



Delft University of Technology

Document Version

Final published version

Licence

CC BY

Citation (APA)

Scavuzzo, L., Aardal, K., Lodi, A., & Yorke-Smith, N. (2026). Machine learning augmented branch and bound for mixed integer linear programming. *Mathematical Programming*, 217(1-2), 123-166. <https://doi.org/10.1007/s10107-024-02130-y>

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

In case the licence states "Dutch Copyright Act (Article 25fa)", this publication was made available Green Open Access via the TU Delft Institutional Repository pursuant to Dutch Copyright Act (Article 25fa, the Taverne amendment). This provision does not affect copyright ownership.
Unless copyright is transferred by contract or statute, it remains with the copyright holder.

Sharing and reuse

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

This work is downloaded from Delft University of Technology.



Machine learning augmented branch and bound for mixed integer linear programming

Lara Scavuzzo¹ · Karen Aardal¹ · Andrea Lodi² · Neil Yorke-Smith¹

Received: 15 December 2023 / Accepted: 24 July 2024 / Published online: 22 August 2024
© The Author(s) 2024

Abstract

Mixed Integer Linear Programming (MILP) is a pillar of mathematical optimization that offers a powerful modeling language for a wide range of applications. The main engine for solving MILPs is the branch-and-bound algorithm. Adding to the enormous algorithmic progress in MILP solving of the past decades, in more recent years there has been an explosive development in the use of machine learning for enhancing all main tasks involved in the branch-and-bound algorithm. These include primal heuristics, branching, cutting planes, node selection and solver configuration decisions. This article presents a survey of such approaches, addressing the vision of integration of machine learning and mathematical optimization as complementary technologies, and how this integration can benefit MILP solving. In particular, we give detailed attention to machine learning algorithms that automatically optimize some metric of branch-and-bound efficiency. We also address appropriate MILP representations, benchmarks and software tools used in the context of applying learning algorithms.

Keywords Mixed integer linear programming · Branch and bound · Machine learning · Metrics · Benchmarks

Mathematics Subject Classification 90C11 · 68T05 · 90C57

This work was partially supported by OPTIMAL, a project funded by the Dutch Research Council (NWO) under grant OCENW.GROOT.2019.015; and by TAILOR, a project funded by EU Horizon 2020 research and innovation programme under grant 952215.

✉ Lara Scavuzzo
L.V.ScavuzzoMontana@tudelft.nl

Karen Aardal
k.i.aardal@tudelft.nl

Andrea Lodi
andrea.lodi@cornell.edu

Neil Yorke-Smith
n.yorke-smith@tudelft.nl

¹ Delft University of Technology, Delft, The Netherlands

² Jacobs Technion-Cornell Institute, Cornell Tech and Technion–IIT, New York, NY, USA

1 Introduction

Mixed Integer Linear Programming (MILP) is a pillar of the field of mathematical optimization. Its widespread use covers multiple application domains such as transportation [14], production planning [85] and energy systems [102]. The theory of modelling and solving MILPs is rich, and has fruitful intersections with several areas of mathematics and computer science [79, 91, 103]. In practice, modern MILP solvers, which all use the tree search method *branch and bound* (B&B) as the main component, are of very high quality and can solve, to provable optimality, problem instances that were considered as completely out of reach only a decade ago. This progress is to a large extent due to remarkable algorithmic advances in the past two decades (see, e.g., Achterberg and Wunderling [2] for an overview).

Continuing these waves of algorithmic progress is key to unlocking new application domains and larger scales, as many areas of potential improvement still exist. For example, several algorithmic decisions that must be made within B&B are made by heuristic rules that have been developed and tuned via computational studies to yield good average performance. For some of these decisions we know that more successful heuristics exist, which are, however, too computationally expensive. Recently, Machine Learning (ML) methodologies have been explored as a potential tool to mimic such time-consuming heuristics for improved performance. This strategy is particularly promising given the increase in data availability, not only in terms of problem instances but also data collected by the solver during the solution process.

This survey addresses precisely the perspective of enhancing key components of B&B by ML so as to integrate machine learning and mathematical optimization as complementary technologies, not as competing ones. We note that collecting data about the solving process and exploiting it to make informed decisions within the algorithm is standard practice in MILP solving. We survey methods that take this idea one step further, in the sense that the mapping between the collected data and the decision is not fixed a priori by an expert, but instead automatically constructed by ‘learning algorithms’ that try to optimize a predefined metric of efficiency.

Scope The survey by Bengio et al. [15] presents the first developments of using machine learning in the context of optimization, with a focus on combinatorial optimization. Our survey is zooming in on the B&B framework in an MILP solver for solving a (class of) MILP. The methods we discuss do not assume any particular problem structure a priori, in the sense that the structure is instead learned. Our focus is on using ML for elements of B&B where choices are being made, such as how to navigate the search tree, how to quickly find good feasible solutions, and how to improve linear relaxations. We do not cover solution methods that replace B&B, such as end-to-end ML approaches. By restricting the attention to the MILP context and its integration with ML, we are able to make a significant step forward in the characterization of choices like the problem representation and of specific aspects like benchmarks and software.

Readership This survey is intended first of all for readers with some background in mathematical programming, and possibly with more limited background in machine

learning. We aim for it to be accessible to those interested in how ML can aid aspects of MILP solvers. Highly detailed knowledge of how contemporary MILP solvers function is not required. We assume some familiarity with B&B.

Outline The remainder of this introduction recalls the definition of an MILP, overviews the main components of a B&B-based MILP solver and presents classical evaluation metrics that will be useful. Section 2 introduces the main concepts of machine learning, with attention to those found most useful for B&B. Section 3 proceeds to survey the tasks in MILP solving where ML can be useful. Section 4 concerns representations of MILP, and in Sect. 5 rounds out the survey by collating instance benchmarks and software used in the literature. Some final remarks and perspectives are given in Sect. 6.

1.1 Mixed integer linear programming

We follow the notation from Wolsey [103], to which the reader is referred to for a general introduction to MILP. We are given a matrix $A \in \mathbb{Q}^{m \times n}$, vectors $c \in \mathbb{Q}^n$ and $b \in \mathbb{Q}^m$, and a partition $(\mathcal{A}, \mathcal{B}, \mathcal{C})$ of the variable index set $\{1, \dots, n\}$. A Mixed Integer Linear Program is the problem of finding

$$\begin{aligned} z^* = \min \quad & c^T x \\ \text{subject to} \quad & Ax \geq b, \\ & x_j \in \mathbb{Z}_{\geq 0} \quad \forall j \in \mathcal{A}, \\ & x_j \in \{0, 1\} \quad \forall j \in \mathcal{B}, \\ & x_j \geq 0 \quad \forall j \in \mathcal{C}. \end{aligned} \tag{1}$$

We write $\mathcal{I} := \mathcal{A} \cup \mathcal{B}$ to denote the set of integer and binary variables. An MILP is said to be mixed binary if $\mathcal{A} = \emptyset$ and binary if $\mathcal{A} = \mathcal{C} = \emptyset$. Relaxing the integrality constraints for variables in $\mathcal{I}$ yields a Linear Program (LP), known as the LP relaxation of the MILP.

1.2 MILP solvers

Modern exact solvers for MILP implement the (LP-based) branch-and-bound algorithm. In our survey, we will stretch the concept of an algorithm to also allow for inclusion of not precisely defined subroutines, such as ‘pick next search node to investigate’. As will become apparent, B&B contains many such subroutines that need to be described in a precise way in order to make it an implementable algorithm.

Notice that throughout the article, we use the term ‘MILP solver’ rather than the, sometimes more common, ‘MIP solver’, the reason being that we especially focus on the learning developed for solving mixed integer *linear* programs, not including the nonlinear case. Nevertheless, most of the technology is used in the solvers for both linear and nonlinear programs. Notably, ML has also been applied to the nonlinear case (see, e.g., [21]).

B&B systematically explores the feasible region by partitioning it into sub-MILPs and obtaining bounds via their LP relaxations, which serve as a mechanism to rule out suboptimal regions. This partitioning scheme can be represented by a search tree, where each node represents a sub-MILP. We use the notation z^{LP_i} to denote the optimal value to the LP relaxation of node i , and we denote x^{LP_i} its optimal solution. At any point during the search, the value of an integer feasible solution provides an upper bound on z^* . The best, i.e., smallest, known upper bound value is called *primal bound*, and we denote it by $\bar{z}$. Similarly, the bounds z^{LP_i} of unprocessed nodes can be used to obtain a lower bound. In particular, $\underline{z} := \min_{i:i \text{ unprocessed}} \{z^{LP_i}\}$ provides a lower bound on z^* and is called *dual bound*. The B&B algorithm ends when all nodes have been processed, when $\bar{z} = \underline{z}$, or when another termination condition (e.g., timeout) applies.

