 2025. Dramatically Faster Partition Crossover for the Traveling Salesman Problem. In *Proceedings of the Genetic and Evolutionary Computation Conference*. ACM, NH Malaga Hotel Malaga Spain, 818–826. doi:10.1145/3712256.3726465
- [39] Fitri Wulandari Ramlan, Gabriel Kronberger, Colm O'Riordan, and James McDermott. 2026. Comparative Analysis of Model Selection Criteria for Symbolic Regression Using Genetic Programming. In *Computational Intelligence*, Francesco Marcelloni, Kurosh Madani, Niki van Stein, and Joaquim Filipe (Eds.). Springer Nature Switzerland, Cham, 91–108. doi:10.1007/978-3-032-15635-8_6

- [40] Cynthia Rudin. 2019. Stop Explaining Black Box Machine Learning Models for High Stakes Decisions and Use Interpretable Models Instead. *Nature Machine Intelligence* 1, 5 (May 2019), 206–215. doi:10.1038/s42256-019-0048-x
- [41] Krzysztof L. Sadowski, Peter A.N. Bosman, and Dirk Thierens. 2016. Learning and Exploiting Mixed Variable Dependencies with a Model-Based EA. In *2016 IEEE Congress on Evolutionary Computation (CEC)*. 4382–4389. doi:10.1109/CEC.2016.7744347
- [42] Rafal Saliustowicz and Jürgen Schmidhuber. 1997. Probabilistic Incremental Program Evolution: Stochastic Search through Program Space. In *Machine Learning: ECML-97*, Maarten van Someren and Gerhard Widmer (Eds.). Springer, Berlin, Heidelberg, 213–220. doi:10.1007/3-540-62858-4_86
- [43] Thalea Schlender, Mafalda Malafaia, Tanja Alderliesten, and Peter Bosman. 2024. Improving the Efficiency of GP-GOMEA for Higher-Arity Operators. In *Proceedings of the Genetic and Evolutionary Computation Conference*. ACM, Melbourne VIC Australia, 971–979. doi:10.1145/3638529.3654118
- [44] E. M. C. Sijben, T. Alderliesten, and P. A. N. Bosman. 2022. Multi-Modal Multi-Objective Model-Based Genetic Programming to Find Multiple Diverse High-Quality Models. In *Proceedings of the Genetic and Evolutionary Computation Conference*. ACM, Boston Massachusetts, 440–448. doi:10.1145/3512290.3528850
- [45] Jinglu Song, Qiang Lu, Bozhou Tian, Jingwen Zhang, Jake Luo, and Zhiguang Wang. 2024. Symbol Graph Genetic Programming for Symbolic Regression. In *Parallel Problem Solving from Nature – PPSN XVIII: 18th International Conference, PPSN 2024, Hagenberg, Austria, September 14–18, 2024, Proceedings, Part I*. Springer-Verlag, Berlin, Heidelberg, 221–237. doi:10.1007/978-3-031-70055-2_14
- [46] Dirk Thierens. 1999. Scalability Problems of Simple Genetic Algorithms. *Evolutionary Computation* 7, 4 (Dec. 1999), 331–352. doi:10.1162/evco.1999.7.4.331
- [47] Dirk Thierens. 2010. The Linkage Tree Genetic Algorithm. In *Parallel Problem Solving from Nature, PPSN XI*, Robert Schaefer, Carlos Cotta, Joanna Kolodziej, and Günter Rudolph (Eds.). Springer, Berlin, Heidelberg, 264–273. doi:10.1007/978-3-642-15844-5_27
- [48] Dirk Thierens and Peter Bosman. 2012. Predetermined versus Learned Linkage Models. In *Proceedings of the 14th Annual Conference on Genetic and Evolutionary Computation (GECCO '12)*. Association for Computing Machinery, New York, NY, USA, 289–296. doi:10.1145/2330163.2330205
- [49] D. Pope Thomas Brooks. 1989. Airfoil Self-Noise. doi:10.24432/C5VW2C
- [50] Renato Tinós, Michal Przewozniczek, Darrell Whitley, and Francisco Chicano. 2023. Genetic Algorithm with Linkage Learning. In *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO '23)*. Association for Computing Machinery, New York, NY, USA, 981–989. doi:10.1145/3583131.3590349
- [51] Renato Tinós, Michal W. Przewozniczek, Darrell Whitley, and Francisco Chicano. 2024. Iterated Local Search with Linkage Learning. *ACM Trans. Evol. Learn. Optim.* 4, 2 (June 2024), 7:1–7:29. doi:10.1145/3651165
- [52] Marco Virgolin, Tanja Alderliesten, Cees Witteveen, and Peter A. N. Bosman. 2017. Scalable Genetic Programming by Gene-Pool Optimal Mixing and Input-Space Entropy-Based Building-Block Learning. In *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO '17)*. Association for Computing Machinery, New York, NY, USA, 1041–1048. doi:10.1145/3071178.3071287
- [53] Marco Virgolin, Tanja Alderliesten, Cees Witteveen, and Peter A. N. Bosman. 2021. Improving Model-based Genetic Programming for Symbolic Regression of Small Expressions. *Evolutionary Computation* 29, 2 (June 2021), 211–237. arXiv:1904.02050 [cs] doi:10.1162/evco_a_00278
- [54] Marco Virgolin and Peter A. N. Bosman. 2022. Coefficient Mutation in the Gene-Pool Optimal Mixing Evolutionary Algorithm for Symbolic Regression. In *Proceedings of the Genetic and Evolutionary Computation Conference Companion*. ACM, Boston Massachusetts, 2289–2297. doi:10.1145/3520304.3534036
- [55] Marco Virgolin, Eric Medvet, Tanja Alderliesten, and Peter A. N. Bosman. 2022. Less Is More: A Call to Focus on Simpler Models in Genetic Programming for Interpretable Machine Learning. arXiv:2204.02046 [cs] doi:10.48550/arXiv.2204.02046
- [56] Ekaterina J. Vladislavleva, Guido F. Smits, and Dick den Hertog. 2009. Order of Nonlinearity as a Complexity Measure for Models Generated by Symbolic Regression via Pareto Genetic Programming. *IEEE Transactions on Evolutionary Computation* 13, 2 (April 2009), 333–349. doi:10.1109/TEVC.2008.926486
- [57] L. Darrell Whitley, Francisco Chicano, and Brian W. Goldman. 2016. Gray Box Optimization for Mk Landscapes (NK Landscapes and MAX-kSAT). *Evolutionary Computation* 24, 3 (2016), 491–519. doi:10.1162/evco_a_00184
- [58] Garnett Wilson and Wolfgang Banzhaf. 2008. A Comparison of Cartesian Genetic Programming and Linear Genetic Programming. In *Genetic Programming*, Michael O'Neill, Leonardo Vanneschi, Steven Gustafson, Anna Isabel Esparcia Alcázar, Ivanoe De Falco, Antonio Della Cioppa, and Ernesto Tarantino (Eds.). Springer, Berlin, Heidelberg, 182–193. doi:10.1007/978-3-540-78671-9_16