

Search Strategy Generation for Branch and Bound Using Genetic Programming

Gwen Maudet¹, Grégoire Danoy²

¹SnT, University of Luxembourg, Esch-sur-Alzette, Luxembourg

²FSTM/DCS, SnT, University of Luxembourg, Esch-sur-Alzette, Luxembourg
{gwen.maudet,gregoire.danoy}@uni.lu

Abstract

Branch-and-Bound (BB) is an exact method in integer programming that recursively divides the search space into a tree. During the resolution process, determining the next subproblem to explore within the tree—known as the search strategy—is crucial. Hand-crafted heuristics are commonly used, but none are effective over all problem classes. Recent approaches utilizing neural networks claim to make more intelligent decisions but are computationally expensive. In this paper, we introduce GP2S (Genetic Programming for Search Strategy), a novel machine learning approach that automatically generates a BB search strategy heuristic, aiming to make intelligent decisions while being computationally lightweight. We define a policy as a function that evaluates the quality of a BB node by combining features from the node and the problem; the search strategy policy is then defined by a best-first search based on this node ranking. The policy space is explored using a genetic programming algorithm, and the policy that achieves the best performance on a training set is selected. We compare our approach with the standard method of the SCIP solver, a recent graph neural network-based method, and handcrafted heuristics. Our first evaluation includes three types of primal hard problems, tested on instances similar to the training set and on larger instances. Our method is at most 2 percents slower than the best baseline and consistently outperforms SCIP, achieving an average speedup of 11.3 percents. Additionally, GP2S is tested on the MIPLIB 2017 dataset, generating multiple heuristics from different subsets of instances. It exceeds SCIP’s average performance in 7 out of 10 cases across 15 times more instances and under a time limit 15 times longer, with some GP2S methods leading on most experiments in terms of the number of feasible solutions or optimality gap.

1 Introduction

Mixed-integer linear programming (MILP) is crucial for optimizing various sectors like transportation (Archetti, Peirano, and Speranza 2022; Yang et al. 2016), production (Tiwari et al. 2015), and e-health services (Rais and Viana 2011), by minimizing or maximizing an objective function under specific constraints. While exact algorithms ensure optimal solutions and quantify proximity to the optimum, they are computationally intensive. Branch-and-Bound (B&B), the standard algorithm for solving discrete

optimization problems, utilizes a “divide and conquer” approach, partitioning the solution space and constructing a tree where simpler subproblems are solved at deeper levels.

B&B has key components that guide decision-making: *cutting* (adding constraints), *branching* (choosing variables to split the search space), and *searching* (determining exploration order within the B&B tree). The solver aims to quickly find high-quality solutions and prove optimality. However, decision-making in B&B remains an open challenge, with no universal method effective across diverse instances (Zarpellon et al. 2021; Parsonson, Laterre, and Barrett 2023). The choice of search strategy (SS) (Ibaraki 1976; Morrison et al. 2016)—the focus of this paper—impacts solution updates and pruning efficiency, significantly affecting overall execution time. Current methods rely on manually crafted heuristics, which lack consistency across MILP instances (Linderoth and Savelsbergh 1999). Recent research has explored ML-based SS policies, often using neural networks to mimic oracle decisions (He, Daume III, and Eisner 2014; Song et al. 2018; Yilmaz and Yorke-Smith 2021; Labassi, Chételat, and Lodi 2022), but these approaches are constrained by the oracle’s limitations and the computational overhead of neural networks.

This paper introduces *GP2S* (Genetic Programming for Search Strategy), a method designed to automatically generate a SS policy using genetic programming (GP), aiming to be both lightweight and efficient. We define node scoring functions based on operations applied to data related to the node and the specific problem. A SS heuristic is then implemented using a best-first search (BFS) approach, which iteratively prioritizes nodes based on their scores. The space of SS heuristics is explored through GP, with the quality of each policy assessed on a training set. The method is presented in Section 4.

Our approach, integrated into the SCIP solver, is benchmarked against SCIP’s native method, a recent a neural network-based imitation learning algorithm (Labassi, Chételat, and Lodi 2022), and three state-of-the-art heuristics. In the first phase, we evaluate performance on three primal hard problem types: Fixed Charge Multicommodity Network Flow (FCMCNF), Maximum Satisfiability (MAXSAT), and Generalized Independent Set (GISP). Each method is evaluated on instances comparable in size to the training set as well as on larger instances. Our method

achieves solving times within 2% of the best baseline for each instance set. Compared to the SCIP, *GP2S* consistently performs equally or better, with an average speedup of 11.3%. In the second phase, we apply *GP2S* to MIPLIB 2017, training multiple GP-based SS heuristics on 50 instances with a 10-second time limit. Seven out of ten heuristics outperform SCIP at a 150-second time limit, with some methods excelling in terms of the number of feasible solutions or optimality gap across various instance sets and time limits. Detailed experimental results are provided in Section 5.

The contributions of this paper are as follows:

- Definition of a space of SS heuristics using a scoring function that integrates node and problem data, including existing handcrafted heuristics.
- Design of a GP algorithm to explore this space, aiming to improve B&B performance on a training set of instances.
- When trained on a specific problem type, *GP2S* performs within 2% of the best baseline and achieves an average speedup of 11.3% over SCIP, even on larger instances than those used for training.
- Evaluation on the MIPLIB 2017 instance set, contrasting with the focus on synthetically generated problems in most evaluations.
- On MIPLIB 2017, *GP2S* surpasses SCIP in 7 out of 10 cases on 15 times more instances than those used for training, with a time limit 15 times longer, with some GP-based methods excelling in terms of feasible solutions or optimality gap on both reduced and full sets across different time limits.