These are the basic principles of the B&B algorithm as implemented in commercial solvers, such as CPLEX [60], Gurobi [52] or Xpress [38], as well as in academic solvers like SCIP [43]. In practice, the execution of the B&B algorithm revolves around some key solver components that handle the different aspects of the solving process. The most important components are preprocessing, the branching rule, the cut management and the primal heuristics.

Preprocessing Most solvers implement a number of procedures that try to reduce the size of the problem and its difficulty, for example by identifying substructures or ways to strengthen the LP relaxation. These routines have been shown to be of crucial importance in speeding the solution process. For an overview of these methods, together with a computational analysis of their effectiveness see [5].

Branching and node selection The procedure of dividing the feasible region is called *branching*. There is a choice to be made with respect to which disjunction is used for branching. The standard¹ is to use single variable disjunctions of the type

$$x_j \leq \lfloor x_j^{LP} \rfloor \vee x_j \geq \lceil x_j^{LP} \rceil,$$

for some $j \in \mathcal{I}$ such that $x_j^{LP} \notin \mathbb{Z}$. Still, the solution to the LP relaxation is likely to violate more than one integrality constraint, which means there is more than one candidate variable for branching. The so-called *branching rule* is the strategy that the solver uses to select a variable for branching. Computational studies have shown that this choice has a critical impact on solver performance [2]. Another important decision, though less critical in terms of solver performance, is *node selection*, where the question is which unprocessed B&B node to consider next.

Cutting planes The LP relaxation can be strengthened by adding valid linear inequalities. These are inequalities that cut off parts of the relaxation but do not exclude any integer feasible solutions. This can, in principle, be done in any node of the B&B tree representing a non-empty feasible region, yet it is standard practice to use cutting planes more heavily, or even exclusively, at the root node. A B&B routine where

¹ More generally, it is also possible to split the problem using multi-variable disjunctions (see, e.g., Karanov and Cornuéjols [61]).

cutting planes are added in other nodes than the root node is typically referred to as branch-and-cut [82]. There is a vast amount of knowledge on cutting planes (or cuts, for short) for MILP (see, e.g., [29, 79, 103]) and most solvers implement a plethora of efficient separation² algorithms. The wide availability of cuts can create an issue: while the addition of cuts strengthens the local relaxation, a large amount of cuts can slow down LP solving and lead to numerical instability. For this reason, a judicious *cut management* strategy [9, 31, 101], which includes separation, selection and removal, is of utmost importance.

Primal heuristics We use the term primal heuristics to refer to routines that try to find feasible solutions in a short amount of time without a success guarantee of doing so. Relying solely on integer feasible LP relaxations to find solutions is most often inefficient. Primal heuristics can provide good solutions early on and help bring down the primal bound $\bar{z}$ more effectively. As with cutting planes, primal heuristics can be used in any node of the tree if desired.

1.3 Evaluation metrics for MILP

In this section, we define a number of metrics that quantify the progress of a branch-and-bound run. We extend the notation just presented in Sect. 1.1 with a variable $t \geq 0$ that represents solving time. We can then define the primal bound $\bar{z}(t) : [0, T_{\max}] \mapsto \mathbb{R}$ and the dual bound $\underline{z}(t) : [0, T_{\max}] \mapsto \mathbb{R}$ as functions of time. We define $\bar{z}(t)$ to be infinity if no integer feasible solution has yet been found at time t . We also make use of a small tolerance value ϵ that is introduced for numerical stability, typically $\epsilon = 10^{-6}$. Notice that many of these metrics make use of the optimal solution value z^* and must therefore be calculated a posteriori, once the instance is solved.

Optimality gap

We define the *optimality gap* as

$$g(t) := \begin{cases} 1 & \text{if no solution has been found yet or } \bar{z}(t) \cdot \underline{z}(t) < 0, \\ \frac{|\bar{z}(t) - \underline{z}(t)|}{\max\{|\bar{z}(t)|, |\underline{z}(t)|, \epsilon\}} & \text{otherwise.} \end{cases} \quad (2)$$

Alternatively, one can track the *integrality gap* defined as $g'(t) = |z^* - \underline{z}(t)|$.

Primal gap and integral

For a given feasible solution $\mathbf{x}$, we define the *primal gap* $\gamma(\mathbf{x})$ as

$$\gamma(\mathbf{x}) := \begin{cases} 1 & \text{if } z^* \cdot \mathbf{c}^T \mathbf{x} < 0, \\ \frac{|z^* - \mathbf{c}^T \mathbf{x}|}{\max\{|z^*|, |\mathbf{c}^T \mathbf{x}|, \epsilon\}} & \text{otherwise.} \end{cases} \quad (3)$$

² The term *separation* refers to the fact that a cutting plane has an effect when it separates a fractional solution of an LP relaxation from the convex hull of (mixed) integer feasible solutions.

We can define a *primal gap function* that maps the solving time to the primal gap of the best solution found up until that point. In particular, denoting by $\mathbf{x}(t)$ the best solution found at time t , we define

$$p(t) := \begin{cases} 1 & \text{if no solution has been found at time } t, \\ \gamma(\mathbf{x}(t)) & \text{otherwise.} \end{cases} \quad (4)$$

The *primal integral* [16] of a process with time limit $T_{max} > 0$ is defined as

$$P(T_{max}) := \int_0^{T_{max}} p(t) dt. \quad (5)$$

An extension of the primal integral is the *confined primal integral* [17] which integrates over an exponentially dampened primal gap function $p_\alpha(t) = p(t)e^{t/\alpha}$ for some $\alpha < 0$. This puts emphasis on the early part of the solution process and avoids a high dependence on the chosen time limit.

Primal-dual integral

One can extend the concept of primal integral to also account for improvements in the dual bound. For this purpose, we integrate over the optimality gap instead of the primal gap, to obtain

$$PD(T_{max}) := \int_0^{T_{max}} g(t) dt. \quad (6)$$

2 A brief introduction to machine learning

This section provides a brief introduction to the key concepts of machine learning that are necessary in order to follow this survey. For a more detailed introduction to machine learning we refer to Mitchell [77].

We are interested in the problem of constructing a mapping from some input data to a desired output space. Let $\mathcal{X}$ be the input data space and let $\mathcal{Y}$ be the output space. It is common to refer to $\mathcal{X}$ as the *feature* space, as it basically represents a set of descriptors of the data samples. The output can be, for example, a prediction based on the input data. The mapping will be constructed by choosing among a family of parameterized functions $f(\cdot, \boldsymbol{\theta}) : \mathcal{X} \rightarrow \mathcal{Y}$ with parameters $\boldsymbol{\theta} \in \Theta$. In short, the objective is to optimize the behaviour of the mapping $f(\cdot, \boldsymbol{\theta})$ by carefully tuning the parameters, based on sampled observations and a progress metric of choice.

To formalize this, we follow standard practices in machine learning and distinguish the following two settings.

Supervised learning The learner has access to a finite collection of pairs $\{(X_i, Y_i)\}_{i=1}^N$, where $X_i \in \mathcal{X}$ and $Y_i \in \mathcal{Y}$. Here, Y_i is the *desired output* to input X_i . It is common to refer to Y_i as the *ground truth* or *label*. The goal is to minimize the *loss function* $l : \mathcal{Y} \times \mathcal{Y} \mapsto \mathbb{R}$, a metric that represents the dissimilarity between

the prediction and the desired output. An example is the mean square error loss (MSE) $\frac{1}{N} \sum_{i=1}^N \left(f(X_i, \theta) - Y_i \right)^2$. Since the true sample distribution is unknown, the loss is minimized with respect to the observed samples

$$\underset{\theta \in \Theta}{\text{minimize}} \sum_{i=1}^N l(Y_i, f(X_i, \theta)).$$

Reinforcement learning Reinforcement learning (RL) is defined in the context of sequential decision making, where actions have long-term consequences and the optimal action is either unknown or too expensive to compute. The learning agent has no access to the ground truth. This methodology is formalized under the framework of Markov Decision Processes (MDPs), see Fig. 1 for a diagram summarizing MDPs. In an MDP, an *agent* interacts with an *environment*. The environment has an associated *state* representing its internal configuration. We denote the state space $\mathcal{S}$. The agent acts on the environment by choosing an action from the action space $\mathcal{A}$ using a *policy* $\pi : \mathcal{S} \mapsto \mathcal{A}$. The agent's action changes the state of the environment, which corresponds to the transition to a new state. The transition mechanism is unknown to the agent.

