

Revisiting local branching with a machine learning lens

Defeng Liu · Matteo Fischetti · Andrea Lodi

Received: date / Accepted: date

Abstract Finding high-quality solutions to mixed-integer linear programming problems (MILPs) is of great importance for many practical applications. In this respect, the refinement heuristic *local branching* (LB) has been proposed to produce improving solutions and has been highly influential for the development of local search methods in MILP. The algorithm iteratively explores a sequence of solution neighborhoods defined by the so-called *local branching constraint*, namely, a linear inequality limiting the distance from a reference solution. For a LB algorithm, the choice of the neighborhood size is critical to performance. In this work, we study the relation between the size of the search neighborhood and the behavior of the underlying LB algorithm, and we devise a learning based framework for predicting the best size for the specific instance to be solved. Furthermore, we have also investigated the relation between the time limit for exploring the LB neighborhood and the actual performance of LB scheme, and devised a strategy for adapting the time limit. We computationally show that the neighborhood size and time limit can indeed be learned, leading to improved performances and that the overall algorithm generalizes well both with respect to the instance size and, remarkably, across instances.

Keywords MILP · Primal Heuristic · Local Branching · Machine Learning

This work is an extension of [37]

Defeng Liu

CERC, Polytechnique Montréal, Canada, E-mail: defeng.liu@polymtl.ca

Matteo Fischetti

Department of Information Engineering, University of Padua, Italy, E-mail: matteo.fischetti@unipd.it

Andrea Lodi

CERC, Polytechnique Montréal, Canada, and Jacobs Technion-Cornell Institute, Cornell Tech and Technion - IIT, USA, E-mail: andrea.lodi@cornell.edu

Mathematics Subject Classification (2020) 90C11 · 90C59 · 90-05 · 68T07

1 Introduction

Mixed-integer linear programming (MILP) is a main paradigm for modeling complex combinatorial problems. The exact solution of a MILP model is generally attempted by a branch-and-bound (or branch-and-cut) [35] framework. Although state-of-the-art MILP solvers experienced a dramatic performance improvement over the past decades, due to the NP-hardness nature of the problem, the computation load of finding a provable optimal solution for the resulting models can be heavy. In many practical cases, feasible solutions are often required within a very restricted time frame. Hence, one is interested in finding solutions of good quality at the early stage of the computation. In fact, it is also appealing to discover early incumbent solutions in the exact enumerate scheme, which improves the primal bound and reduces the size of the branch-and-bound tree by pruning more nodes [9].

In this respect, the concept of heuristic is well rooted as a principle underlying the search of high-quality solutions. In the literature, a variety of heuristic methods have proven to be remarkably effective, e.g., *local branching* [17], *feasibility pump* [16], *RINS* [13], *RENS* [10], *proximity search* [19], *large neighborhood search* [22], etc. More details of these developments are reviewed in [8, 18, 22]. In this paper, we focus on local branching, a *refinement heuristic* that iteratively produces improved solutions by exploring suitably predefined solution neighborhoods.

Local branching (LB) was one of the first methods using a generic MILP solver as a black-box tool inside a heuristic framework. Given an initial feasible solution, the method first defines a solution neighborhood through the so-called *local branching constraint*, then explores the resulting subproblem by calling a black-box MILP solver. For a LB algorithm, the choice of neighborhood size is crucial to performance. In the original LB algorithm [17], the size of neighborhood is mostly initialized by a small value, then adjusted in the subsequent iterations. Although these conservative settings have the advantage of yielding a series of easy-to-solve subproblems, each leading to a small progress of the objective, there is still a lot of space for improvement. As discussed in [19], a significantly better performance can be potentially achieved with an ad-hoc tuning of the size of the neighborhood. Our observation also shows that the “best” size is strongly dependent on the particular MILP instance. To illustrate this, the performance of different LB neighborhood size settings for two MILP instances are compared in Figure 1. In principle, it is desirable to have neighborhoods to be relatively small to allow for an efficient exploration, but still large enough to be effective for finding improved solutions. Nonetheless, it is reasonable to believe that the size of an ideal neighborhood is correlated with the characteristics of the particular problem instance.




Fig. 1 Evaluation of the size of LB neighborhood on a set covering instance (sc-0) and a maximum independent set instance (mis-1). The neighborhood size k is computed as $k = r \times N$, where N is the number of binary variables, and $r \in [0, 1]$. A time limit is imposed for each neighborhood exploration.

Furthermore, it is worth noting that, in many applications, instances of the same problem are solved repeatedly. Problems of real-world applications have a rich structure. While more and more datasets are collected, patterns and regularities appear. Therefore, problem-specific and task-specific knowledge can be learned from data and applied to the corresponding optimization scenario. This motivates a broader paradigm of learning to guide the neighborhood search in refinement heuristics.

In this paper, we investigate a learning framework for sizing the search neighborhood of local branching. In particular, given a MILP instance, we exploit patterns in both the structure of the problem and the information collected from the solving process to predict the size of the LB neighborhood and the time limit for exploring the neighborhood, with the aim of maximizing the performance of the underlying LB algorithm. We computationally show that the neighborhood size and the time limit can indeed be learned, leading to improved performances, and that the overall algorithm generalizes well both with respect to the instance size and, more surprisingly, across instances.

We note that a shorter conference version of this paper appeared in [37]. Our initial conference paper did not consider the effect of the time limit for solving the LB neighborhood on the performance of overall algorithm. This extended paper addresses this issue and provides additional analysis on applying our refined LB algorithms as a primal heuristic within the MILP solver.

The paper is organized as follows. In Section 2, we give a review of the related works in the literature. In Section 3, we introduce the basic local branching scheme and some relevant concepts. In Section 4, we present our methodology for learning to search in the local branching scheme. In Section 5 describes the setup of our experiments and reports the results to validate our approach. In Section 6, we apply local branching as a primal heuristic and provide two possible implementations of how our local branching scheme interacts with the MILP solver. Section 7 concludes the paper and discusses future research.

2 Related work

Recently, the progress in machine learning (ML) has stimulated increasing research interest in learning algorithms for solving MILP problems. These works can be broadly divided into two categories: *learning decision strategies within MILP solvers*, and *learning primal heuristics*.

The first approach investigates the use of ML to learn to make decisions inside a MILP solver, which is typically built upon a general branch-and-bound framework. The learned policies can be either cheap approximations of existing expensive methods, or more sophisticated strategies that are to be discovered. Related works include: learning to select branching variables [3, 21, 32], learning to select branching nodes [26], learning to select cutting planes [47], and learning to optimize the usage of primal heuristics [11, 27, 33].

The *learning primal heuristics* approach is to learn algorithms to produce primal solutions for MILPs. Previous works in this area typically use ML methods to develop *large neighborhood search* (LNS) heuristics. Within an LNS scheme, ML models are trained to predict “promising” solution neighborhoods that are expected to contain high-quality solutions. In [14], the authors trained neural networks to directly predict solution values of binary variables, and then applied the LB heuristic to explore the solution neighborhoods around the predictions. The work [40] also uses neural networks to predict partial solutions. The subproblems defined by fixing the predicted partial solutions are solved by a MILP solver. The work [44] proposes a LNS heuristic based on a “learn to destroy” strategy, which frees part of the current solution. The variables to be freed are selected by trained neural networks using imitation learning. Note that their methods rely on parallel computation, which makes the outcome of the framework within a non-parallel environment less clear. In [43], the authors proposed a decomposition-based LNS heuristic. They used imitation learning and reinforcement learning to decompose the set of integer variables into subsets of fixed size. Each subset defines a subproblem. The number of subsets is fixed as a hyperparameter.