2 Background

SS Policies in B&B A MILP can be formulated as $\min_{x \in N_0} \{c^T x : Ax \geq b\}$, where $N_0 = \mathbb{Z}^k \times \mathbb{R}^{n-k}$, $A \in \mathbb{R}^{m \times n}$, $b \in \mathbb{R}^m$ and $c \in \mathbb{R}^n$ are vectors of constants. The B&B algorithm addresses this problem through the following steps: (i) relaxation: solve the relaxed problem over $x \in N_i$ (initially $N_i = N_0$). Let $x_i = (x_i^j)_{1 \leq j \leq n}$ denote the optimal solution and z_i represent the objective value; (ii) feasibility check: if x_i satisfies the integrality constraints, update the best objective value to $z^* = \min(z^*, z_i)$; if x_i does not meet the integrality constraints, z_i serves as a lower bound (LB); if z_i exceeds z^* ($z^* < z_i$), the problem is *pruned*; (iii) branching: if the problem is not pruned, select a variable x^j (for $1 \leq j \leq k$) that fails to meet integrality constraints; generate two subproblems: $N_i^{j-} = N_i$, $x^j \leq \lfloor x_i^j \rfloor$ and $N_i^{j+} = N_i$, $x^j \geq \lfloor x_i^j \rfloor + 1$; additionally, denote d_i as the total number of times a variable has been selected as a branching variable to define N_i , so $d_i^{j+} = d_i + 1$ and $d_i^{j-} = d_i + 1$; here, d_i represents the depth of the node in the tree; (iv) exploration: continue by selecting a new subproblem from the remaining unexplored domains N_i and terminate once all subproblems are either explored or pruned. The process of selecting the next node to explore

is known as the *search strategy* (SS). The efficiency of finding the best-known solution z^* is significantly impacted by the sequence in which nodes are explored, with faster convergence to high-quality solutions enhancing the effectiveness of subsequent pruning. A SS policy can be characterized by two components. First, the *node-comparing* method establishes a relationship (dominance or equivalence) between two nodes. Second, the *node-selection* method determines which node to explore, typically selecting the highest-ranked node (defining a BFS) or through a more complex paradigm. This structure is exemplified in SCIP (Bolusani et al. 2024), which uses a similar decomposition.

GP Framework An evolutionary algorithm is a population-based metaheuristic where a population of individuals is evaluated and evolved using operators to enhance their quality assessed via a fitness function. A key feature of GP is that each individual is defined as a computer program. When these individuals represent heuristics, GP explores the heuristic space, thus classifying it as a generative hyper-heuristic (Dokeroglu, Kucukyilmaz, and Talbi 2024). An evolutionary algorithm, including GP, follows a generic structure for defining and evolving individuals within a population. After generating the initial population, the population is refined through multiple iterations, or generations, using three primary evolutionary processes, typically in the following order: (i) *selection*, which determines which individuals are chosen for the next processes according to their fitness and how many offspring each selected individual produces; (ii) *crossover*, which combines two or more individuals to generate new ones; and (iii) *mutation*, which modifies an individual. The evolutionary process concludes based on criteria such as the number of generations, elapsed time, or convergence rate, after which the best individual(s) are returned. For a more detailed understanding of evolutionary methods, refer to (Talbi 2009). In GP, individuals are often structured as trees, which simplifies the representation of mutations and crossovers; the programs are reassembled recursively (Koza 1989). The individuals in this structure include two main types of components: *terminals*, found at the leaves of the tree and acting as the building blocks of the program, and *operators*, located at non-leaf nodes, which perform operations on one or more terminals depending on their arity. Typically, all terminals are of the same type, and operators produce results of the same type.

3 Related Work

Handcrafted SS Heuristics The literature on SS for B&B includes various handcrafted heuristics employing different methods for node comparison and selection (Lawler and Wood 1966; Morrison et al. 2016; Tomazella and Nagano 2020). Typically, a node is evaluated based on features derived from the relaxation of the problem, with node comparison involving these feature-based scores. A common approach is to prioritize nodes with the smallest LB z_i , assuming that if the relaxed problem yields a promising solution, the original problem is likely to follow suit. An alternative

method introduced by (Benichou et al. 1971) provides an estimate of the objective function’s value. For a branching variable x^j , its fractional part is $f_i^j = x_i^j - \lfloor x_i^j \rfloor$. The objective values of the related subdomains N_i^{j-} and N_i^{j+} are denoted as z_i^{j-} and z_i^{j+} , respectively. Pseudocosts are computed as $P_i^{j-} = (z_i^{j-} - z_i)/f_i^j$ and $P_i^{j+} = (z_i^{j+} - z_i)/(1 - f_i^j)$. The metric, known as the *best estimate* at node N_i , is given by $BE_i = z_i + \sum_{1 \leq j \leq k} \min(P_i^{j-} f_i^j, P_i^{j+} (1 - f_i^j))$. For node selection, the best node is generally chosen from a set of open nodes, with ranking determined by the node-comparing method. The set of open nodes can include all nodes for BFS, offspring nodes relative to the active node for depth-first search (DFS), or siblings for breadth-first search, with potential transitions from DFS to BFS upon reaching a leaf node.

No single method dominates the others for node evaluation or node selection. For example, BFS is theoretically more efficient but can significantly increase the open node list size, leading to scheduling and memory issues that slow down the process. (Linderoth and Savelsbergh 1999) evaluated 13 handcrafted SS heuristics, integrating the node-comparing and node-selection methods discussed. These heuristics are highly interpretable and computationally efficient but are manually constructed, indicating that automating their generation could potentially enhance performance.

ML for MILP Recent advancements have seen ML being increasingly applied to MILP, as summarized in several surveys. (Lodi and Zarpellon 2017) reviews ML contributions to B&B SS policies, focusing on approaches using imitation learning and reinforcement learning. (Bengio, Lodi, and Prouvost 2020) provides an in-depth survey on the intersection of ML and combinatorial optimization, exploring various integrations between the two fields and different ML learning objectives. (Cappart et al. 2022) discusses the applications of graph neural networks (GNN) to optimization problems. Furthermore, (Scavuzzo et al. 2024) examines ML’s role in developing primal heuristics, branching strategies, SS policies, cutting planes, and other solver configurations.

Classification Method Applied to Node Selection The node-comparing method can be framed as a classification problem with three potential outcomes: one node dominating another, the other node dominating, or equivalence between the nodes. The decision should be made based on the current state, which includes the compared nodes, the problem’s features, and the current progress of the B&B process. Classification models are trained on previously solved instances with the objective of identifying the dominant node that leads to faster resolution, using imitation learning on an oracle that is presumed to be the perfect solver. These methods are applied to obtain a fixed number of feasible solutions before switching to a handcrafted method, such as best-estimate BFS, to complete the resolution. (He, Daume III, and Eisner 2014) introduced a selection method and a pruning technique based on support vector machines and the DAGger algorithm. (Yilmaz and Yorke-Smith 2021; Song et al. 2018) trained a multilayer perceptron model

with the same feature as (He, Daume III, and Eisner 2014), while (Labassi, Chételat, and Lodi 2022) employed a GNN. Each of these methods has its specific characteristics, but the GNN approach is generally more versatile, enabling it to handle problems of different sizes. However, imitation learning is confined to predefined methods within the solver. Additionally, methods based on neural networks require significant computational resources for each node comparison and function as black boxes. Furthermore, these methods are only applied during the initial stages of tree exploration, and thus cannot be considered comprehensive SS policies.