Apart from the environment's state, the agent can observe a *reward function*. We consider the episodic case, where this interaction between agent and environment happens sequentially in discrete time-steps until a *terminal state* is reached and the interaction ends. A realization of such agent-environment interaction is called *trajectory*

$$\tau := (S_0, A_0, R_1, \dots, S_{T-1}, A_{T-1}, R_T, S_T),$$


where T denotes the episode length. The goal of the agent is to find a policy that maximizes the expected cumulative reward, known as the *value function* and formally defined as

$$V_\pi := \mathbb{E}_\tau \left[\sum_{t=1}^T \gamma^{t-1} R_t \right], \quad (7)$$

where $\gamma \in [0, 1]$ is the *discount factor*. This parameter controls the greediness of the policy: for γ close to zero, the agent will prioritize obtaining immediate rewards, whereas for γ close to one the agent is encouraged to follow a strategy that pays off in the long term. There is a variety of RL methodologies (see, e.g., [92]) that provide ways to train a parametrized function $\pi(s, \theta)$ with the goal of maximizing V_π . Notice that the trajectory distribution depends both on the agent's policy and the unknown transition mechanism of the environment. For this reason, RL algorithms use trajectory sampling as a way to estimate this expectation. Note that $-V_\pi$ can be seen as a loss function.

The MDP formulation can also be used in a supervised learning setting. In this case, the reward signal is substituted by an expert that tells the agent what is the optimal action. The expert is typically expensive to query, and for this reason we want the agent to learn a policy that imitates the expert, but at a lower cost. This methodology is known as *imitation learning*.

Fig. 1 Markov decision process



2.1 Mapping features to predictions

A fundamental step in the design of a ML scheme is the choice of function space parameterized by θ . This is commonly known as the *architecture*, and it is independent of the choice of learning methodology. Modern machine learning favors neural networks as the computational representation and mechanism of f . A simple form of neural network is the feed-forward neural network that consists of a series of linear transformations, each followed by a non-linear transformation called an *activation function*.

Definition 1 (*Feed-forward neural network*) Let $X \in \mathbb{R}^d$ be the input data and let d^L be the desired output dimension. An L -layer *feed-forward neural network* is a function $f : \mathbb{R}^d \mapsto \mathbb{R}^{d^L}$ that defines a mapping from X to $a^L \in \mathbb{R}^{d^L}$ through the recursive relation

$$\begin{aligned} z^l &= W^l a^{l-1} + b^l \\ a^l &= \sigma(z^l) \end{aligned}$$

for $l = 1, \dots, L$, where

- $a^0 = X, d^0 = d$,
- $W^l \in \mathbb{R}^{d^l \times d^{l-1}}$ and $b^l \in \mathbb{R}^{d^l}$ are learnable parameters, and
- σ is the activation function.

Some common activation functions are the Rectified Linear Unit (ReLU) or the sigmoid function, defined, respectively, as

$$\sigma_{\text{ReLU}}(x) = \max(0, x),$$

$$\sigma_s(x) = \frac{1}{1 + e^{-x}}.$$

Notice that these functions are applied componentwise.

Another prominent architecture is the graph neural network (GNN) that will be of particular interest to the discussion in Sect. 4. Consider an undirected graph $G = (V, E)$ with vertex set V and edge set E . For a node $v \in V$, we denote $N(v) \subseteq V$ the set of neighbors of v .




Fig. 2 Embedding computation with a Graph Neural Network. Here, we use the abbreviation ξ_v^t for $\xi^t(G, v)$. To update the embedding ξ_v^t of node v at time t , the embeddings of neighboring nodes are added and then combined with the current embedding using the `comb` function. The result is ξ_v^{t+1}

Definition 2 (*Graph embedding*) A (d -dimensional) *graph embedding* ξ is a function that takes in a graph $G = (V, E)$ and a node $v \in V$ and returns an element $\xi(G, v) \in \mathbb{R}^d$.

Definition 3 (*Graph neural network*) A *Graph Neural Network* is a function that takes as input a graph $G = (V, E)$ and an initial embedding ξ^0 and defines a recursive embedding ξ^t over the vertices of G . This function is characterized by

- A combination function $\text{comb}: \mathbb{R}^{2d} \rightarrow \mathbb{R}^d$, and
- An update rule $\xi^{t+1}(G, v) = \text{comb}\left(\xi^t(G, v), \sum_{u \in N(v)} \xi^t(G, u)\right)$.

Figure 2 depicts one iteration of this recursive embedding mechanism. It is common to refer to a single iteration as a *message-passing operation*. A possible combination function is a feed-forward neural network over the two concatenated vectors. We can extend graph embeddings to edges. This is, for $e \in E$, we can define $\xi(G, e) \in \mathbb{R}^d$. In the presence of edge embeddings, we can redefine a GNN update rule as

$$\xi^{t+1}(G, v) = \text{comb}\left(\xi^t(G, v), \sum_{u \in N(v)} \text{aggr}\left(\xi^t(G, u), \xi(G, \{u, v\})\right)\right) \quad (8)$$

with an aggregation function $\text{aggr}: \mathbb{R}^{2d} \rightarrow \mathbb{R}^d$ that handles the combination of node and edge embeddings. An example of aggregation function is a simple component-wise sum.

2.2 Elements of the learning process

The learning process is itself one of (inexact) optimization, termed *training*.

Optimization algorithms The goal of the training process is to optimize the behaviour of the mapping $f(\cdot, \theta)$ over the parameter space Θ with respect to the loss function. This is done with algorithms for continuous non-convex optimization, which take gradient

descent as their base. Some examples of commonly-used algorithms are AdaGrad [35] or Adam [67]. For reinforcement learning, the situation is more complex, in the sense that several learning paradigms are available for approximating the value function and learning the optimal policy.³

Hyper-parameters There is a number of additional parameters that influence the learning process that are not part of the parameters θ to be learned. These meta-level parameters are called *hyper-parameters*. Examples of hyper-parameters are the number of iterations of the optimization algorithm, the learning rate, the size of the learned function, etc. Tools exist to automatically tune hyper-parameters given a validation dataset (see below).

Train, validation and test datasets The data used to tune the parameters of mapping f is called the *training data*. The data used to estimate the performance of f in an unbiased way, and to tune the hyper-parameters, is called the *validation data*. Finally, the data used to evaluate the final trained and optimised mapping is called the *test set*. These sets must be disjoint.

Data collection The learner makes use of the available training data to improve its behaviour. This data can be gathered all at once during an initial data collection phase, before training starts. Alternatively, several data collection and training steps can be run in an alternating fashion.

Overfitting We say a learned model suffers from *overfitting* when the performance on the training set is distinctively better than the performance on the test set. In other words, the model is unable to generalise to data unseen during the training phase. This phenomenon occurs when the chosen model is too large (in terms of number of parameters) for the task, or the amount of relevant data is too limited for the size of the model being learned. See Goodfellow et al. [48] for a discussion of methods to avoid overfitting.

Online versus offline learning Whenever we can distinguish separate training and execution phases we speak of *offline learning*. That is, in offline learning, we perform the data collection and learning as a separate, preliminary, once and for all process. After that, the learned function is fixed and used without further tuning. This training phase can be computationally costly but this cost is not reflected in the execution time (often called inference time) of the algorithm, which is typically the metric of interest to the end user. In fact, this can be seen as doing the heavy work upfront while alleviating the effort at the moment of execution.

In contrast, in an *online learning* setting, data becomes available during execution. Learning must therefore happen dynamically as each data sample becomes available, and in parallel to the main process. This has the great advantage of generating a function that is highly adaptive. On the other hand, it entails adding the cost of learning to the cost of execution and further there is usually less available data.

³ A general discussion of those paradigms is outside the scope of the article and special cases will be discussed while surveying specific results from the literature, see Sect. 3.