Note that the learning-based LNS methods listed above directly operate on the integer variables, i.e., the predictions of ML models are at a variable-wise level, which still encounters the intrinsic combinatorial difficulty of the problem and limits their generalization performances on generic MILPs. Moreover, the learning of these heuristics is mostly based on the extraction of static features of the problem, the dynamic statistics of the heuristic behavior of the solver being barely explored. In our work, we aim at avoiding directly making predictions on variables. Instead, we propose to guide the (local) search by learning how to control the neighborhood size at an instance-wise level. To identify promising solution neighborhoods, our method exploits not only the static features of the problem, but also the dynamic features collected during the solution process as a sequential approach.

In the literature, there has also been an effort to learn algorithms for solving specific combinatorial optimization problems [4, 29, 31, 34, 38, 41]. For a detailed overview of “learn to optimize”, see [5].

3 Preliminaries

3.1 Local branching

We consider a MILP problem with 0–1 variables of the form

$$(P) \quad \min \quad \mathbf{c}^T \mathbf{x} \quad (1)$$

$$\text{s.t.} \quad \mathbf{A}\mathbf{x} \leq \mathbf{b}, \quad (2)$$

$$x_j \in \{0, 1\}, \quad \forall j \in \mathcal{B}, \quad (3)$$

$$x_j \in \mathbb{Z}^+, \quad \forall j \in \mathcal{G}, \quad x_j \geq 0, \quad \forall j \in \mathcal{C}, \quad (4)$$

where the index set of decision variables $\mathcal{N} := \{1, \dots, n\}$ is partitioned into $\mathcal{B}, \mathcal{G}, \mathcal{C}$, which are the index sets of binary, general integer and continuous variables, respectively.

Note that we assume the existence of binary variables, as one of the basic building blocks of our method—namely, the local branching heuristic—is based on this assumption. However, this limitation can be relaxed and the local branching heuristic can be extended to deal with general integer variables, as proposed in [7].

Let $\bar{\mathbf{x}}$ be a feasible *incumbent* solution for (P) , and let $\bar{\mathcal{S}} = \{j \in \mathcal{B} : \bar{x}_j = 1\}$ denote the binary support of $\bar{\mathbf{x}}$. For a given positive integer parameter k , we define the neighborhood $N(\bar{\mathbf{x}}, k)$ as the set of the feasible solutions of (P) satisfying the *local branching constraint*

$$\Delta(\mathbf{x}, \bar{\mathbf{x}}) = \sum_{j \in \mathcal{B} \setminus \bar{\mathcal{S}}} x_j + \sum_{j \in \bar{\mathcal{S}}} (1 - x_j) \leq k. \quad (5)$$

In the relevant case in which solutions with a small binary support are considered (for example, in the famous traveling salesman problem only n or the $O(n^2)$ variables take value 1), the asymmetric form of local branching constraint is suited, namely

$$\Delta(\mathbf{x}, \bar{\mathbf{x}}) = \sum_{j \in \bar{\mathcal{S}}} (1 - x_j) \leq k. \quad (6)$$

The local branching constraint can be used in an exact branching scheme for (P) . Given the incumbent solution $\bar{\mathbf{x}}$, the solution space with the current branching node can be partitioned by creating two child nodes as follows:

$$\text{Left: } \Delta(\mathbf{x}, \bar{\mathbf{x}}) \leq k, \quad \text{Right: } \Delta(\mathbf{x}, \bar{\mathbf{x}}) \geq k + 1.$$

3.2 The neighborhood size optimization problem

For a *neighborhood size* parameter $k \in \mathbb{Z}^+$, the LB algorithm obtained from choosing k can be denoted as \mathcal{A}_k . Given a MILP instance \mathbf{p} , with its incumbent

solution $\bar{\mathbf{x}}$, the neighborhood size optimization problem over k for one iteration of \mathcal{A}_k is defined as

$$\min \quad \mathcal{C}(\mathbf{p}, \bar{\mathbf{x}}; \mathcal{A}_k) \quad (7)$$

$$\text{s.t.} \quad k \in \mathbb{Z}^+, \quad (8)$$

where $\mathcal{C}(\mathbf{p}, \bar{\mathbf{x}}; \mathcal{A}_k)$ measures the ‘‘cost’’ of \mathcal{A}_k on instance $(\mathbf{p}, \bar{\mathbf{x}})$ as a trade-off between execution speed and solution quality.

In practice, a run of the LB algorithm consists of a sequence of LB iterations. To maximize the performance of the LB algorithm, a series of the above optimization problems need to be solved. Since the cost function \mathcal{C} is unknown, those problems cannot be solved analytically. In general, the common strategy is to evaluate some trials of k and select the most performing one. This is often done by using black-box optimization methods [1]. As the evaluation of each setting involves a run of \mathcal{A}_k and the best k is instance-specific, those methods are not computationally efficient enough for online use. That is why the original LB algorithm initializes k with a fixed small value and adapts it conservatively by a deterministic strategy.

Currently, learning from experiments and transferring the learned knowledge from solved instances to new instances is of increasing interest and somehow accessible. In the next section, we will introduce a new strategy for selecting the neighborhood size k by using data-driven methods.

4 Learning methods

Next, we present our framework for learning the neighborhood size in the LB scheme. The original LB algorithm chooses a conservative value for k as default, with the aim of generating a easy-to-solve subproblem for general MILPs. However, as discussed in Section 1, our observation shows that the ‘‘best’’ k is dependent on the particular MILP instance. Hence, in order to optimize the performance of the LB heuristic, we aim at devising new strategies to learn how to tailor the neighborhood size for a specific instance. In particular, we investigate the dependencies between the state of the problem (defined by a set of both static and dynamic features collected during the LB procedure, e.g., context of the problem, incumbent solution, solving status, computation cost, etc.) and the size of the LB neighborhood.

Our framework consists of a two-phase strategy. In the first phase, we define a regression task to learn the neighborhood size for the first LB iteration. Within our method, this size is predicted by a pretrained regression model. For the second phase, we leverage reinforcement learning (RL) [45] and train a policy to dynamically adapt the neighborhood size at the subsequent LB iterations. The exploration of each LB neighborhood is the same as in the original LB framework, and a generic black-box MILP solver is used to update the incumbent solution. The overall scheme is exact in nature although turning it into a specialized heuristic framework is trivial (and generally preferable).

4.1 Scaled regression for local branching

For intermediate LB iterations, the statistics of previous iterations (e.g., value of the previous k , solving statistics, etc.) are available. One can then take advantage of this information and exploit the learning methods based on dynamic programming (e.g., reinforcement learning). Section 4.2 will address this case. However, for the first LB iteration, there is no historical information available as input. In this section, we will show how to define a regression task to learn the first k from the context of the problem and the incumbent solution.

Let \mathcal{S} denote the set of available features of the MILPs before the first LB iteration. We aim to train a regression model $f : \mathcal{S} \rightarrow \mathbb{R}$ that maps the features of a MILP instance \mathbf{s} to k_0^* , the label of best k_0 . However, the label k_0^* is unknown, and we do not have any existing method to compute the exact k_0^* . To generate labels, we first define a metric for assessing the performance of a LB algorithm \mathcal{A}_k , and then use black-box optimization methods to produce approximations of k_0^* as labels.