GP Based on Node Scoring GP is also a branch of ML, though it remains relatively unexplored in combination with MILP; no works related to GP are reported in the following surveys (Lodi and Zarpellon 2017; Bengio, Lodi, and Prouvost 2020; Cappart et al. 2022; Scavuzzo et al. 2024). However, there are preliminary works proposing the application of GP for SS in B&B. In two papers (Kostikas and Fragakis 2004; Morikawa, Nagasawa, and Takahashi 2019), GP is used to define a function that evaluates the quality of a node, enabling the construction of a SS policy by iteratively selecting the node with the highest quality. (Kostikas and Fragakis 2004) introduced a two-phase B&B approach. The first phase involves a traditional exploration of the search space, followed by a second phase where a GP-based method is employed to develop a new SS policy. The fitness function for this GP approach is defined based on improvements in integer infeasibilities during dives, according to the tree constructed in the first phase. (Morikawa, Nagasawa, and Takahashi 2019) developed a GP scoring method using terminals specific to job shop scheduling.

Positioning Our methodology leverages a generic hyper-heuristic GP approach akin to (Kostikas and Fragakis 2004), with key distinctions. Unlike their two-phase resolution process tailored to individual instances, our GP build heuristics that generalize effectively across diverse instance sets. While their implementation in GLPK omits critical details—such as the initial SS, phase-transition criteria, and GP configurations—we integrate our method with the modern SCIP solver and provide an open-source implementation. Experimentally, their approach shows limited performance gains and is at least twice as slow as contemporary state-of-the-art heuristics, highlighting the inefficiency of two-phase strategies. Unlike the neural network-based approaches proposed by (Yilmaz and Yorke-Smith 2021; Song et al. 2018; Labassi, Chételat, and Lodi 2022), our focus is on constructing an SS policy that is both computationally lightweight for node comparison and interpretable. Since our heuristic space encompasses most handcrafted heuristics, the goal is to develop a policy that is smarter than those previously discovered. This approach is classified under “Learning to configure algorithms” with a “multi-instance” learning objective, as defined in (Bengio, Lodi, and Prouvost 2020).

4 GP-based Node Scoring Function

We now introduce *GP2S*, an ML method that generate a SS policy for B&B through GP. A policy is represented as the scoring of a B&B node, coupled with a BFS approach that always selects the node with the best score. The scoring function has a tree structure, and we explore the space of scoring functions using the GP method. Firstly, we delineate the SS policy space, followed by an explanation of our choices for the evolutionary algorithm tuning.

Definition of the SS Policy Space

We define the space of scoring functions using standard operations on features derived from both the node and the problem. The SS is then constructed by combining these functions with a BFS-based node-selection method. Since the scoring function is real-valued, we utilize terminals of floating-point types and operators tailored for real number operations, as summarized in Table 1. The operators are selected as standard arithmetic functions over real numbers. To integrate our approach within the SCIP solver, we employ real-type variables that are either node-specific or model-specific—commonly used attributes for characterizing nodes and models. Node-related variables vary from node to node, while model-related variables differ between instances but remain constant within a single instance. Additionally, we introduce the constant M , a sufficiently large positive value. This allows for the definition of multiple levels of scoring, particularly when the scoring function consists of the sum of several components, each weighted by different powers of M . For instance, the SS policy combining DFS with the lowest LB can be represented in this way: depth is prioritized first, and in the case of a tie, the node with the lowest LB is selected. The corresponding minimization function is expressed as $f(N_i) = z_i - M \times d_i$. This construction significantly expands the SS policy space beyond manually defined heuristics. Given α as the number of 2-arity operators and β as the number of terminals, $\beta^{2^r} \times \alpha^{2^r - 1}$ perfect trees of size r can be generated. As a numerical example, using the operators and terminals detailed in Table 1, approximately 10^{11} different trees of depth 3 can be obtained.

OPERATORS		TERMINALS	
Sym	Meaning	Sym	Name
			Node-related
+	Add two inputs	d_i	Depth
-	Subtract two inputs	BE_i	Best estimate
		z_i	Lower bound
			Model-related
×	Multiply two inputs	z_0	Dual bound at the root
	Divide with protection: if division by 0, return 1	m	Number of constraints
		n	Number of variables
			Constant
		M	Big positive constant

Table 1: Operators and terminals.

GP Framework

We base our GP approach on the framework outlined in (Baeck, Fogel, and Michalewicz 2000), which includes pa-

rameters such as population size p , number of generations g , and probabilities for crossover $P_{\text{mate}} \in [0, 1]$ and mutation $P_{\text{mutate}} \in [0, 1]$. The process begins with the creation of an initial population of size p . Each generation starts with the selection phase, where individuals are chosen (with possible duplication) to maintain the population size. During the evolution phase, individuals are paired for crossover with probability P_{mate} . Following crossover, mutation occurs where each individual in the population has a probability of P_{mutate} of undergoing mutation.

Fitness Function A crucial aspect of defining an evolutionary algorithm is the characterization of an individual’s fitness, which signifies the desired traits and is refined through the evolutionary process, serving as a discriminating factor during selection. In our case, we may aim to define an SS policy for either complete or incomplete instance resolution (typically under a time limit), which leads to two distinct methods for evaluating the policy’s performance on an instance. For each performance measure defined on the training set instances, the fitness function is based on the 1-shifted geometric mean, a standard metric for assessing average performance on a benchmark (Achterberg 2007). For the complete resolution objective, as examined in the first part of the simulation, the focus is on solving instances as quickly as possible, so solving time is used as the criterion. In the case of incomplete resolution, explored in the second part of the simulation, the performance of the best-found solution (if any) is assessed using the *optimality gap*. This gap measures the relative difference between the best integer solution and the lowest LB among unexplored nodes. The lowest LB is defined as $\text{Best LB} = \min_{i, N_i \text{ unexplored}}(z_i)$, and the gap is calculated as $\text{Gap} = |(z^* - \text{Best LB}) / (\min(z^*, \text{Best LB}))|$. If the optimal solution is found, then $\text{Gap} = 0$. By convention, if no solution is found, the optimality gap is set to $1e + 20$ (as returned by the SCIP solver). This fitness function is structured to first minimize the number of infeasible solutions and then to speed up convergence toward a better solution.