3 Learning tasks

So far, we have introduced the elements of the B&B algorithm for MILP, and introduced general terminology in machine learning. This section presents different ways to formulate a learning task within the B&B algorithm. We split this discussion by considering each major solver component separately.

Before getting into the details, two considerations are required. First, throughout the section we will use the symbol X to abstractly denote a representation of an MILP instance. This representation may include data coming from the problem description (such as in Eq. 1) and from the B&B process. Section 4 discusses, in more detail, approaches to build such a representation. Second, it is important to highlight that the learning tasks that we describe next might be associated with different levels of required generalization. More precisely, the ML models are often trained on data belonging to MILP instances in the same class, for example, set covering, knapsack, etc. In other words, the characteristics of the constraint matrix in (1) are leveraged to obtain accurate predictions, so the generalization power of the resulting models is generally restricted to that specific MILP class used for training. However, we will also review cases in which learning happens instance by instance or extends outside of a known MILP distribution. We will highlight these differences throughout the entire review of the learning tasks.

3.1 Primal heuristics

Primal heuristics play a crucial role in quickly finding feasible solutions and consequently improving the primal bound $\bar{z}$ in the early stages of the solve. In the context of primal heuristics, ML techniques can be of interest to leverage common structures in the instances. A number of methodologies have been proposed for this purpose (Table 1). Conceptually, they can be split into three main approaches: (a) guiding a heuristic search with a starting predicted solution, (b) solution improvement via a learned neighborhood selection criterion, and (c) learning a schedule to pre-existing heuristic routines. These methodologies are summarized in Fig. 3 and will be discussed in more detail in the following. At the end of the section, Table 2 presents a summary of the reviewed work and their main characteristics.

An important class of primal heuristics are the so-called Large Neighborhood Search (LNS) [94] heuristics. The idea is to optimize an auxiliary MILP of smaller size, constructed by reducing the feasible region of the original MILP. This is typically done by fixing the value of some of the variables and optimizing the rest. Another strategy is to search over the neighborhood of a solution $\hat{x}$ by imposing the constraint

$$\sum_{j:\hat{x}_j=0} x_j + \sum_{j:\hat{x}_j=1} (1 - x_j) \leq \eta, \quad (9)$$

with η the parameter that controls the neighborhood size. This is known as *local branching* [39]. The reader is referred to Fischetti and Lodi [40] for a survey on the use of LNS heuristics in MILP.




Fig. 3 Three learning problems related to primal heuristics: **a** predict a reference solution and search in its neighborhood, **b** neighborhood selection—which and/or how many variables to unfix and re-optimize, **c** heuristic scheduling—which heuristics to run and/or for how long

3.1.1 Solution prediction to guide the search

Some authors explore the idea of using predictions of the optimal solution. The goal is to produce a (partial) assignment of the binary variables⁴ in a binary or mixed binary MILP, that can then be used to guide the search. This idea has been implemented in different ways, both in terms of how to obtain the predictions and how to use them. Let us start by discussing the latter.

Ding et al. [34] use a local branching constraint (see Eq. 9) with respect to predicted values of a subset $J \subseteq \mathcal{B}$ of the binary variables. This restricts the search to a neigh-

⁴ Note that the focus on binary variables is justified because, in most well-established instance collections, binary variables account for the great majority of integer variables. As an example, more than 68% of the instances in the MIPLIB [76] are purely binary, and in the remaining ones, more than 90% of the integer variables are in fact binary. See, e.g., Nair et al. [78] for extensions to integer variables.

Table 1 Three methods to define the target probability distribution p_T , given a collection $D(X) = \{\hat{x}^{(1)}, \hat{x}^{(2)}, \dots, \hat{x}^{(K)}\}$ of feasible solutions to problem X

Authors	Target
Ding et al. [34]	$p_T(x_j = 1 X) = \hat{x}_j^{(1)}$ for every j s.t. $\hat{x}_j^{(k)} = \hat{x}_j^{(1)}$ for all $k \in \{1, \dots, K\}$
Nair et al. [78]	$p_T(\hat{x}^{(k)} X) = e^{-c^T \hat{x}^{(k)}} / \sum_{i=1}^K e^{-c^T \hat{x}^{(i)}}$ and $p_\theta(\hat{x} X) := \prod_{j=1}^n p_\theta(x_j = \hat{x}_j X)$
Khalil et al. [65]	$p_T(x_j = 1 X) = \frac{1}{K} \sum_{k=1}^K \hat{x}_j^{(k)}$

The goal is to learn a function $p_\theta(x_j = 1|X)$ that represents the probability that variable x_j takes value 1 given problem X by imitating the target probability $p_T(x_j = 1|X)$

borhood of the predicted partial solution. In contrast, both Nair et al. [78] and Khalil et al. [65] propose to fix the variables in J to their predicted value and hand over this partial assignment to an MILP solver that optimizes over the remaining variables. In MILP terminology, this corresponds to a warm start. Khalil et al. [65] further use the predictions to guide node selection (see Sect. 3.4).

The question still remains as to how to obtain these predictions. This problem is naturally formulated as a supervised learning task, where the desired output is the optimal solution. However, optimal solutions to be used as labels in the training phase are costly to obtain and do not capture information about the region where they lie. The aforementioned work of Ding et al. [34], Nair et al. [78], Khalil et al. [65] make use of a set of feasible (not necessarily optimal) solutions to learn predictions. This is, the learning process starts with a data collection phase that, for each instance X , gathers a set of feasible solutions $D(X) = \{\hat{x}^{(1)}, \hat{x}^{(2)}, \dots, \hat{x}^{(K)}\}$. Once the data collection phase is finished, the prediction task can be formulated as learning a probability distribution $p_\theta(x_j = 1|X)$ that represents the probability that variable x_j takes value 1 given problem X , and is parametrized by θ . The parameters are tuned to make the behavior of these functions resemble as closely as possible that of a target probability distribution $p_T(x_j = 1|X)$. Table 1 summarizes the proposed targets.

3.1.2 Solution improvement via neighborhood selection

Alternatively to solution prediction, one may be interested in learning a *destroy heuristic* criterion. That is, given an initial feasible solution $\hat{x}$, we select a subset of the integer variables to be re-optimized, leaving the remaining integer variables fixed. This process can be run iteratively. The goal is to identify substructures of the problem that can be used to decompose it into smaller, more manageable sub-problems.

Following this line, Song et al. [95] learn to partition the set of integer variables $\mathcal{I}$ into K disjoint subsets $\{S_i\}_{i=1}^K$ such that $\mathcal{I} = S_1 \cup \dots \cup S_K$. They iteratively unfix and reoptimize the variables in each subset, fixing the rest to the value in the best known solution. The variable subsets are of fixed size, which means that the variables simply need to be classified into the subset they belong to. Alternatively, Wu et al. [104], Sonnerat et al. [96] and Huang et al. [58] propose a method to select a single flexible-sized

subset of variables to unfix. Their approach resembles the value-prediction methodology discussed in the previous section. However, instead of predicting the value that a variable takes in an optimal solution, they aim at predicting whether the variable is already assigned to its optimal value in the current best solution $\hat{x}$. The number of unfixings is learnt implicitly through the conditionally-independent probabilities $p_\theta(\hat{x}_j = x_j^* | X)$ from which they sample variable unfixings. While Wu et al. [104] use a reinforcement learning algorithm to train their policy, Sonnerat et al. [96] explicitly calculate the best solution at most η unfixings away from $\hat{x}$ and use it as a target unfixing policy. Huang et al. [58] follow a similar approach, while also providing the learner with negative examples. This means variable unfixings that do not lead to a sufficiently large improvement, as compared to the best. This provides additional information for the learner to distinguish good and bad unfixings.

The methods of Song et al. [95], Wu et al. [104], Sonnerat et al. [96] and Huang et al. [58] were compared against continuously running the MILP solvers they use as a subroutine. Given the same amount of time, the ML-assisted methods are able to find better solutions. Interestingly, using a random selection of the variables to unfix often results in a performance improvement compared to running the solver continuously. This highlights the fact that MILP solvers are typically tuned to optimize different performance metrics, such as the optimality gap, which means that a significant part of the computational effort is spent in, e.g., obtaining good dual bounds. Nonetheless, the proposed methods provide a relevant methodology to make use of the MILP solver in a black-box manner when the goal is to obtain high-quality but not necessarily optimal solutions in a short amount of time.