4.1.1 Approximation of the best k_0

To define a cost metric for \mathcal{A}_k , we consider two factors. The first factor is the computational effort (e.g., CPU time) to solve the sub-MILP defined by the LB neighborhood, while the second factor is the solution quality (e.g., the objective value of the best solution). To quantify the trade-off of speed and quality, the cost metric can be defined as

$$c^{k_0} = \alpha * t_{scaled}^{k_0} + (1 - \alpha) * o_{scaled}^{k_0}, \quad (9)$$

where $t_{scaled}^{k_0} \in [0, 1]$ is the scaled computing time for solving the sub-MILP, $o_{scaled}^{k_0} \in [0, 1]$ is the scaled objective of sub-MILP, and α is a constant.

Given c^{k_0} , the label k_0^* can be defined as

$$k_0^* = \underset{k_0 \in \mathbb{Z}^+}{\operatorname{argmin}} c^{k_0}, \quad (10)$$

and is usually evaluated through black-box optimization methods: Given the time limit and a collection of training instances of interest, one evaluates the LB algorithm introduced in Section 3 with different values of k_0 , and k_0^* is estimated by choosing the value with the best performance assessed by (9), which is typically the largest k_0 such that the resulting sub-MILP can still be solved to optimality within the time limit. Since the evaluation process for each instance is quite expensive, we propose to approximate it through regression.

4.1.2 Regression for learning k_0

With a collected dataset $\mathcal{D} = (\mathbf{s}_i, k_{0i}^*)_{i=1}^N$ with N instances, a regression task can be analyzed to learn a mapping from the state of the problem to the estimated k_0^* . The regression model $f_\theta(\mathbf{s})$ can be obtained by solving

$$\theta^* = \operatorname{argmin}_{\theta \in \Theta} \frac{1}{N} \sum_{i=1}^N \mathcal{L}(f_\theta(\mathbf{s}_i), k_{0i}^*), \quad (11)$$

where $\mathcal{L}(f_\theta(\mathbf{s}_i), k_{0i}^*)$ defines the loss function, a typical choice for regression task being the mean squared error.

4.1.3 The scaled regression task

Let \mathbf{x}' be the optimal linear programming (LP) fractional solution without local branching constraint, and let k' be the value of the left-hand side of the local branching constraint evaluated using \mathbf{x}' . Specifically, k' is computed by

$$k' = \Delta(\mathbf{x}', \bar{\mathbf{x}}). \quad (12)$$

As discussed in [19], any $k \geq k'$ is likely to be useless as the LP solution after adding the LB constraint would be unchanged. Hence, k' provides an upper bound for k .

We can therefore parametrize k as

$$k = \phi k', \quad (13)$$

where $\phi \in (0, 1)$. Now, we define the regression task over a scaled space $\phi \in (0, 1)$ instead of directly over k .

Given k_0^* and k_0' , the label is easily computed by

$$\phi_0^* = \frac{k_0^*}{k_0'}. \quad (14)$$

The regression problem reduces to

$$\theta^* = \operatorname{argmin}_{\theta \in \Theta} \frac{1}{N} \sum_{i=1}^N \mathcal{L}(f_\theta(\mathbf{s}_i), \phi_{0i}^*), \quad (15)$$

where $\mathcal{L}(f_\theta(\mathbf{s}_i), \phi_{0i}^*)$ defines the loss function.

MILP representation We represent the state \mathbf{s} as a bipartite graph $(\mathbf{C}, \mathbf{E}, \mathbf{V})$ [21]. Given a MILP instance, let n be the number of variables with d features for each variable, m be the number of constraints with q features for each constraint. The variables of the MILP, with $\mathbf{V} \in \mathbb{R}^{n \times d}$ being their feature matrix, are represented on one side of the graph. On the other side are nodes corresponding to the constraints with $\mathbf{C} \in \mathbb{R}^{m \times q}$ being their feature matrix. A constraint node i and a variable node j are connected by an edge (i, j) if variable i appears in constraint j in the MILP model. Finally, $\mathbf{E} \in \mathbb{R}^{m \times n \times e}$ denotes the tensor of edge features, with e being the number of features for each edge.

Algorithm 1 LB with scaled regression

Input: instance dataset $\mathcal{P} = \{\mathbf{p}_i\}_{i=1}^M$
for instance $\mathbf{p}_i \in \mathcal{P}$ **do**
 0. initialize the state \mathbf{s} with an initial solution $\bar{\mathbf{x}}$
 1. solve the LP relaxation and get solution \mathbf{x}'
 2. compute $k' = \Delta(\mathbf{x}', \bar{\mathbf{x}})$
 3. predict $\phi_0 = f_\theta(\mathbf{s})$ by the regression model
 4. compute $k_0 = \phi_0 k'$
 5. apply k_0 to execute the first LB iteration
 6. update the incumbent $\bar{\mathbf{x}}$ and continue LB algorithm with its default setting
 repeat
 | execute the next LB iteration
 until termination condition is reached;
end

Regression model Given that states are represented as graphs, with arbitrary size and topology, we propose to use *graph neural networks* (GNNs) [24, 25] to parameterize the regression model. Indeed, GNNs are size-and-order invariant to input data, i.e., they can process graphs of arbitrary size, and the ordering of the input elements is irrelevant. Another appealing property of GNNs is that they exploit the sparsity of the graph, which makes GNNs an efficient model for embedding MILP problems that are typically very sparse [21].

Our GNN architecture consists of three modules: the input module, the convolution module, and the output module. In the input layer, the state \mathbf{s} is fed into the GNN model. The input module embeds the features of the state \mathbf{s} . The convolution module propagates the embedded features with graph convolution layers. In particular, our graph convolution layer applies the message passing operator, defined as

$$\mathbf{v}_i^{(h)} = f_\theta^{(h)} \left(\mathbf{v}_i^{(h-1)}, \sum_{j \in \mathcal{N}(i)} g_\phi^{(h)} \left(\mathbf{v}_i^{(h-1)}, \mathbf{v}_j^{(h-1)}, \mathbf{e}_{j,i} \right) \right), \quad (16)$$

where $\mathbf{v}_i^{(h-1)} \in \mathbb{R}^d$ denotes the feature vector of node i from layer $(h-1)$, $\mathbf{e}_{j,i} \in \mathbb{R}^m$ denotes the feature vector of edge (j, i) from node j to node i of layer $(h-1)$, and $\mathbf{f}_\theta^{(h)}$ and $\mathbf{g}_\phi^{(h)}$ denote the embedding functions in layer h .

For a bipartite graph, a convolution layer is decomposed into two half-layers: one half-layer propagates messages from variable nodes to constraint nodes through edges, and the other one propagates messages from constraint nodes to variable nodes. The output module embeds the features extracted from the convolution module and then applies a pooling layer, which maps the graph representation into a single neuron. The output of this neuron is the prediction of ϕ_0 .

4.1.4 LB with scaled regression

Our refined LB heuristic, *LB with scaled regression*, is obtained when k_0 is predicted by the regression model. The pseudocode of the algorithm is outlined in Algorithm 1.

4.2 Reinforced neighborhood search

In this section, we leverage reinforcement learning (RL) to adapt the neighborhood size iteratively. We first formulate the problem as a Markov Decision Process (MDP) [30]. Then, we propose to use policy gradient methods to train a policy model.

4.2.1 Markov Decision Process

Given a MILP instance with an initial feasible solution, the procedure can be formulated as a MDP, wherein at each step, a neighborhood size is selected by a policy model and applied to run a LB iteration. In principle, the state space \mathcal{S} is the set of all the features of the MILP model and its solving statistics, which is combinatorial and arbitrarily large. To design an efficient RL framework for this problem, we choose a compact set of features from the solving statistics to construct the state. These features characterize the progress of the optimization process and are instance-independent, allowing broader generalization across instances.