GP Population Initialization Initially, a population of p individuals, represented as trees, is generated. The generation of an individual begins at the root node, whose nature is determined by selecting uniformly at random from the set of operators and terminals. If an operator is chosen, its child nodes are generated, and the process continues. If a terminal is selected, the branch is completed. Each branch is constrained to have a minimum depth of $D_{\text{init-min}}$ and a maximum depth of $D_{\text{init-max}}$.

Selection For selection, we use a double tournament method as defined in (Luke and Panait 2002), designed to avoid convergence towards excessively deep trees. Deep trees can incur higher costs during node score calculations and are particularly prone to overfitting, whereas simpler formulations tend to be more general. In the first tournament, of size k , we randomly select k individuals from the GP population and retain the best one based on one of the previously defined fitness criteria. After performing two fitness-based tournaments, the two winning candidates proceed to a sec-

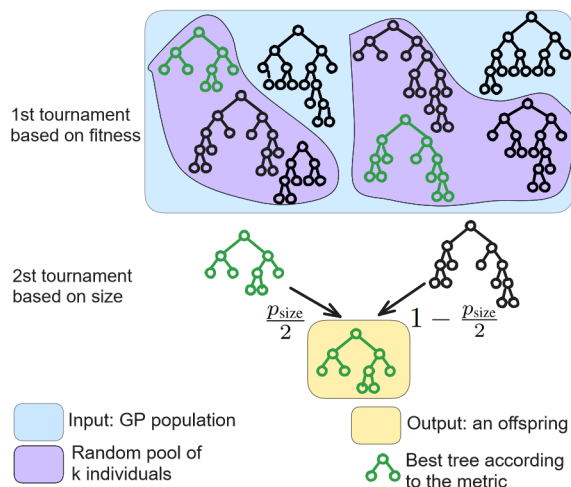


Figure 1: Selection process based on double tournament: this step is repeated p times to generate p offspring.

ond round, where the deciding factor is the total number of nodes in the tree. for this second round, to favor the selection of smaller trees, the smaller tree has a probability $\frac{P_{size}}{2}$ of being selected, where $P_{size} \in [1, 2]$. An illustration of the double tournament selection process is presented in Fig. 1.

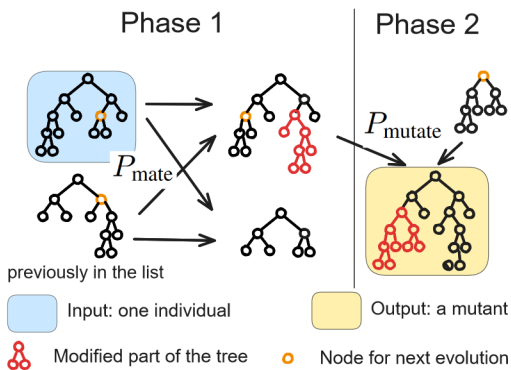


Figure 2: Crossover and mutation processes: the whole population completes phase 1, then moves to phase 2.

Crossover and Mutation After the selection process comes the crossover, then mutation. If a crossover operation is performed for an individual (with probability P_{mate}), another individual (typically the previous one in the selection list) is chosen from the population. A one-point crossover is then executed by selecting a node in each tree and swapping the subtrees rooted at these nodes, thereby modifying both offspring. If no crossover occurs, the individual remains unchanged. Next, we perform the mutation process on the modified population. With probability P_{mutate} (otherwise unchanged with probability $1 - P_{mutate}$), a uniform mutation is applied: the new tree is defined with leaves bounded by the limits $D_{mut-min}$ and $D_{mut-max}$. A node is chosen randomly such that the new tree grafts at this point, replacing the previ-

ous subtree. The crossover and mutation processes are summarized in Fig. 2.

5 Experiments

Our experiments are conducted in two phases. In the first phase, *GP2S* is evaluated on three synthetic problem types, testing performance and generalizability on instances similar in size to the training set and larger instances. In the second phase, *GP2S* heuristics are generated from a subset of MIPLIB 2017 instances under a solving time limit and evaluated on the full dataset with extended time limits. This phase examines the heuristic’s generalization to a heterogeneous dataset and efficiency within constrained time. All simulations use the SCIP solver (Bolusani et al. 2024; Maher et al. 2016), with complete code available on the Git repository: <https://gitlab.com/uniluxembourg/snt/pcog/ultrabo/search-strategy-generation-for-branch-and-bound-using-genetic-programming.git>. Detailed integration of our method into SCIP is also provided within the framework.

As we are evaluating performance within SCIP, we use the default SCIP SS policy as the primary baseline, making no modifications to the solving core (denoted as *SCIP* in the performance evaluation table). We further compare our approach to several manually crafted heuristics. Specifically, we include the method that ranks nodes by their LB in a BFS (*LB BFS*), as well as the best estimate in BFS (*BE BFS*) and DFS (*BE DFS*). Additionally, for the first part of the experiment, we include a recent GNN-based SS method (Labassi, Chételat, and Lodi 2022). The original GNN SS policy is implemented in two phases: the trained GNN is used until two solutions are found, after which the policy switches to *BE BFS*. For comparison, we include this approach (*GNN 2 dives*), as well as the variant where the trained GNN-based node-comparison method is applied throughout the entire resolution process (*GNN full*). This method requires the complete resolution of training instances so that the GNN learns to replicate decisions leading to the best solutions. However, because complete B&B resolution is not guaranteed for MIPLIB 2017 instances, this method is not included in the comparisons for the second part of the experiment.