Related to this line of work, Liu et al. [74] point out that the optimal value of the neighborhood size parameter η is strongly dependent on the class of instances being solved. They work in the context of local branching and propose to learn two policies. A first policy $f_{init}(X)$, obtained in a supervised way, determines an appropriate neighborhood size η_0 for a first iteration. Another policy $f_a(X)$, obtained by RL, decides how to *adapt* the neighborhood size in successive iterations given the previous solver statistics. Compared to the simple neighborhood size selection proposed in Fischetti and Lodi [39], these two policies close the primal gap (see Eq. 3) faster, especially when combined together. Their computational experiments demonstrate that these selection rules have great generalization ability across different instance sizes and types, especially for $f_a(X)$, also across different classes of instances, proving the wide applicability of their approach.

3.1.3 Learning to schedule heuristics

Whether they use ML or not, many successful primal heuristics for MILP have been proposed. Experimentally, we observe that no single heuristic dominates the others on all problems [55]. Their performance is highly dependent on the problem, and even on the solving stage. A new decision problem arises: given a collection of primal heuristics, which one should be run? In a broader sense, heuristic scheduling also tries to answer questions such as how frequently should heuristics be run, or under what computational budget. In this section, we review the line of research on using learning methods to answer such questions.

Hendel [54] proposes Adaptive Large Neighborhood Search (ALNS), a heuristic that, whenever called, makes a choice among eight Large Neighborhood Search methods. This choice is framed as a Multi Armed Bandit problem [23], an online learning methodology that learns a selection policy per instance. The learned policy is based on the observed (a posteriori) performance of the chosen heuristics. This simple formulation encapsulates the classical exploitation versus exploration dilemma: we must balance running heuristics that have performed well in the past with running heuristics whose performance is unknown because they have not been selected a sufficient number of times. The author uses the α -UCB algorithm [23] with a reward function that combines several aspects of a heuristic's performance: whether an incumbent was found, the objective improvement and the computational time the heuristic needs. The experiments on MIPLIB 2017 [46] show a considerable improvement in the primal integral (see Eq. 5), as well as a speed-up. The Multi Armed Bandit formulation was also used by Hendel et al. [55] and Chmiela et al. [27], who extend it to new types of heuristics.

Chmiela et al. [26] take a different perspective: instead of a selection policy per instance, they create a prioritization order that applies to a given class of instances. In fact, these authors' method constructs a schedule that assigns both a priority and a computational budget to each heuristic. In contrast with the abovementioned work of Hendel et al. [55] and Chmiela et al. [27], which use online learning, the schedule is now crafted offline, after a data collection phase where different heuristics are evaluated. They formulate a MIQP that minimizes the computational budget while finding feasible solutions for a large fraction of the B&B nodes, and they solve this using a greedy heuristic.

The proposed method produces schedules for a number of different heuristics, which result in improvements in the primal integral compared to default SCIP and even a version of the solver that has been manually tuned for the considered class of instances.

Other than deciding which heuristics to run and for how long, there is the question of *when* to run them. The work discussed so far uses conventional rules to decide at which nodes heuristics will be run. The methodology proposed by Khalil et al. [64] is to instead build a mapping between B&B nodes and a yes/no decision for running each heuristic.

This decision problem is challenging. Even with perfect knowledge of when a heuristic will be successful, a run-when-successful rule does not necessarily minimize solving time. Khalil et al. [64] analyze the competitive ratio of such a rule compared to an optimal offline decision policy. They then imitate the imperfect run-when-successful rule by learning to predict when a heuristic will succeed.

Their computational study shows the effectiveness of this method: heuristics are called less often, but with a higher success rate, resulting in a better primal integral. This effect is even stronger when the policy is trained on data coming from the same instance class.

Table 2 Summary of different learning approaches for primal heuristic management

	ML paradigm	Online/offline	Model
Ding et al. [34]	SL	Offline	GNN
Nair et al. [78]	SL	Offline	GNN
Khalil et al. [65]	SL	Offline	GNN
Song et al. [95]	SL & RL	Offline	PCA [49]+NN
Wu et al. [104]	RL, policy gradient	Offline	GNN
Sonnerat et al. [96]	SL, imitation	Offline	GNN
Huang et al. [58]	SL, contrastive	Offline	GNN
Liu et al. [74]	SL & RL	Offline	GNN
Hendel [54]	Multi-Armed Bandit	Online	–
Chmiela et al. [27]	Multi-Armed Bandit	Online	–
Chmiela et al. [26]	Greedy heuristic	Offline	–
Khalil et al. [64]	SL, classification	Offline	Logistic regression

We use the acronyms SL for supervised learning and RL for reinforcement learning

3.1.4 Future outlook

The success of most of the approaches discussed in Sect. 3.1 is rooted in the ability of learning within the distribution of specific classes of MILP instances. Generalizing outside of a specific class, has proven so far difficult. This is certainly the main challenge for the integration of ML-augmented primal heuristics within MILP solvers. Studies like the one in Liu et al. [74] indicate that some significant generalization is achievable, especially in the context of algorithms that sequentially adapt while exploring the solution space as in classical RL schemes. This is a promising direction that, however, conflicts with the MILP solver's general need of executing (extremely) fast primal heuristics. Further, in the context of solution prediction, the case of general integer variables has been comparatively less studied; see Nair et al. [78] for a discussion of how to treat this case.

3.2 Branching

Branching is one of the core mechanisms on which the B&B algorithm operates. The branching rule, i.e., the criterion used to select a variable for branching, has been identified as having a critical impact on performance [2, 73]. Computational studies have served to identify a number of metrics that are good indicators of how a variable will perform. Notably, state-of-the-art branching rules look at the change in objective value in the resulting children nodes. More specifically, let z^{LP_i} be the objective value of the current node i and let $z^{LP_{i+}}$ and $z^{LP_{i-}}$ be the objective values of the two nodes resulting from branching on candidate variable x_k . This variable will be scored using a combination⁵ of $\Delta^+ := z^{LP_{i+}} - z^{LP_i}$ and $\Delta^- := z^{LP_{i-}} - z^{LP_i}$. These values can be

⁵ For example, the variable's score can be computed as $\max\{\Delta^+, 10^{-8}\} \cdot \max\{\Delta^-, 10^{-8}\}$. See, e.g. Achterberg [1] for a discussion.

explicitly calculated for each candidate at the moment of branching (thus solving two LPs per candidate). Such strategy, known as *strong branching*, was introduced in the context of the travelling salesman problem [10] and later standardised by CPLEX.

Strong branching is known to produce small B&B trees, but at high computational cost per branching. Alternatively, one can attempt to estimate the objective change based on past values, once they become naturally available through node processing. In particular, solvers typically store the values of Δ^+ and Δ^- normalized by the variable's fractionality, and keep track of the per-variable average, known as *pseudocosts*. The so-called *reliability branching* rule [3] performs strong branching at the top of the tree as an initialization phase, and then switches to using pseudocosts as soon as a variable has been branched on enough times. The initialization phase is not only important to build a branching history, but also because branching decisions have the most impact at the top of the tree. This is because judicious branching decisions here can lead to much smaller tree sizes, due to the earlier finding of feasible solutions and the stronger pruning of nodes.

In this section, we will discuss different approaches to learning to branch. Their common goal is to learn a function that maps a description of the candidate variables to scores. We note that alternative strategies exist, such as learning a weighting scheme for a portfolio of branching rules (see, e.g., Balcan et al. [12]). Ultimately, the objective is to minimize the solving time, which typically entails a favorable balance among different sub-targets. For example, (i) being computationally cheap, (ii) generating small trees as a result of their scoring, and (iii) adapting to the different situations that may arise. These objectives are often at odds with each other. At the end of the section, Table 3 presents a summary of the reviewed work and their main characteristics.

3.2.1 A first approach to learning from strong branching

There is a well-established body of research on fast approximations of the strong branching rule. This idea was first explored by Alvarez et al. [7] who propose to learn a prediction of the strong branching score of each variable. The predictor is learned offline (see Sect. 2) and tested on both heterogeneous and homogeneous instance collections. The experiments of Alvarez et al. [7] show that the method achieves the desired objective of imitating strong branching decisions without the large computational overhead. Indeed, when compared with strong branching on a fixed number of nodes, the closed gap is only moderately worse, indicating that the branching decisions are of high quality. Simultaneously, for a fixed time, the method explores a much larger number of nodes and closes a greater proportion of the gap, therefore achieving an overall better trade-off. However, reliability branching still outperforms the learned branching rule in closed optimality gap. The experimental results seem to indicate that the former makes smarter decisions, and is on average faster in making them. It is interesting to note that the method of Alvarez et al. [7] performs better on homogeneous instance collections, i.e., when the problems in the training and the test set are of the same type.