For the action space $\text{State}(\mathcal{A})$, instead of directly selecting a new k , we choose to adapt the value of k of the last LB iteration. The set of possible actions consists of four options

$$\{+k_{step}, 0, -k_{step}, reset\}, \quad (17)$$

where “ $+k_{step}$ ” means increasing k by $k = k + k_{step} * k$, “ $-k_{step}$ ” means decreasing k by $k = k - k_{step} * k$, “0” denotes keeping k without any change, and “*reset*” means resetting k to a default value. The policy π_k maps a state to one of the four actions. The step size k_{step} is a hyperparameter of the algorithm.

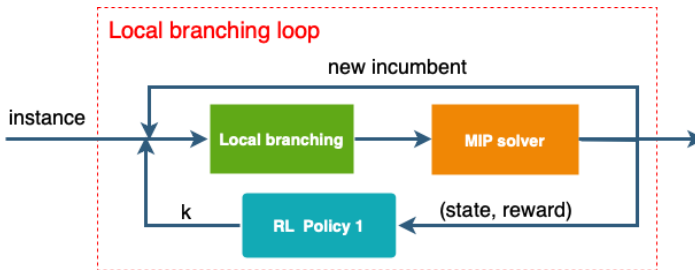


Fig. 2 RL framework for adapting k

Algorithm 2 Reinforced neighborhood search for adapting k

```

Input: instance dataset  $\mathcal{P} = \{p_i\}_{i=1}^M$ 
for instance  $p_i \in \mathcal{P}$  do
  0. initialize the state  $s$  with an initial solution  $\bar{x}$ 
  1. compute  $k_0$  by the procedure in Algorithm 1 or set  $k_0$  by a default value
  2. apply  $k_0$  to execute the first LB iteration
  3. collect the new state  $s$  and the incumbent  $\bar{x}$ 
  repeat
    update  $k$  by policy  $\pi_k(s)$ 
    apply  $k$  and execute the next LB iteration
    collect the new state  $s$  with the incumbent  $\bar{x}$ 
  until termination condition is reached;
end

```

The compact description of states and actions offers several advantages. First of all, it simplifies the MDP formulation and makes the learning task easier. In addition, it allows for the use of simpler function approximators, which is critical for speeding up the learning process. In Section 5, we will show—by training a simple linear policy model using the off-the-shelf policy gradient method—that the resulting policy can significantly improve the performance of the LB algorithm.

By applying the updated k , the next LB iteration is executed with time limit t_{limit} . Then, the solving sub-MILP statistics are collected to create the next state. In principle, the reward r_k is formulated according to the outcome of the last LB iteration, e.g., the computing time and the quality of the incumbent solution.

To maximize the objective improvement and minimize the solution time of the LB algorithm, we define the combinatorial reward as

$$r_k = o_{imp} * (t_{max} - t_{elaps}), \quad (18)$$

where o_{imp} denotes the objective improvement obtained from the last LB iteration, t_{max} is the global time limit of the LB algorithm, and t_{elaps} is the cumulated running time.

The definition above is just one possibility to build a MDP for the LB heuristic. Actually, defining a compact MDP formulation is critical for constructing efficient RL algorithms for this problem.

4.2.2 Learning strategy

For training the policy model, we use the *reinforce* policy gradient method [46], which allows a policy to be learned without any estimate of the value functions.

The refined LB heuristic, *reinforced neighborhood search for adapting k* is obtained when the neighborhood size k , is dynamically adapted by the RL policy. The pseudocode of the algorithm is outlined in Algorithm 2.

4.3 Further improvement by adapting LB node time limit

In the previous section, we suppose that the time limit for solving each local branching neighborhood is given. Indeed, the setting of time limit for each “LB node” decides how much computational effort is spent on each LB node, and therefore customizes how the overall time limit is split. E.g., a larger “LB node” time limit allows more computation for solving that LB node, which potentially leads to a larger k . However, it should be noted that the relation between the value of “LB node” time limit and the actual performance of LB scheme is not known. Therefore, we seek to address this issue by leveraging reinforcement learning again to learn policies for tailoring the time limit for each LB node. This new “RL-t” loop is built on top of the inner “RL-k” loop for adapting the LB neighborhood size. The scheme is shown in Figure 3.

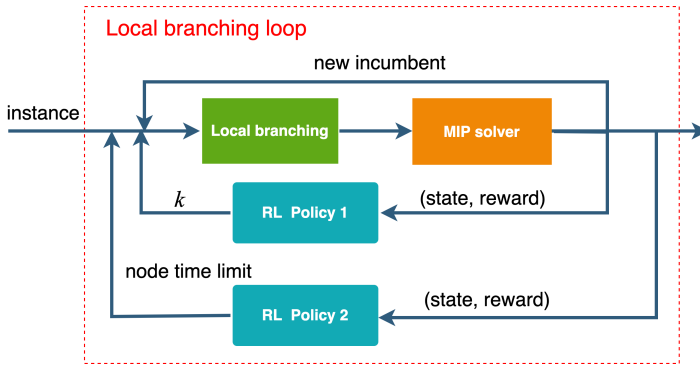


Fig. 3 RL framework for adapting the time limit for solving the LB subproblem.

To design the action space for adapting the time limit of each LB node, we apply a similar setting as used for adapting k . The set of possible actions consists of four options

$$\{ *t_{step}, 0, /t_{step}, reset \}, \quad (19)$$

where “ $*t_{step}$ ” means increasing t by $t = t * t_{step}$, “ $/t_{step}$ ” means decreasing t by $t = t / t_{step}$, “0” denotes keeping t without any change, and “reset” means resetting t to a default value. The policy π_t maps a state to one of the four actions.

For the reward design of the “RL-t” loop, the reward also takes into account both the cost of computation spent on solving the LB subproblem and the solution quality. Specifically, the reward r_t for the outer RL-t loop consists of two components. The first component r_1 maintains the same reward r_k as used in the “RL-k” loop, which attempts to maximize the improvement of the objective by the last LB iteration while minimizing the computing time. In addition, the LB environment explicitly returns a reward signal as a second component to penalize those LB iterations where the subproblem is too large

Algorithm 3 Hybrid reinforced neighborhood search for adapting k and t

Input: instance dataset $\mathcal{P} = \{p_i\}_{i=1}^M$
for instance $p_i \in \mathcal{P}$ **do**
 0. initialize the state s with an initial solution \bar{x}
 1. compute k_0 by the procedure in Algorithm 1 or set k_0 by a default value
 2. apply k_0 to execute the first LB iteration
 3. collect the new state s and the incumbent \bar{x}
 repeat
 update k by policy $\pi_k(s)$
 update t by policy $\pi_t(s)$
 apply k, t and execute the next LB iteration
 collect the new state s with the incumbent \bar{x}
 until *termination condition is reached*;
end

to solve and there is no objective improvement within the time limit. The penalty reward r_p is a binary signal defined as

$$r_p = \begin{cases} 1 & \text{if the LB subproblem is not solved and no improving} \\ & \text{solution is returned,} \\ 0 & \text{otherwise.} \end{cases}$$

Formally, the reward r_t is a combination of two components, defined as

$$r_t = \beta_1 r_1 + \beta_2 r_2$$

where $r_1 = r_k$, $r_2 = r_p$, $\beta_1 > 0$, $\beta_2 > 0$.