GP2S is a parametric approach that requires defining evolutionary features. We have chosen classical parameters that are widely used in the development of such methods (Ahvanooy 2019; Luke and Panait 2002; Duffo et al. 2019): the probabilities for crossover and mutation are set to $P_{mate} = 0.9$ and $P_{mutate} = 0.1$; the depth of population initialization is restricted by minimum and maximum limits of $D_{init-min} = 1$ and $D_{init-max} = 17$; the depth of branches generated during mutation is bounded by $D_{mut-min} = 1$ and $D_{mut-max} = 5$; the parameter related to the selection of the smallest tree size is set to $P_{size} = 1.2$. The parameters governing the number of individuals and generations are detailed in each part of the experiment, as these factors greatly influence the computational cost of the solution. To visually compare the performance of each baseline across different sets of instances and evaluation metrics, we color the text in the performance table: green indicates the best results, red

PROBLEM	FCMCNF		MAXSAT		GISP	
	TEST	TRANSFER	TEST	TRANSFER	TEST	TRANSFER
BE DFS	2.9±1.5	17.4±1.9	5.2±1.9	9.6±1.9	3.4±1.5	34.3±1.7
BE BFS	3.0±1.5	18.6±1.9	6.3±1.8	12.6±1.8	2.7±1.6	26.2±1.7
LB BFS	3.0±1.5	20.0±2.0	4.6±1.8	10.2±2.1	2.9±1.6	25.6±1.5
GNN 2 dives	3.8±1.4	20.0±1.9	5.7±1.8	11.0±2.1	3.6±1.5	25.6±1.5
GNN full	4.0±1.6	27.3±2.4	4.9±1.9	10.5±2.3	3.5±1.6	31.3±1.5
SCIP	3.3±1.5	20.0±1.9	5.8±1.7	10.3±1.8	2.4±1.6	20.9±1.6
GP2S	3.0±1.5	17.7±1.9	4.5±1.8	8.9±2.1	2.4±1.6	20.9±1.6

Table 2: Solving time with standard deviation (1-shifted geometric mean and geometric standard deviation) for our proposal and baselines on three problem types, for both training-similar (test) and larger instances (transfer).

marks results that are more than 40% worse than the optimum, and a color gradient is applied for intermediate results.

Solving Time for Specific Types of Problem

Types of Problem Solved We benchmark our method against the SS policy from (Labassi, Chételat, and Lodi 2022) using similar experimental settings, including problem types and instance sizes, that enable indirect comparisons with methods evaluated in their study (He, Daume III, and Eisner 2014; Yilmaz and Yorke-Smith 2021; Song et al. 2018). We evaluate three challenging primal problem types, each with 50 instances: a training set for development, a test set with similar instances for effectiveness, and a transfer set with larger instances for generalizability. The problem types are: (i) FCMCNF (Hewitt, Nemhauser, and Savelsbergh 2009), with $n = 15$ nodes for training and testing, and $n = 20$ for transfer, with $m = 1.5n$ commodities (Chmiela et al. 2021); (ii) MAXSAT, with $n \in [60, 70]$ for training and testing, and $n \in [70, 80]$ for transfer (Béjar et al. 2009); (iii) GISP, with $n \in [60, 70]$ for training and testing, and $n \in [80, 100]$ for transfer (Chmiela et al. 2021). All problems are defined on Erdos-Rényi random graphs (Erdos 1961) with $p = 0.3$ for FCMCNF and $p = 0.6$ for MAXSAT and GISP. Each instance is solved to optimality.

Hardware and GP2S Setup In the first phase, a neural network-based method is used as the baseline, requiring GPU-equipped hardware. Each run, involving solving an instance with a baseline SS policy, is executed on a node with 7 CPU cores, 192 GB of RAM, and an NVIDIA V100 GPU with 16 GB of RAM.

For the *GP2S* setup, the method uses 50 individuals over 50 generations with a tournament size of $k = 5$. The fitness function, as detailed in Section 4, evaluates the GP-based SS policy using the geometric mean solving time across 50 training instances for each problem. GP convergence plots and scoring functions are presented in Appendix B.

Evaluation and Discussion For each method, problem type, and instance size, we record the solving times across all instances. The average solving times, calculated using the 1-shifted geometric mean, are presented in Table 2.

Among the handcrafted heuristics, none demonstrate consistent performance across all studied problem types. Each heuristic tends to be relatively effective on two problems but

slower on the third. For the GNN-based methods, *GNN 2 dives* is preferable, as the *GNN full* approach can significantly slow down solving times, particularly for the FCMCNF transfer instances. The *GNN 2 dives* method offers stable results for the transfer partition, avoiding extreme performance deviations, making it a more reliable SS policy compared to handcrafted heuristics. Similar results are observed for SCIP, which achieves notably better performance on GISP problems. Regarding our proposed method, *GP2S*, the solving time is at most 2% slower than the best SS baseline. It either surpasses or matches SCIP’s performance across all categories, achieving an average speedup of 11.3% across the six categories. Thus, *GP2S* outperforms all other SS policies considered when trained on a specific problem type. It is important to note that solving an instance takes several seconds, making the total time for the GP algorithm to develop a SS policy for one problem $O(50^3)$ seconds, which translates to several days of computation. This computational effort is a one-time investment and is beneficial, as it significantly accelerates resolution for problems of such types compared to the other baselines. Furthermore, as our simulation results show, *GP2S* can be trained on simpler problems while maintaining high efficiency on more complex ones.

Time-Limited Performance on MIPLIB2017

Partitions and Time Limits We conducted experiments using the MIPLIB 2017 collection (Gleixner et al. 2021), excluding instances unlikely to yield meaningful results under time constraints. We removed instances tagged as *feasibility*, *numerics*, *infeasible*, *no solution*, as well as those with presolve times exceeding 300 seconds, those unsolved within 600 seconds, or solved at the root, resulting in 758 instances. GP methods were trained with a 10-second solving time limit. To focus on informative cases, we selected a subset of 166 instances where SCIP, in standalone mode, explores more than one node within 10 seconds. Experiments were conducted with time limits ranging from 10 seconds to 150 seconds for both subsets.