3.2.2 Online learning to branch

The results presented in Alvarez et al. [7] call for several reflections. First, experiments seem to indicate that performing strong branching at the top of the tree, where branching decisions have the most impact, is highly advantageous. Second, the authors already point towards more adaptation to the problem structure as a promising direction of improvement.

These ideas are studied in Khalil et al. [63] and the follow-up work of Alvarez et al. [6]. Independently, these authors proposed to use online learning with a strong branching initialization at the top of the tree. One key difference is that, while Alvarez et al. [6] continue to frame learning to branch as a prediction task (i.e., predicting the strong branching score of a variable), Khalil et al. [63] formulate the task as that of *ranking*. Indeed, one does not necessarily need to accurately predict the variables' scores, but rather which variables have relatively better ones. The latter task is easier from a learning perspective.

Framing an online learning task allows to learn a specialized ML model per instance, in this case a simple linear function. Khalil et al. [63] analyse their learned models by studying the weight assigned to each of the features. The first question they ask is: are the learned models obtained for each of the instances similar? Their analysis concludes that there is only a weak correlation among the models. This supports the idea that adaptation plays a key role. In spite of the learned models being quite different, they were able to identify some features that are often given high importance (large absolute weight), such as the product of the pseudocosts, and data related to the constraint matrix. For a more detailed discussion of this analysis see Khalil [62].

3.2.3 Offline learning with structure specialization

Alvarez et al. [6] and Khalil et al. [63] use a linear mapping from variable features to output. This offers the advantage of interpretability and low computational cost. The more complex GNN model proposed by Gasse et al. [44] (see Sect. 4 for a description) has proven to be remarkably effective in representing MILPs for the task of variable selection and beyond. The GNN architecture that we presented in Sect. 2.1 consists of a number of parametric function compositions that enable learning more complex relations between input and output. They also require a larger amount of training data, which rules out the online learning methodology previously discussed. In order to still ensure some level of specialization, Gasse et al. [44] propose a middle point: they argue that in many realistic cases instances of the same *class* are routinely solved. This justifies learning a branching rule per instance type, as a sensible trade-off between a completely general rule (such as the one in Alvarez et al. [7]) and a completely instance-tailored rule (like in Alvarez et al. [6], Khalil et al. [63]).

Gasse et al. [44] propose training this GNN to imitate strong branching via behavioral cloning [86]. In short, this means that we again disregard the actual variable scores and focus on learning relative magnitudes among them. Through this approach the authors were able to outperform reliability branching, marking a breakthrough in the learning to branch literature. Building on this work, Gupta et al. [50, 51] propose modifications of the original loss function that further improve the performance of the

learned model. Nair et al. [78] test the methodology of Gasse et al. [44] on a variety of instance types, including heterogeneous sets. Seyfi et al. [93] additionally propose a mechanism that incorporates information about past branchings into the scoring system.

3.2.4 Towards a general branching rule

We have discussed methodologies that specialize to certain combinatorial structures (such as Gasse et al. [44]) or that yield a custom strategy per instance (e.g., Khalil et al. [63]). As discussed at the beginning of this section, these methodologies are a response to the great difficulty of learning one unique policy on an heterogeneous set of instances. Zarpellon et al. [106] argue that, while learning such a general policy poses a big challenge, it is possible to overcome it by using a representation of the search tree to inform variable selection. Their hypothesis is that there is a higher-order shared structure among MILPs, even among those with different combinatorial structure, and that this shared structure can be captured in the space of B&B trees. To test this they define a set of features describing the state of the search which, together with variable descriptors, are mapped into scores. We discuss these features in more detail in Sect. 4. The experiments in Zarpellon et al. [106] show that, while the model of Gasse et al. [44] struggles to learn over a heterogeneous data distribution, adding the tree context is beneficial to the generalisation performance of the model. Lin et al. [72] take this idea one step further and propose to keep a record of the features in [106] at each node where branching was performed. At every step, branching decisions are informed by this historical data, which is carefully aggregated and combined with the descriptions of the variables.

The work in [72, 106] highlights the potential of using solver statistics to influence branching decisions. It is unclear how the features proposed in Zarpellon et al. [106] are used to score branching candidates, since neural networks lack explainability. Nonetheless, this information certainly opens the door to branching rules that switch among different behaviors at different parts of the tree, or stages of the solving process. It is important to note that reliability branching still outperforms both Zarpellon et al. [106] and Lin et al. [72], perhaps because of the overhead of computing and processing these comprehensive descriptors. Nonetheless, this work calls for further research on exploiting tree information.

3.2.5 Expert-free learning to branch

So far, we have discussed methods for building fast approximations of strong branching. The general consensus is that strong branching yields relatively small B&B trees compared to other classic branching strategies, and it is therefore advantageous to try to imitate it. This idea can be challenged with two arguments. First, as pointed out by Gamrath and Schubert [42], standard implementations of strong branching benefit from using data obtained as a by-product of the score calculation, such as bound tightenings and other statistics. A rule that *imitates* strong branching cannot profit from such side-effects, which means that, even in the case of perfect emulation, the imitator's performance would be worse than expected. Second, strong branching relies on

Table 3 Summary of different learning approaches for branching

	ML paradigm	Online/offline	Model
Alvarez et al. [7]	SL, regression	Offline	ExtraTrees
Alvarez et al. [6]	SL, regression	Online	Linear
Khalil et al. [63]	SL, ranking	Online	Linear
Gasse et al. [44]	SL, imitation	Offline	GNN
Etheve et al. [37]	RL, Q-learning	Offline	NN
Scavuzzo et al. [90]	RL, policy gradient	Offline	GNN
Zarpellon et al. [106]	SL, imitation	Offline	NN
Lin et al. [72]	SL, imitation	Offline	Transformer

We use the acronym SL for supervised learning and RL for reinforcement learning

the LP relaxation for scoring variables, which can provide little information in cases when the optimal LP objective value does not change with branching. In such cases strong branching cannot be considered a reliable expert. In the absence of a better alternative to strong branching, the following question arises: can we learn branching rules without expert knowledge? This question is addressed by Etheve et al. [37] and Scavuzzo et al. [90], who use RL to learn branching rules from scratch. Etheve et al. [37] prove that when using a depth-first-search (DFS) node traversal, branching to minimize the size of each subtree is equivalent to minimizing the size of the entire B&B tree. Scavuzzo et al. [90] build upon this and define a new framework for learning to branch, the *tree MDP*, which encodes the intuition that branching decisions have a larger impact on the subtree they create than in other parts of the B&B tree. The computational study in Scavuzzo et al. [90] compares one such RL methodology against the imitation learning method of Gasse et al. [44]. The experiments show that on instances where strong branching performs very well, the imitation method of Gasse et al. [44] is superior. Yet, when strong branching struggles, the RL-based method is able to find a better branching strategy, proving the point that expert-free learning is of great interest in certain cases.

The branching strategies proposed in Etheve et al. [37] and Scavuzzo et al. [90] do not follow demonstrations by another rule. One may therefore wonder if there is an interpretation for the learned behavior. It is unfortunately difficult to answer this question, because the architectures they use are not interpretable.

3.2.6 Future outlook

Adaptiveness The computational studies suggest that no single branching rule outperforms others universally across all instances. Consequently, a desirable approach involves a rule that dynamically adapts its behavior to the specific characteristics of each situation. One way in which one can introduce such adaptiveness is by controlling the distribution of data samples that the model uses for learning. Some authors propose to specialize to each given instance. This means that a set of parameters θ is generated for each new instance, obtained by allowing the ML model to only see data coming from that instance. Another possibility is to use data samples coming from

instances of the same problem class. This gives us parameters θ that specialize to a given class and work well, on average, on different instances with shared combinatorial structure. We can therefore achieve adaptiveness on different scales. We can also understand adaptiveness of a branching rule as some sort of mechanism to change its behavior on different parts of the search process. Some progress has also been made in this regard by investigating different statistical measures that can inform branching (see Sect. 3.2.4). Yet, little is understood about how these metrics are used or in which ways we can further enhance performance without sacrificing speed.