For training the policy $\pi_t(s)$, we use the same RL method as used for training the policy $\pi_k(s)$. Note that we fix the pretrained $\pi_k(s)$ policy while tuning $\pi_t(s)$ to make the learning process more stable.

A new LB algorithm, *hybrid reinforced neighborhood search*, is obtained when the neighborhood size k and the time limit t for each LB node are dynamically adapted by the RL policies. The pseudocode of the algorithm is outlined in Algorithm 3.

5 Experiments

In this section, we present the details of our experimental results over five MILP benchmarks. We compare different settings of our approach against the original LB algorithm, using SCIP [20] as the underlying MILP solver.

5.1 Data collection

5.1.1 MILP instances

We evaluate on five MILP benchmarks: set covering (SC) [2], maximum independent set (MIS) [6], combinatorial auction (CA) [36], generalized independent set problem (GISP) [12,28], and MIPLIB 2017 [23]. The first three benchmarks are used for both training and evaluation. For SC, we use instances with

5000 rows and 2000 columns. For MIS, we use instances on Barabási–Albert random graphs with 1000 nodes. For CA, we use instances with 4000 items and 2000 bids. In addition, to evaluate the generalization performance on larger instances, we also use a larger dataset of instances with doubled size for each benchmark, denoted by LCA, LMIS, LCA. The larger datasets are only used for evaluation.

For GISP, we use the public dataset from [11]. For MIPLIB, we select binary integer linear programming problems from MIPLIB 2017. Instances from GISP and MIPLIB are only used for evaluation.

For each instance, an initial feasible solution is required to run the LB heuristic. We use two initial incumbent solutions: (1) the first solution found by SCIP; (2) an intermediate solution found by SCIP, typically the best solution obtained by SCIP at the end of the root node computation, i.e., before branching.

5.1.2 Data Collection for regression

To collect data for the scaled regression task, one can use black-box optimization methods to produce the label ϕ_0^* . As the search space has only one dimension, we choose to use the grid search method. In particular, given a MILP instance, an initial incumbent $\bar{\mathbf{x}}$, the LP solution \mathbf{x}' , and a time limit for a LB iteration, we evaluate ϕ_0 from $(0, 1)$ with a resolution limit 0.01. For each ϕ_0 , we compute the actual neighborhood size by $k_0 = k' \phi_0$, where $k' = \Delta(\mathbf{x}', \bar{\mathbf{x}})$. Then, k_0 is applied to execute an iteration of LB. From all the evaluated ϕ_0 , the one with best performance is chosen as a label ϕ_0^* .

The state \mathbf{s} consists of context features of the MILP model and the incumbent solution. The state \mathbf{s} together with the label k^* construct a valid data point (\mathbf{s}, k^*) .

5.2 Experimental setup

5.2.1 Datasets

For each reference set of SC, MIS and CA problems, we generate a dataset of 200 instances, and split each dataset into training (70%), validation (10%), and test (20%) sets. For larger instances, we generate 40 instances of LSC, LMIS and LCA problems, separately. The GISP dataset contains 35 instances. For MIPLIB, we select 29 binary MILPs that are also evaluated by the original LB heuristic [17].

5.2.2 Model architecture and feature design

For the regression task, we apply the GNNs described in the paper with three modules. For the input module, we apply 2-layer perceptron with the *rectified linear unit* (ReLU) activation function to embed the features of nodes. For

the convolution module, we use two half-layers, one from nodes of variables to nodes of constraints, and the other one from nodes of constraints to nodes of variables. For the output module, we also apply 2-layer perceptron with the ReLU activation function. The pooling layer uses the *sigmoid* activation function. All the hidden layers have 64 neurons.

Given the input $x \in \mathbb{R}$, the Sigmoid and ReLU functions are defined as

$$\text{Sigmoid}(x) = \frac{\exp(x)}{\exp(x) + 1}, \quad (20)$$

$$\text{ReLU}(x) = \max(0, x). \quad (21)$$

For the bipartite graph representation, we reference the model used in [21]. The features in the bipartite graph are listed in Table 1. In practice, if the instance is a pure binary MILP, one can choose a more compact set of features to accelerate the training process—for example, the features describing the type and bound of the variables can be removed.

Tensor	Feature	Description
C	bias	Bias value, normalized with constraint coefficients.
E	coef	Constraint coefficient, normalized per constraint.
	coef	Objective coefficient, normalized.
	binary	Binary type binary indicator.
	integer	Integer type indicator.
V	imp_integer	Implicit integer type indicator
	continuous	Continuous type indicator.
	has_lb	Lower bound indicator.
	has_ub	Upper bound indicator.
	lb	Lower bound.
	ub	Upper bound.
	sol_val	Solution value.

Table 1 Description of the features in the bipartite graph $s = (\mathbf{C}, \mathbf{E}, \mathbf{V})$.

For the two RL policies, we apply a linear model with seven inputs and four outputs. The set of features used by the RL policies is listed in Table 2.

5.2.3 Training and evaluation

For the regression task, the model learns from the features of the MILP formulation and the incumbent solution. We use the mean squared error as the loss function. We train the regression model with two scenarios: the first one trains the model on the training set of SC, MIS and CA separately, the other

Feature	Description
optimal	Sub-MILP is solved and the incumbent is updated, indicator.
infeasible	Sub-MILP is proven infeasible, indicator.
improved	Sub-MILP is not solved but the incumbent is updated, indicator.
not_improved	Sub-MILP is not solved and no solution is found, indicator.
diverse	No improved solution is found for two iterations, indicator.
t_available	time available before the time limit of the sub-MILP is reached
obj_improve	improvement of objective.

Table 2 Description of the input features of the RL policy.

one trains a single model on a mixed dataset of the three training sets. The models trained from the two scenarios are compared on the three test sets.

For the RL task, since we only use the instance-independent features selected from solving statistics, the RL policy is only trained on the training set of SC, and evaluated on all the test sets.

To further evaluate the generalization performance with respect to the instance size and the instance type, the RL policies (trained on the SC dataset) and the regression model (trained on the SC, MIS and LCS datasets) were evaluated on GISP and MIPLIB datasets.

5.2.4 Evaluation metrics

We use two measures to compare the performance of different heuristic algorithms. The first indicator is the *primal gap*. Let $\tilde{\mathbf{x}}$ be a feasible solution, and $\tilde{\mathbf{x}}_{opt}$ be the optimal (or best known) solution. The *primal gap* (in percentage) is defined as

$$\gamma(\tilde{\mathbf{x}}) = \frac{|c^T \tilde{\mathbf{x}}_{opt} - c^T \tilde{\mathbf{x}}|}{|c^T \tilde{\mathbf{x}}_{opt}|} \times 100,$$

where we assume the denominator is nonzero.