Hardware and GP2S Setup The method in (Labassi, Chételat, and Lodi 2022) is not considered in this phase, as it requires a training set of pre-solved instances, which is not available for all MIPLIB 2017 instances. A CPU-only architecture is used, with a dedicated node featuring 28 cores

	REDUCED MIPLIB 2017: 166 instances						ALL MIPLIB 2017: 758 instances					
	10 seconds		50 seconds		150 seconds		10 seconds		50 seconds		150 seconds	
	INF	GAP	INF	GAP	INF	GAP	INF	GAP	INF	GAP	INF	GAP
BE BFS	34	0.72±4.0	33	0.62±3.5	26	0.45±3.8	455	2.08±5.8	289	0.75±4.4	249	0.69±4.4
BE DFS	34	0.75±4.1	28	0.44±3.9	21	0.33±4.5	391	1.06±4.5	297	0.77±4.4	240	0.72±4.5
LB BFS	51	1.26±4.4	30	0.39±4.1	30	0.4±4.0	398	1.03±4.3	298	0.74±4.4	304	1.14±4.8
SCIP	38	0.76±4.1	26	0.41±4.0	26	0.41±4.1	393	1.06±4.5	289	0.76±4.4	293	1.14±4.8
GP2S 2	35	0.73±4.0	27	0.42±3.9	26	0.34±4.3	454	2.08±5.8	297	0.75±4.4	254	0.66±3.9
GP2S 4	49	1.25±4.4	30	0.47±3.8	30	0.38±4.2	394	1.0±4.3	299	0.76±4.3	304	1.16±4.7
GP2S 6	34	0.71±4.1	27	0.44±3.9	25	0.36±4.2	392	1.04±4.5	297	0.77±4.3	253	0.67±3.9
GP2S 8	33	0.76±4.1	24	0.47±3.8	21	0.42±3.9	390	1.07±4.5	285	0.79±4.3	290	1.19±4.7
GP2S 10	37	0.73±4.1	28	0.44±3.9	24	0.35±4.3	395	1.03±4.4	298	0.77±4.4	249	0.67±4.0
GP2S 12	35	0.75±4.0	22	0.49±3.8	16	0.41±4.0	393	1.04±4.3	283	0.78±4.3	232	0.75±4.3
GP2S 14	36	0.76±4.1	29	0.48±3.8	29	0.5±3.7	392	1.05±4.4	297	0.8±4.3	251	0.75±4.4
GP2S 16	35	0.72±4.1	30	0.39±4.1	29	0.28±5.0	450	2.07±5.8	300	0.75±4.4	246	0.68±4.4
GP2S 18	36	0.77±4.1	25	0.48±3.8	24	0.39±4.2	394	1.07±4.5	362	1.48±5.1	296	1.18±4.7
GP2S 20	37	0.75±4.0	27	0.5±3.8	21	0.38±4.2	450	2.09±5.8	284	0.77±4.4	236	0.7±4.4

Table 3: Number of instances with no feasible solution (Inf) and geometric mean gap (Gap) for instances where all SS baselines found a solution. GP-based functions are labeled as GP2S+seed. Results are shown for reduced and full MIPLIB 2017 sets at time limits of 10, 50, and 150 seconds.

and 128 GB of RAM to solve instances using SS baseline policies.

From the reduced set, we randomly select 50 instances with a 10-second time limit for the GP process. To account for the heterogeneity of MIPLIB 2017, we perform multiple independent GP runs to vary the training set through different seeds. To manage computational demands, we set the population size to 20, limit generations to 20, and use a fitness tournament size of $k = 3$. The fitness function, detailed in Section 4, uses the 1-shifted geometric mean gap (assigning $1e + 20$ when no integer solution is found). The resulting scoring functions are presented in Appendix C.

Evaluation and Discussion To assess solution quality, we use two metrics: the number of instances with no feasible solution (where $z^* = +\infty$) and the 1-shifted geometric mean optimality gap for instances where all SS baselines found a feasible solution. Results for 10 GP methods are presented with time limits of $\{10, 50, 100\}$ in Table 3.

No single method consistently outperforms all others across both metrics for any given time limit. However, *GP2S* surpasses *SCIP* in at least 7 out of 10 instances, both in terms of infeasibility and optimality gap, for both the reduced and full sets at a 150-second time limit. The handcrafted heuristics generally perform well, with *BE DFS* showing particularly stable results, never performing worse than 32% relative to the best performance, making it the top method in this regard. Some GP-based methods demonstrate remarkable effectiveness: *GP2S 12* has a number of unresolved instances only 5% higher than the best-performing solution, while *GP2S 16*, except for the entire MIPLIB at a 10-second time limit, achieves an optimality gap within 1% of the best-known solution; these methods are the best according to these metrics. Given that GP-based SS policies are trained on a small pool of diverse instances, their performance can be highly variable. Yet, despite being automatically generated in just a few days of computation without any assumptions, and tested on 15 times more instances with a time limit

up to 15 times longer, they perform very well compared to expert-engineered methods refined over years.

6 Conclusion and Perspectives

In this paper, we have examined the SS component in B&B, a crucial factor affecting the solving time of MILP solvers. We presented *GP2S*, a method that automatically generates a heuristic from a training set, implemented in the open-source solver *SCIP*. We trained *GP2S* on well-known synthetic problem types and demonstrated its significant superiority over the baselines, even when tested on instances larger than those used in training. When applied to the MIPLIB 2017 collection, *GP2S* generated a variety of heuristics based on different training instance pools. Generally, these heuristics outperformed *SCIP*, with some methods surpassing all baseline approaches. Given the complexity of generating a single heuristic for all MILP problems—due to challenges such as defining the training pool and achieving generalizability—a promising direction for future work would be to define multiple problem subsets and generate a heuristic tailored to each specific subset. Another significant highlighted of this paper is the effective use of the GP paradigm to define features within an MILP solver. While current trends in optimization often rely on neural networks, these approaches have notable drawbacks, including their dependency on oracles, their opaque “black-box” nature, and their inconsistent performance, which may arise from the expansive search space and significant computational overhead. In contrast, methods like ours, with their simpler and more constrained search spaces, offer faster convergence to high-quality solutions. Additionally, our approach is designed to be highly efficient, imposing minimal additional cost on the overall solving process. Future work could explore the application of automatic heuristic search methods to other solver components, such as branching strategies or cutting plane selection.




Figure 3: 1-shifted geometric mean solving time over the training set of the best SS policy in the population across generations for: (a) FCMCNF, (b) MAXSAT, and (c) GISP.

Problem	GP function
FCMCNF	$\frac{BE_i}{d_i}$
MAXSAT	$\frac{BE_i + M}{d_i}$
GISP	$\frac{d_i}{BE_i} - d_i$

Table 4: Scoring functions produced by GP from training sets for three problem types.

Acknowledgements

This research was funded by the Agence Nationale de la Recherche (grant ANR-22-CE46-0011) and the Luxembourg National Research Fund (grant INTER/ANR/22/17133848) as part of the UltraBO Project. The computational experiments were conducted using the University of Luxembourg’s HPC facility.