Expert guidance The strong branching heuristic has been used by many as an expert from which we can learn effective decision-making. The claim that strong branching is a desirable strategy to follow has recently been challenged, with some notable examples of instances where strong branching scores provide no useful information. Interesting research directions include identifying new experts, new strategies to better imitate them, or, conversely, more efficient approaches to learning without expert knowledge. The latter case is what is referred to in Bengio et al. [15] as *experience*: there is no clear mathematical understanding of what should be statistically learned (the expert), so exploration should be performed. In turn, this clearly calls for RL methods that are also natural candidates to extend the work in Zarpellon et al. [106], Lin et al. [72] and exploit tree information.

New directions The work we surveyed showcases the potential in mixing the extensive body of domain knowledge in variable selection with new learning techniques. Still a lot of open questions remain. For example, little attention has been directed towards highlighting important subsets of variables, as opposed to choosing a single one at each node. Khalil et al. [66] propose an approach to finding such important subsets, in this case the so-called *backdoors* [33], and show promise in using them as prioritized branching candidates. Another relevant gap is the absence of expert knowledge for certain classes of MILPs. In any case, it is clear that new ML-based methods need to build upon the pre-existing knowledge on variable selection to achieve a fruitful combination.

3.3 Cutting planes

Cutting plane routines are another essential part of modern MILP solvers. They tend to work in *rounds*, also called *separation rounds*. Given an LP-relaxation solution x^{LP} , one round consists of generating a number of valid cuts from different families, selecting a subset of them via a selection criterion, adding them and finally resolving the LP relaxation, where x^{LP} will now be infeasible. A good selection criterion is critical to improving the LP relaxation while avoiding an excessive number of cuts, which would slow down LP solving as well as lead to numerical instability [9, 31]. Several metrics have been proposed for the purpose of scoring cuts (see, e.g., Wesselmann and Stuhl [101] for an overview). For example, the *objective parallelism*, measured as the cosine of the angle between the objective function and the cut, or the *cutoff distance*, measured as the distance between the cut and the LP-relaxation solution. More recently

the question of cut selection has been addressed with ML-driven predictions, which is the topic of this section. We review different work in this area and, at the end, present in Table 4 a summary of the different methodologies and their main characteristics. For further discussion of ML for cut selection in MILP and beyond we refer to Deza and Khalil [32].

As noted in Sect. 1.2, cuts are typically more heavily applied at the root node, and for this reason the work that we survey focuses on cut selection at the root node. Still, there is no obstacle to applying these methods in other nodes. However, it is unclear whether using cuts outside the root node is computationally beneficial (see Berthold et al. [19] for a discussion of this topic).

3.3.1 Single-cut selection

Tang et al. [97] and Paulus et al. [84] frame the task of cut selection as an MDP. At each step k , a single cut c_k is selected from a cutpool $\mathcal{C}_k$, after which the LP relaxation is resolved. In particular, let $\mathcal{C}$ be a collection of cuts, and let $z(\mathcal{C})$ be the result of solving the root LP relaxation after adding all cuts in $\mathcal{C}$. The metric of interest to these authors is the LP-bound improvement attained by a cut c at step k , defined as

$$\Delta_k(c) := z(\{c_1, \dots, c_{k-1}, c\}) - z(\{c_1, \dots, c_{k-1}\}) .$$

This MDP model is summarized in Fig. 4a, where we use the abbreviation $z_k = z(\{c_1, \dots, c_k\})$.

Paulus et al. [84] use imitation learning. Their expert is the result of explicitly calculating $\Delta_k(c)$ for each possible cut c in the cutpool $\mathcal{C}_k$, and then picking the cut with highest Δ_k . This means that their approach is greedy: they look at bound improvement one step ahead. Their computational studies show that this greedy heuristic that they aim to imitate is in fact very effective in improving the LP bound after T cuts have been added, compared to other selection heuristics.

Another approach is that of Tang et al. [97], who use RL with reward $R_k = \Delta_k(c_k)$. Due to the discount factor (see Eq. 7), this RL strategy offers the possibility to learn less greedy policies, and doing so without explicitly computing $\Delta_k(c)$ for each k and $c \in \mathcal{C}_k$. However, many RL algorithms are known to suffer from sample inefficiency and lack of generalization [36, 68]. Paulus et al. [84] compare their approach to the one of Tang et al. [97] and to SCIP's v.7.0.2 default rule. They use the sum of gaps as a metric (lower value is better)

$$SG := \sum_{k=1}^T \frac{z^* - z_k}{z^* - z_0} , \quad (10)$$

with $T = 30$ being the total number of cuts added, and z^* being the pre-computed optimal solution. Notice that this metric is constructed in a way that might favour greedy policies.⁶ The computational results favor the method of Paulus et al. [84]. The

⁶ This can be seen through the following example. Without loss of generality, let $z^* = 1$ and $z_0 = 0$. Consider the following cases using two cuts: (i) one with $z_1 = 1/2$ and $z_2 = 3/4$, and (ii) one with $z_1 = 0$




Fig. 4 Three models for learning to cut: **a** Tang et al. [97], Paulus et al. [84], **b** Wang et al. [100], **c** Turner et al. [98]. Here $PD(t_{max})$ refers to the primal-dual integral (see Sect. 1.3)

authors also show promising results in solving time when their method is incorporated into SCIP and instances are solved to optimality, though they do not include root node processing time.

3.3.2 Multi-cut selection

A potential criticism to the cut selection models is that solvers usually add more than one cut per round, in order to reduce the number of times the LP needs to be resolved. In fact, historically this proved crucial to the efficiency of cutting plane routines [11, 30]. Having information about the LP solution after the addition of each cut is therefore unrealistic. Paulus et al. [84] do not include root node processing times in their report, a metric under which their method is likely unfavored. Furthermore, metrics like the one in Eq. 10 encourage greedy bound improvements, whereas in practice cuts can work together to achieve a better bound improvement at the end of the round. In other words, the optimality gap closed after a full separation round is likely a more informative metric.

To address the potential interactions among cuts, Wang et al. [100] propose a policy that decides the fraction of cuts from the pool to be selected, and scores ordered subsets of this size. See Fig. 4b for a summary of this selection model. The authors train the policy with an RL algorithm and use end-of-run statistics like solving time as the reward. This requires solving an MILP to optimality for each training sample. While this allows to learn a mapping from cut selection to actual solver performance (instead of the root node bound improvement, which is just a proxy for performance) this sample collection comes at a great computational cost.

and $z_2 = 1$. In the first case, we have $SG = 3/4$, and in the latter, $SG = 1$. The first choice is therefore preferred, even though the second one closes the gap completely.

A third model to learning cut selection is proposed in the work of Turner et al. [98] (see Fig. 4c). The procedure builds upon SCIP's default strategy, a rule that has been carefully curated through computational studies [1, 101]. This rule combines four cut scoring functions $s_i : \mathbb{R}^{n+1} \rightarrow \mathbb{R}_+$, $i = 1, \dots, 4$ via a convex combination

$$s(c, \mu) = \mu_1 s_1(c) + \mu_2 s_2(c) + \mu_3 s_3(c) + \mu_4 s_4(c),$$

$$\sum_{i=1}^4 \mu_i = 1, \quad \mu_i \geq 0, \quad i \in \{1, 2, 3, 4\}.$$

Cuts are then selected sequentially based on their score (the higher the better) and their parallelism to already selected cuts (too 'parallel' cuts are discarded). We refer to Turner et al. [98] for the definition of the scoring functions s_i , $i = 1, \dots, 4$. The problem of choosing a good set of parameters $\mu \in \mathbb{R}^4$ has been studied from a learning theoretical perspective by Balcan et al. [13]. As a next step, Turner et al. [98] argue, through both theoretical and computational arguments, in favor of adapting the coefficients μ to the instance being solved, as opposed to a unique, static choice. One of their experiments consists of finding a custom set of parameters μ per instance through grid search. While impractical, this experiment uncovers the large potential for improvement when adapting the value of μ . In order to exploit this potential in a more realistic way, Turner et al. [98] devise a policy that, given an instance, chooses custom parameters μ , and they train it via RL. The computational study shows that the learned policy is competent in its task, outperforming random selection in terms of closed optimality gap at the root node. However, and perhaps surprisingly, their policy does not perform consistently better than SCIP's default settings when it comes to final solving time.