For the second measure, we use the *primal integral* [9], which takes into account both the quality of solutions and the solving time required to find them. To define the primal integral, we first consider a *primal gap function* $p(t)$ as a function of time, defined as

$$p(t) = \begin{cases} 1, & \text{if no incumbent until time } t, \\ \bar{\gamma}(\tilde{\mathbf{x}}(t)), & \text{otherwise,} \end{cases}$$

where $\tilde{\mathbf{x}}(t)$ is the incumbent solution at time t , and $\bar{\gamma}(\cdot) \in [0, 1]$ is the *scaled primal gap* defined by

$$\bar{\gamma}(\tilde{\mathbf{x}}) = \begin{cases} 0, & \text{if } c^T \tilde{\mathbf{x}}_{opt} = c^T \tilde{\mathbf{x}} = 0, \\ 1, & \text{if } c^T \tilde{\mathbf{x}}_{opt} \cdot c^T \tilde{\mathbf{x}} < 0, \\ \frac{|c^T \tilde{\mathbf{x}}_{opt} - c^T \tilde{\mathbf{x}}|}{\max\{|c^T \tilde{\mathbf{x}}_{opt}|, |c^T \tilde{\mathbf{x}}|\}}, & \text{otherwise.} \end{cases}$$

Let $t_{\max} > 0$ be the time limit. The primal integral of a run is then defined as

$$P(t_{\max}) = \int_0^{t_{\max}} p(t) dt.$$

5.2.5 Details of experimental settings

To tune the learning rate for training the regression model, we have experimented different learning rates from 10^{-5} to 10^{-1} and have chosen a learning rate of 10^{-4} . We trained the model with a limit of 300 epochs.

For training the RL policies, we used the same method to tune the learning rate and have chosen a learning rate of 10^{-2} for π_k , and 10^{-1} for π_t . We trained the RL policies with a limit of 300 epochs.

For the hyperparameters, we have chosen $\alpha = 0.5$ as a trade-off between speed and quality for the cost metric for k_0 . We used $\beta_1 = \beta_2 = 1$ for the two components of r_t . We used $k_{step} = 0.5$ for the action design in Section 4.2 and $t_{step} = 2$ for the action design in Section 4.3. We set a time limit of 10 seconds for each LB iteration for all the compared algorithms.

Our code is written in Python 3.7 and we use Pytorch 1.60 [42], Pytorch Geometric 1.7.0 [15], PySCIPOpt 3.1.1 [39], SCIP 7.01 [20] for developing our models and solving MILPs.

5.3 Results

In order to validate our approach, we first implement LB as a heuristic search strategy for improving a certain incumbent solution using the MILP solver as a black-box.

5.3.1 Local branching search with adapting k

We perform the evaluations of our framework on the following four settings:

- *lb-sr*: Algorithm 1 with regression model trained by a homogenous dataset of *SC*, *MIS*, *CA*, separately;
- *lb-srm*: Algorithm 1 with regression model trained by a mixed dataset of *SC*, *MIS*, *CA*;
- *lb-rl*: Algorithm 2 with setting k_0 by a default value;
- *lb-srml*: combined algorithm using regression from Algorithm 1 (with regression model trained by mixed dataset of *SC*, *MIS*, *CA*) and RL from Algorithm 2.

We use the original local branching algorithm as the baseline. All the algorithms use SCIP as the underlying MILP solver and try to improve the initial incumbent with a time limit of 60s. Our code and more details of the experiment environment are publicly available¹.

¹ <https://github.com/pandat8/ML4LB>

Algo.	SC		MIS		CA	
	first	root	first	root	first	root
lb-base	57.013	5.697	5.713	4.893	9.296	3.308
lb-sr	3.643	1.751	1.537	3.499	6.358	2.025
lb-srm	4.338	1.818	1.419	3.654	6.556	2.120
lb-rl	17.304	4.513	2.616	2.616	4.450	2.042
lb-srml	2.991	1.567	1.244	2.452	3.202	1.454

Table 3 Primal integral (geometric means) for SC, MIS, CA problems.

Algo.	SC		MIS		CA	
	first	root	first	root	first	root
lb-base	1411.624	1.390	0.326	0.193	3.307	1.218
lb-sr	1.425	1.102	0.220	0.210	2.460	0.877
lb-srm	1.645	1.118	0.247	0.232	2.466	0.971
lb-rl	6.876	1.621	0.449	0.263	0.558	0.229
lb-srml	1.335	0.550	0.275	0.242	0.465	0.311

Table 4 Final primal gap (geometric means in percentage) for SC, MIS, CA problems.

The evaluation results for the basic SC, MIS, CA datasets are shown in Table 3 and Table 4. Our first observation is that the learning based algorithms of our framework significantly outperform the original LB algorithm. Both the primal integral and the final primal gap of the four LB variants are smaller than those of the baseline over most datasets, showing improved heuristic behavior. Note that, although the regression model trained by using supervised learning and the policy model trained by RL can be used independently, they benefit from being used together. As a matter of fact, the hybrid algorithm *lb-srml* combining both methods achieves a solid further improvement and outperforms the other algorithms for most cases.

We also evaluate the impact of the choice of training set for the regression model. By comparing *lb-sr* and *lb-srm*, we observe that the regression model trained on a mixed dataset of *SC*, *MIS*, *CA* exhibit a performance very close to that of the model trained on a homogeneous dataset. Indeed, the GNN networks we used embed the features of the MILP problem and its incumbent solution. In particular, one significant difference between our method and those of previous works is that, instead of training a separate model for each class of instances, our method is able to train a single model yielding competitive generalization performances across instances. This is because our models predict the neighborhood size at a instance-wise level, rather than making predictions on variables.

Broader Generalization Next, we evaluate the generalization performance with respect to the size and the type of instances. Recall that our regression model is trained on a randomly mixed dataset of SC, MIS and CA problems, and the RL policy is only trained on the training set of SC. We evaluate the trained models on larger instances (LSC, LMIS, LCA) and new MILP problems (GISP, MIPLIB). The results of evaluation on larger instances are shown in

Algo.	LSC		LMIS		LCA	
	first	root	first	root	first	root
lb-base	59.303	19.147	26.230	27.228	30.610	13.187
lb-srm	3.642	2.350	4.746	9.175	19.639	6.869
lb-rl	45.658	7.889	5.476	7.929	17.338	7.081
lb-srmrl	2.790	1.876	1.819	5.807	10.078	4.311

Table 5 Primal integral (geometric means) for LSC, LMIS, LCA problems.

Table 5 and Table 6, whereas the results on GISP and MIPLIB datasets, are shown in Table 7 and Table 8.

Algo.	LSC		LMIS		LCA	
	first	root	first	root	first	root
lb-base	8105.973	4.221	18.633	12.342	26.429	5.156
lb-srm	2.379	1.242	0.230	0.659	15.882	1.961
lb-rl	136.401	4.216	0.362	0.206	3.341	0.987
lb-srmrl	1.326	0.777	0.195	0.205	2.007	0.152

Table 6 Final primal gap (geometric means in percentage) for LSC, LMIS, LCA problems.

Algo.	GISP		MIPLIB	
	first	root	first	root
lb-base	19.833	16.692	12.318	9.003
lb-srm	13.359	10.173	11.319	6.863
lb-rl	18.949	15.429	12.676	8.318
lb-srmrl	13.739	9.865	10.172	7.151

Table 7 Primal integral (geometric means) for GISP and MIPLIB problems.

Algo.	GISP		MIPLIB	
	first	root	first	root
lb-base	22.244	17.171	44.794	11.893
lb-srm	10.307	7.704	32.061	8.730
lb-rl	20.030	14.402	42.676	11.869
lb-srmrl	11.641	5.929	20.575	9.696

Table 8 Final primal gap (geometric means in percentage) for GISP and MIPLIB problems.

Overall, all of our learning-based LB algorithms outperform the baseline, and the hybrid algorithm *lb-srmrl* achieves the best performance on most datasets. These results show that our models, trained on smaller instances,

generalize well both with respect to the instance size and, remarkably, across instances.

5.3.2 Local branching search with adapting both k and t

In this section, we analyze the results of our approach on adapting both the neighborhood size k and time limit t for each LB node. Again, local branching is implemented as an independent heuristic search scheme for improving a certain incumbent solution using a black-box MILP solver.