We express our gratitude to Mark Turner for his support with SCIP and the MIPLIB 2017 library, and to Abdel Labassi for his guidance in using his framework. We are also deeply grateful to Pierre Talbot for the valuable informal discussions that greatly improved the quality of this work. Finally, we appreciate the constructive feedback from the anonymous reviewers, which significantly contributed to this revised version.

Appendices

A Hardware and Software Settings

We used SCIP 9.1 along with the DEAP 1.4.1 package as the framework for implementing our GP algorithm. Additionally, we integrated the method from (Labassi, Chételat, and Lodi 2022) as a baseline. For their method, we retained all parameters used in the original paper. This baseline relies on PyTorch, and we used PyTorch 2.4.0 for our implementation.

To specify the hardware setup used in each simulation phase, we utilized the following:

- For the first phase, requiring GPU resources: 1 Xeon Gold 6132 @ 2.6GHz [14c/140W], 2 Tesla V100 SXM2 16G.
- For the second phase, which only required CPU resources: either 2 Xeon E5-2680v4 @ 2.4GHz

[14c/120W] or 2 Xeon Gold 6132 @ 2.6GHz [14c/140W] nodes.

B GP Behavior for Specific Problem Types

We applied GP methods to the FCMCNF, MAXSAT, and GISP problems. Due to the extensive runtime of each GP method, we did not conduct robustness tests by repeatedly running GP on the same problem. However, the SS policies derived for each problem type can be considered effective, as they demonstrate strong performance compared to the other baselines.

In Fig. 3, we present convergence graphs showing the fitness function (1-shifted geometric mean solving time) of the best individual among the 50 considered for each generation. These graphs allow us to assess the convergence speed of the method towards the optimal solution. Additionally, in Table 4, we display the scoring function that performed best on the training set for each problem studied. Notably, as seen in Fig. 3(a), the best scoring method for FCMCNF converges very quickly with little subsequent improvement. This result is due to the final method representing a best estimate BFS; the convergence is consistent, as seen in Table 2, where this heuristic is the best-performing baseline. For the other methods, the final function is slightly more complex (using three terminals rather than two) and is found after a greater number of generations. Interestingly, for MAXSAT and GISP, the *GP2S* method prioritizes higher values of the best estimate (since they are divisors in the formula).

C GP Scoring Functions for MIPLIB 2017

We now present the functions generated by GP across the 10 seeds in Table 5. It is important to note once again that the set of instances has a significant impact on the convergence of the GP method. We observe that, due to the increased heterogeneity of the instances (including varying instance sizes and problem types), the final scoring functions are considerably more complex, incorporating a greater number of problem-related variables.

Seed	GP Function
1	$z_i + \frac{n}{d_i - n}$
2	$BE_i + z_i + m$
4	$\frac{z_i}{z_0} + n \times d_i$
6	$(BE_i + z_i) \times (BE_i + z_i)$
8	$z_0 \times (M + M - z_i) - BE_i - d_i \times z_i$
10	$\frac{z_i}{\frac{m}{d_i}}$
12	$\frac{z_i}{d_i \times (n + BE_i + d_i)}$
14	$z_i \times \frac{z_i}{z_i} \times (BE_i - n)$
16	$z_i + z_i$
18	$M - z_i \times d_i + BE_i$
20	$\frac{n - m}{BE_i}$

Table 5: Functions generated by GP from training sets based on MIPLIB 2017 instances.

References

Achterberg, T. 2007. *Constraint Integer Programming*. Ph.D. thesis, ZIB Berlin.

Ahvanooey. 2019. A Survey of Genetic Programming and Its Applications. *KSII Transactions on Internet and Information Systems*, 13(4).

Archetti, C.; Peirano, L.; and Speranza, M. G. 2022. Optimization in multimodal freight transportation problems: A Survey. *European Journal of Operational Research*, 299(1): 1–20.

Baeck, T.; Fogel, D. B.; and Michalewicz, Z. 2000. *Evolutionary Computation 1: Basic Algorithms and Operators*. CRC Press. ISBN 978-0-7503-0664-5. Google-Books-ID: 4HMYCq9US78C.

Bengio, Y.; Lodi, A.; and Prouvost, A. 2020. Machine Learning for Combinatorial Optimization: a Methodological Tour d’Horizon. ArXiv:1811.06128 [cs, stat].

Benichou, M.; Gauthier, J. M.; Girodet, P.; Hentges, G.; Ribiere, G.; and Vincent, O. 1971. Experiments in mixed-integer linear programming. *Mathematical Programming*, 1(1): 76–94.

Bolusani, S.; Besançon, M.; Bestuzheva, K.; Chmiela, A.; Dionisio, J.; Donkiewicz, T.; Doormalen, J. v.; Eifler, L.; Ghannam, M.; Gleixner, A.; Graczyk, C.; Halbig, K.; Hedtke, I.; Hoen, A.; Hojny, C.; Hulst, R. v. d.; Kamp, D.; Koch, T.; Kofler, K.; Lentz, J.; Manns, J.; Mexi, G.; Mühmer, E.; Pfetsch, M. E.; Schlösser, F.; Serrano, F.; Shinano, Y.; Turner, M.; Vigerske, S.; Weninger, D.; and Xu, L. 2024. The SCIP Optimization Suite 9.0.

Béjar, R.; Cabiscol, A.; Manyà, F.; and Planes, J. 2009. Generating Hard Instances for MaxSAT. In *2009 39th International Symposium on Multiple-Valued Logic*, 191–195. ISSN: 2378-2226.

Cappart, Q.; Chételat, D.; Khalil, E.; Lodi, A.; Morris, C.; and Veličković, P. 2022. Combinatorial optimization and reasoning with graph neural networks. ArXiv:2102.09544 [cs, math, stat].

Chmiela, A.; Khalil, E. B.; Gleixner, A.; Lodi, A.; and Pokutta, S. 2021. Learning to Schedule Heuristics in Branch-and-Bound. ArXiv:2103.10294 [cs, math].

Dokeroglu, T.; Kucukyilmaz, T.; and Talbi, E.-G. 2024. Hyper-heuristics: A survey and taxonomy. *Computers & Industrial Engineering*, 187: 109815.

Duflo, G.; Kieffer, E.; Brust, M. R.; Danoy, G.; and Bouvry, P. 2019. A GP Hyper-Heuristic Approach for Generating TSP Heuristics. In *2019 IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW)*, 521–529.