We compare the *lb-base* and our best algorithm *lb-srmrl* outlined in the previous section with a new algorithm:

- *lb-srmrl-adapt-t*: Combined RL algorithm using regression from Algorithm 1 (with regression model trained by mixed dataset of *SC*, *MIS*, *CA*) and Hybrid RL from Algorithm 3.

In order to validate the effectiveness of our strategy for adapting t , we conducted the experiments on the two difficult MIP datasets (GISP, MIPLIB) with a longer global time limit, 600s.

The results are shown in Table 9 and Table 10. In order to demonstrate the outcome of our LB algorithms on the solving progress of the instances over the running time, we plot the evolution of the average *primal integral* on the MIPLIB dataset in Figure 4. The plot on the left reports the results of the instances initialized by the first solution found by SCIP, whereas the plot on the right reports the results of the same instances initialized by the best solution obtained at the end of the root node of the B&B tree.

Algo.	GISP		MIPLIB	
	first	root	first	root
lb-base	116.882	101.013	59.916	35.747
lb-srm-rl	108.345	91.600	54.116	33.038
lb-srm-rl-adapt-t	81.888	68.042	50.843	28.558

Table 9 Primal integral (geometric means) for GISP and MIPLIB problems with a time limit of 600s for each instance.

Algo.	GISP		MIPLIB	
	first	root	first	root
lb-base	12.440	11.039	36.860	3.389
lb-srm-rl	12.230	11.991	20.532	5.273
lb-srm-rl-adapt-t	5.880	5.328	16.723	4.396

Table 10 Final primal gap (geometric means in percentage) for GISP and MIPLIB problems with a time limit of 600s for each instance.




Fig. 4 Evolution of the *primal integral* (geometric means) over time on binary MIPLIB dataset. Left / right: using the first / root solution to start LB.

From these results, we observe that: 1) All our learning-based LB algorithms converge faster than the LB baseline, showing improved heuristic behavior; 2) With our new strategy for adapting the LB node time limit t , the hybrid RL algorithm further improves its baseline with only a single RL policy for adapting k . The improvement becomes more and more significant as the solving time increases.

6 Local branching as a primal heuristic within a MILP solver

Local branching can also be implemented as a refinement heuristic within a generic MILP solver. In this section, we present two possible implementations of the LB algorithms outlined in the previous section, to be used as a primal heuristic. One major difference between these implementations and those of the previous section is how the global MILP search is structured. Instead of applying LB as a metaheuristic strategy, we use LB within the MILP solver to improve the incumbent at certain nodes of B&B tree. In particular, we considered the following two possible implementations:

- executing LB primal heuristic only at the root node: the MILP solver calls local branching only at the root node or at the node where the first incumbent solution is found;
- executing LB primal heuristic at multiple nodes: the MILP solver checks if there is a new incumbent or not for every f nodes in the B&B tree, and calls local branching if that is the case.

To configure the frequency f of executing the LB primal heuristic for the second implementation, we have conducted a simple hyperparameter search for f from the set of $\{1, 10, 100, 1000\}$, and the results showed that 100 performed the best. Thus, we set $f = 100$.

We evaluate the two primal heuristic implementations on the MIPLIB binary dataset and report the evolution of the average *primal integral* in Figure 5. Our observation is that the SCIP solver is improved by adding our learning

based LB into SCIP as a primal heuristic. The results also suggest that the best strategy of executing LB is only to call it at the root node of the B&B tree (or at the node where the first incumbent solution is found).




Fig. 5 Evolution of *primal integral* (geometric means) over time on binary MIPLIB dataset (1200s).

7 Discussion

In this work, we have looked at the local branching paradigm by using a machine learning lens. We have considered the neighborhood size as a main factor for quantifying high-quality LB neighborhoods. We have presented a learning based framework for predicting and adapting the neighborhood size for the LB heuristic. The framework consists of a two-phase strategy. For the first phase, a *scaled regression* model is trained to predict the size of the LB neighborhood at the first iteration through a regression task. In the second phase, we leverage reinforcement learning and devise a *reinforced neighborhood search* strategy to dynamically adapt the size at the subsequent iterations. Furthermore, we have also investigated the relation between “LB node” time limit t and the actual performance of LB scheme, and devised a strategy for adapting t . We have computationally shown that the neighborhood size and LB node time limit can indeed be learned, leading to improved performances and that the overall algorithm generalizes well both with respect to the instance size and, remarkably, across instances.

Our framework relies on the availability of an initial solution, thus it can be integrated with other refinement heuristics. For future research, it would be interesting to design more sophisticated hybrid frameworks that learn to optimize multiple refinement heuristics in a more collaborative way.

Acknowledgements We would like to thank Maxime Gasse, Didier Chételat and Elias Khalil for their helpful discussions on the project. This work was supported by Canada Excellence Research Chair in Data Science for Real-Time Decision-Making at Polytechnique Montréal. The work of Matteo Fischetti was partially supported by MiUR, Italy.

References

1. Audet, C., Hare, W.: *Derivative-Free and Blackbox Optimization*. Springer (2017)
2. Balas, E., Ho, A.: Set covering algorithms using cutting planes, heuristics, and subgradient optimization: a computational study. In: *Combinatorial Optimization*, pp. 37–60. Springer (1980)
3. Balcan, M.F., Dick, T., Sandholm, T., Vitercik, E.: Learning to branch. In: *International conference on machine learning*, pp. 344–353. PMLR (2018)
4. Bello, I., Pham, H., Le, Q.V., Norouzi, M., Bengio, S.: Neural combinatorial optimization with reinforcement learning. *arXiv preprint arXiv:1611.09940* (2016)
5. Bengio, Y., Lodi, A., Prouvost, A.: Machine learning for combinatorial optimization: a methodological tour d’horizon. *European Journal of Operational Research* **290**(2), 405–421 (2021)
6. Bergman, D., Cire, A.A., Van Hoes, W.J., Hooker, J.: *Decision diagrams for optimization*, vol. 1. Springer (2016)
7. Bertacco, L., Fischetti, M., Lodi, A.: A feasibility pump heuristic for general mixed-integer problems. *Discrete Optimization* **4**(1), 63–76 (2007)
8. Berthold, T.: Primal heuristics for mixed integer programs. Ph.D. thesis, Zuse Institute Berlin (ZIB) (2006)
9. Berthold, T.: Measuring the impact of primal heuristics. *Operations Research Letters* **41**(6), 611–614 (2013)
10. Berthold, T.: Rens. *Mathematical Programming Computation* **6**(1), 33–54 (2014)
11. Chmiela, A., Khalil, E., Gleixner, A., Lodi, A., Pokutta, S.: Learning to schedule heuristics in branch-and-bound. *arXiv preprint arXiv:2103.10294* (2021)
12. Colombi, M., Mansini, R., Savelsbergh, M.: The generalized independent set problem: Polyhedral analysis and solution approaches. *European Journal of Operational Research* **260**(1), 41–55 (2017)
13. Danna, E., Rothberg, E., Pape, C.L.: Exploring relaxation induced neighborhoods to improve mip solutions. *Mathematical Programming* **102**(1), 71–90 (2005)
14. Ding, J.Y., Zhang, C., Shen, L., Li, S., Wang, B., Xu, Y., Song, L.: Accelerating primal solution findings for mixed integer programs based on solution prediction. *Proceedings of the AAAI Conference on Artificial Intelligence* **34**(02), 1452–1459 (2020)
15. Fey, M., Lenssen, J.E.: Fast graph representation learning with pytorch geometric. *arXiv preprint arXiv:1903.02428* (2019)
16. Fischetti, M., Glover, F., Lodi, A.: The feasibility pump. *Mathematical Programming* **104**(1), 91–104 (2005)
17. Fischetti, M., Lodi, A.: Local branching. *Mathematical programming* **98**(1-3), 23–47 (2003)
18. Fischetti, M., Lodi, A.: Heuristics in mixed integer programming. *Wiley Encyclopedia of Operations Research and Management Science* (2010)
19. Fischetti, M., Monaci, M.: Proximity search for 0-1 mixed-integer convex programming. *Journal of Heuristics* **20**(6), 709–731 (2014)
20. Gamrath, G., Anderson, D., Bestuzheva, K., Chen, W.K., Eifler, L., Gasse, M., Gemander, P., Gleixner, A., Gottwald, L., Halbig, K., Hendel, G., Hojny, C., Koch, T., Le Bodic, P., Maher, S.J., Matter, F., Miltenberger, M., Mühmer, E., Müller, B., Pfetsch, M.E., Schlösser, F., Serrano, F., Shinano, Y., Tawfik, C., Vigerske, S., Wegscheider, F., Weninger, D., Witzig, J.: *The SCIP Optimization Suite 7.0*. ZIB-Report 20-10, Zuse Institute Berlin (2020). URL <http://nbn-resolving.de/urn:nbn:de:0297-zib-78023>
21. Gasse, M., Chételat, D., Ferroni, N., Charlin, L., Lodi, A.: Exact combinatorial optimization with graph convolutional neural networks. In: *Advances in Neural Information Processing Systems*, pp. 15554–15566 (2019)

22. Gendreau, M., Potvin, J.Y.: Handbook of metaheuristics, vol. 2. Springer (2010)
23. Gleixner, A., Hendel, G., Gamrath, G., Achterberg, T., Bastubbe, M., Berthold, T., Christophel, P., Jarck, K., Koch, T., Linderoth, J.: Miplib 2017: data-driven compilation of the 6th mixed-integer programming library. *Mathematical Programming Computation* pp. 1–48 (2021)
24. Gori, M., Monfardini, G., Scarselli, F.: A new model for learning in graph domains. In: Proceedings. 2005 IEEE International Joint Conference on Neural Networks, 2005., vol. 2, pp. 729–734. IEEE (2005)
25. Hamilton, W.L., Ying, R., Leskovec, J.: Representation learning on graphs: Methods and applications. arXiv preprint arXiv:1709.05584 (2017)
26. He, H., Daume III, H., Eisner, J.M.: Learning to search in branch and bound algorithms. *Advances in neural information processing systems* **27**, 3293–3301 (2014)
27. Hendel, G.: Adaptive large neighborhood search for mixed integer programming. *Mathematical Programming Computation* **14**(2), 185–221 (2022)
28. Hochbaum, D.S., Pathria, A.: Forest harvesting and minimum cuts: a new approach to handling spatial constraints. *Forest Science* **43**(4), 544–554 (1997)
29. Hottung, A., Tierney, K.: Neural large neighborhood search for the capacitated vehicle routing problem. arXiv preprint arXiv:1911.09539 (2019)
30. Howard, R.A.: Dynamic programming and markov processes. MIT Press, Cambridge (1960)
31. Khalil, E., Dai, H., Zhang, Y., Dilkina, B., Song, L.: Learning combinatorial optimization algorithms over graphs. arXiv preprint arXiv:1704.01665 (2017)
32. Khalil, E., Le Bodic, P., Song, L., Nemhauser, G., Dilkina, B.: Learning to branch in mixed integer programming. *Proceedings of the AAAI Conference on Artificial Intelligence* **30**(1) (2016). URL <https://ojs.aaai.org/index.php/AAAI/article/view/10080>
33. Khalil, E.B., Dilkina, B., Nemhauser, G.L., Ahmed, S., Shao, Y.: Learning to run heuristics in tree search. In: IJCAI, pp. 659–666 (2017)
34. Kool, W., Van Hoof, H., Welling, M.: Attention, learn to solve routing problems! arXiv preprint arXiv:1803.08475 (2018)
35. Land, A.H., Doig, A.G.: An automatic method for solving discrete programming problems. In: 50 Years of Integer Programming 1958-2008, pp. 105–132. Springer (2010)
36. Leyton-Brown, K., Pearson, M., Shoham, Y.: Towards a universal test suite for combinatorial auction algorithms. In: Proceedings of the 2nd ACM conference on Electronic commerce, pp. 66–76 (2000)
37. Liu, D., Fischetti, M., Lodi, A.: Learning to search in local branching. *Proceedings of the AAAI Conference on Artificial Intelligence* **36**(4), 3796–3803 (2022). DOI 10.1609/aaai.v36i4.20294. URL <https://ojs.aaai.org/index.php/AAAI/article/view/20294>
38. Liu, D., Lodi, A., Tanneau, M.: Learning chordal extensions. *Journal of Global Optimization* **81**(1), 3–22 (2021)
39. Maher, S., Miltenberger, M., Pedroso, J.P., Rehfeldt, D., Schwarz, R., Serrano, F.: PySCIPOpt: Mathematical programming in python with the SCIP optimization suite. In: *Mathematical Software – ICMS 2016*, pp. 301–307. Springer International Publishing (2016). DOI 10.1007/978-3-319-42432-3_37
40. Nair, V., Bartunov, S., Gimeno, F., von Glehn, I., Lichocki, P., Lobov, I., O’Donoghue, B., Sonnerat, N., Tjandraatmadja, C., Wang, P.: Solving mixed integer programs using neural networks. arXiv preprint arXiv:2012.13349 (2020)
41. Nazari, M., Oroojlooy, A., Snyder, L., Takác, M.: Reinforcement learning for solving the vehicle routing problem. arXiv preprint arXiv:1802.04240 (2018)
42. Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L.: Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems* **32**, 8026–8037 (2019)
43. Song, J., Yue, Y., Dilkina, B.: A general large neighborhood search framework for solving integer linear programs. arXiv preprint arXiv:2004.00422 (2020)
44. Sonnerat, N., Wang, P., Ktena, I., Bartunov, S., Nair, V.: Learning a large neighborhood search algorithm for mixed integer programs. arXiv preprint arXiv:2107.10201 (2021)
45. Sutton, R.S., Barto, A.G.: Introduction to reinforcement learning, vol. 135. MIT Press, Cambridge (1998)

46. Sutton, R.S., McAllester, D.A., Singh, S.P., Mansour, Y.: Policy gradient methods for reinforcement learning with function approximation. In: Advances in neural information processing systems, pp. 1057–1063 (2000)
47. Tang, Y., Agrawal, S., Faenza, Y.: Reinforcement learning for integer programming: Learning to cut. In: International Conference on Machine Learning, pp. 9367–9376. PMLR (2020)

Statements and Declarations

Funding This work was supported by Canada Excellence Research Chair in Data Science for Real-Time Decision-Making at Polytechnique Montréal. The work of Matteo Fischetti was partially supported by MiUR, Italy.

Conflict of interest The authors declare that have no conflict of interest.

Competing Interests The authors have no relevant financial or non-financial interests to disclose.

Author Contributions All authors contributed to the study and development of this work. All authors read and approved the final manuscript.

Data Availability The datasets used the current study are either public, or available from the corresponding author on reasonable request.

Code Availability The code is publicly available at <https://github.com/pandat8/ML4LB>