Erdos, P. 1961. On the evolution of random graphs. *Bulletin of the Institute of International Statistics*, 38: 343–347.

Gleixner, A.; Hendel, G.; Gamrath, G.; Achterberg, T.; Bastubbe, M.; Berthold, T.; Christophel, P.; Jarck, K.; Koch, T.; Linderoth, J.; Lübbecke, M.; Mittelmann, H. D.; Ozyurt, D.; Ralphs, T. K.; Salvagnin, D.; and Shinano, Y. 2021. MIPLIB 2017: data-driven compilation of the 6th mixed-integer programming library. *Mathematical Programming Computation*, 13(3): 443–490.

He, H.; Daume III, H.; and Eisner, J. M. 2014. Learning to Search in Branch and Bound Algorithms. In *Advances in Neural Information Processing Systems*, volume 27. Curran Associates, Inc.

Hewitt, M.; Nemhauser, G. L.; and Savelsbergh, M. W. P. 2009. Combining Exact and Heuristic Approaches for the Capacitated Fixed-Charge Network Flow Problem. *INFORMS Journal on Computing*. Publisher: INFORMS.

Ibaraki, T. 1976. Theoretical comparisons of search strategies in branch-and-bound algorithms. *International Journal of Computer & Information Sciences*, 5(4): 315–344.

Kostikas, K.; and Fragakis, C. 2004. Genetic Programming Applied to Mixed Integer Programming. In Keijzer, M.; O’Reilly, U.-M.; Lucas, S.; Costa, E.; and Soule, T., eds., *Genetic Programming*, Lecture Notes in Computer Science, 113–124. Berlin, Heidelberg: Springer. ISBN 978-3-540-24650-3.

Koza, J. R. 1989. Hierarchical genetic algorithms operating on populations of computer programs. In *Proceedings of the Eleventh International Joint Conference on Artificial Intelligence IJCAI-89*. Proceedings of the Eleventh International Joint Conference on Artificial Intelligence IJCAI-89.

Labassi, A. G.; Chételat, D.; and Lodi, A. 2022. Learning to Compare Nodes in Branch and Bound with Graph Neural Networks. *NeurIPS*.

Lawler, E. L.; and Wood, D. E. 1966. Branch-and-Bound Methods: A Survey. *Operations Research*. Publisher: INFORMS.

Linderoth, J. T.; and Savelsbergh, M. W. P. 1999. A Computational Study of Search Strategies for Mixed Integer Programming. *INFORMS Journal on Computing*, 11(2): 173–187.

Lodi, A.; and Zarpellon, G. 2017. On learning and branching: a survey. *TOP*, 25(2): 207–236.

Luke, S.; and Panait, L. 2002. Fighting Bloat with Nonparametric Parsimony Pressure. In Guervós, J. J. M.; Adamidis,

- P.; Beyer, H.-G.; Schwefel, H.-P.; and Fernández-Villacañas, J.-L., eds., *Parallel Problem Solving from Nature — PPSN VII*, 411–421. Berlin, Heidelberg: Springer. ISBN 978-3-540-45712-1.
- Maher, S.; Miltenberger, M.; Pedroso, J. P.; Rehfeldt, D.; Schwarz, R.; and Serrano, F. 2016. PySCIPOpt: Mathematical Programming in Python with the SCIP Optimization Suite. In Greuel, G.-M.; Koch, T.; Paule, P.; and Sommese, A., eds., *Mathematical Software – ICMS 2016*, 301–307. Cham: Springer International Publishing. ISBN 978-3-319-42432-3.
- Morikawa, K.; Nagasawa, K.; and Takahashi, K. 2019. Job Shop Scheduling by Branch and Bound Using Genetic Programming. *Procedia Manufacturing*, 39: 1112–1118.
- Morrison, D. R.; Jacobson, S. H.; Sauppe, J. J.; and Sewell, E. C. 2016. Branch-and-bound algorithms: A survey of recent advances in searching, branching, and pruning. *Discrete Optimization*, 19: 79–102.
- Parsonson, C. W. F.; Laterre, A.; and Barrett, T. D. 2023. Reinforcement Learning for Branch-and-Bound Optimisation Using Retrospective Trajectories. *Proceedings of the AAAI Conference on Artificial Intelligence*, 37(4): 4061–4069. Number: 4.
- Rais, A.; and Viana, A. 2011. Operations Research in Healthcare: a survey. *International Transactions in Operational Research*, 18(1): 1–31. eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1475-3995.2010.00767.x>.
- Scavuzzo, L.; Aardal, K.; Lodi, A.; and Yorke-Smith, N. 2024. Machine Learning Augmented Branch and Bound for Mixed Integer Linear Programming. ArXiv:2402.05501 [cs, math].
- Song, J.; Lanka, R.; Zhao, A.; Yue, Y.; and Ono, M. 2018. Learning to Search via Retrospective Imitation.
- Talbi, E.-G. 2009. *Metaheuristics: From Design to Implementation*. John Wiley & Sons. ISBN 978-0-470-49690-9. Google-Books-ID: SIsa6zi5XV8C.
- Tiwari, A.; Hoyos, P. N.; Hutabarat, W.; Turner, C.; Ince, N.; Gan, X.-P.; and Prajapat, N. 2015. Survey on the use of computational optimisation in UK engineering companies. *CIRP Journal of Manufacturing Science and Technology*, 9: 57–68.
- Tomazella, C. P.; and Nagano, M. S. 2020. A comprehensive review of Branch-and-Bound algorithms: Guidelines and directions for further research on the flowshop scheduling problem. *Expert Systems with Applications*, 158: 113556.
- Yang, X.; Li, X.; Ning, B.; and Tang, T. 2016. A Survey on Energy-Efficient Train Operation for Urban Rail Transit. *IEEE Transactions on Intelligent Transportation Systems*, 17(1): 2–13. Conference Name: IEEE Transactions on Intelligent Transportation Systems.
- Yilmaz, K.; and Yorke-Smith, N. 2021. A Study of Learning Search Approximation in Mixed Integer Branch and Bound: Node Selection in SCIP. *AI*, 2(2): 150–178. ArXiv:2007.03948 [cs, math].
- Zarpellon, G.; Jo, J.; Lodi, A.; and Bengio, Y. 2021. Parameterizing Branch-and-Bound Search Trees to Learn Branching Policies. *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(5): 3931–3939. Number: 5.