3.3.3 Beyond scoring

Cut scoring for selection is an essential part of cut management. Yet, there are other important decisions. Wang et al. [100] incorporate the number of cuts added as a decision that the policy must make. Very recently, Li et al. [71] defined a learning task for separator configuration. The objective is to select a subset of the available separators. Only the selected separators will then be active and contribute to the cut-pool, meaning that this selection step happens *before* the cut selection phase. Li et al. [71] propose a methodology to overcome the high dimensionality of the configurations space, which is $2^{M \cdot R}$, with M the number of separators and R the number of cutting rounds. Their experimental findings show a lot of promise. More research into adapting other parametric choices could provide further insights.

3.3.4 Future outlook

Measuring performance What is the purpose of the cutting routines? One is inclined to believe that cuts should strengthen the LP relaxation, hence bringing the LP bound closer to the optimal value. However, will this result in a faster solve? Turner et al.

Table 4 Summary of different learning approaches for cut management

	ML paradigm	Online/offline	Model
Tang et al. [97]	RL, evolutionary strategies	Offline	Attention + LSTM
Paulus et al. [84]	SL, imitation	Online	GNN
Wang et al. [100]	RL, policy gradient	Offline	LSTM + NN + Pointer
Turner et al. [98]	RL, policy gradient	Offline	GNN
Li et al. [71]	Contextual bandits	Offline	GNN

We use the acronym SL for supervised learning and RL for reinforcement learning

[98] experimentally measure the (kind of folklore) fact that a better root LP bound does not always translate into a shorter solving time. Other than the clear influence of the number of added cuts, many solver components can be affected by the cut choice resulting in different performance. Wang et al. [100] address this by incorporating solving time as a reward signal, instead of root LP bound. However, observing final performance comes at great computational cost, which could be prohibitive for larger instances. More research is needed on how to efficiently navigate this trade-off.

Multi-cut rounds The selection model of Tang et al. [97] and Paulus et al. [84] adds one cut at a time, resolving the LP at each iteration. Their work constitutes an important step towards learning cut selection rules. However, in a practical setting, such procedure is unlikely to outperform models that do limited LP resolving by adding multiple cuts at once. Going forward, models like the one of Wang et al. [100] or Turner et al. [98] have more potential for improvement.

3.4 Node selection

Primal heuristics have the clear goal of improving the best known MILP-objective value (primal bound). Analogously, branching rules and cutting routines are typically designed to improve the dual bound. Node selection policies have the difficult task of balancing both goals, which are often at odds. As usual, a better node selection rule is one that results in the shortest solving time. This is typically associated with smaller search trees. For that, one needs to avoid processing nodes that could be pruned if the optimal solution was known in advance. Finding a good (or even optimal) solution fast makes that task easier to accomplish.

One strategy is to first process nodes with the best (lowest) known lower bound. This is called best first search (BFS) and has the benefit of quickly improving the dual bound. Note that for all B&B trees there is a node selection policy of BFS type⁷ that minimizes the number of processed nodes [1]. However, it is well known that such an approach delays finding good (or even optimal) solutions. The depth first search (DFS) strategy prioritizes children or siblings of the node that was last processed. This approach sequentially provides more constrained sub-MIPs therefore increasing the chances of finding feasible solutions both naturally and via primal

⁷ There can be more than one BFS policy because of ties.

Table 5 Summary of different learning approaches for node selection

	ML paradigm	Online/offline	Model
He et al. [53]	SL, imitation	Offline	Linear
Yilmaz and Yorke-Smith [105]	SL, imitation	Offline	NN
Labassi et al. [70]	SL, classification	Offline	GNN
Khalil et al. [65]	SL, classification	Offline	GNN

We use the acronyms SL for supervised learning

heuristics. Furthermore, DFS has the added benefit of faster node processing times, on account of the similarity between sub-problems that are solved consecutively, which usually differ in one variable bound change. In practice, node selection rules alternate between both behaviors, while considering other statistics about branching that allow for estimating the cost of integrality.⁸

Different methodologies have been proposed for learning a node selection strategy over a homogeneous instance set (see Table 5 for an overview). He et al. [53] propose to learn from a node selection oracle that always chooses to process the node that is on a path to the optimal solution. This requires knowledge of the optimal solution during the training process, but not at test time. Similarly, Yilmaz and Yorke-Smith [105] prioritize nodes that contain high-quality solutions using a policy that always picks a child of the current node and uses an ML-based prediction to choose among these children. Once the dive is finished, they propose different ways in which the next node can be selected. Labassi et al. [70] learn a function that compares any two nodes in the tree. This can be used to substitute lower bounds as the proxy of a node's potential. The learned comparison function can be used in combination with different selection strategies, such as picking the child node with highest potential. Finally, Khalil et al. [65] guide the search based on a prediction of the optimal solution. They look at the fixed variables at each node and measure the similarity between the fixed and predicted values. This favors nodes where the partial solution aligns with the predicted solution.

3.4.1 Future outlook

The experimental results of the papers surveyed above show promise, yet the margins of improvement remain small. Often, an effective heuristic schedule and branching strategy are much more crucial and, when chosen correctly, make the impact of the node selection strategy relatively small. Still, it is interesting to observe that ML has opened new opportunities in an area where research has been pretty much inexistent for decades. This suggests that there is some potential for looking at an “old” problem from a different perspective and with new tools. For example, one could pair different node selection strategies with restarts [8], i.e., changing the node selection in a more dramatic way over time.

⁸ The literature in node selection is not extensive, good discussions can be found for example in [1, 4].

3.5 Configuration decisions

MILP solvers are highly parametric. To illustrate this, consider SCIP version 8.0.0: it has more than 2000 parameters that the user can tune. A good parameter configuration that suits the instances being solved (for example a certain class of instances) can have a crucial effect on the solving process. Again, we can look at this problem with a Machine Learning lens: we can base certain parametric decisions on a prediction given by an ML model. One can see this as falling under the realm of Automatic Algorithm Configuration (AAC). However, AAC methods typically entail configuring a large number of parameters at the same time (see, e.g. Hutter et al. [59]). Here, we are interested in the use of ML to answer a single parametric question or, at least, one question at a time. Furthermore, and contrary to other AAC methodologies, the ML models make use of a description of the instance. These models are trained on a heterogeneous collection of instances and allow for instance-dependent parameter prescriptions, as opposed to a single configuration for the given instance distribution. Some work in this area is summarized in Table 6. Notice that the aforementioned work of Turner et al. [98] on cut selection can also be seen under this light. The work of Hendel et al. [55] can be further highlighted as a dynamic parameter setting scheme. They update the pricing rule throughout the search depending on the observed runtime in the already processed nodes.

This avenue of research has already fostered considerable success. Notably, the method presented in Berthold and Hendel [18] is used by default in FICO Xpress v.8.9. The great potential of customised configurations is highlighted in problems where the preprocessing techniques have a broader impact, such as in mixed integer nonlinear programming. A prime example of this is Bonami et al. [21, 22], where the authors prescribe for each mixed integer quadratic programming instance if the quadratic objective function should be linearised or not. However, note that preprocessing has been shown to be the single most impactful component of MILP solvers [2, 5], hence the use of ML to configure the MILP algorithmic decisions based on the characteristics of an instance or a class of instances has a strong potential.

3.5.1 Future outlook

Related to the work discussed in this section is the line of research dedicated to predicting search completion (see, e.g., Fischetti et al. [41], Hendel et al. [56]) by using a range of solver statistics. This is an established field with contributions to state-of-the-art solvers, e.g., the incorporation of a search completion estimation in SCIP version 7. These predictions can be used to trigger a restart [8], a technique that has gained a lot of attention and that allows to *reconfigure* algorithmic decisions based on the evolution of the B&B. Following this line of thought, one can envision that, in the future, ML models such as the ones in Table 6 can be used not just during preprocessing, but also to prescribe a change in strategy during the solve. This is especially appealing because before a restart some (sometimes) extensive data collection has happened, data to be leveraged by ML.

Along the same lines, it is worth mentioning the recent attempts to leverage data to better solve sequences of MILP instances that differ very little one from another. This

Table 6 Five examples of using ML to set a solver parameter

	Component	Question	Options
Kruber et al. [69]	General	Should the Dantzig-Wolfe decomposition be used?	Yes / no
Hendel et al. [55]	LP solver	Which simplex pricing rule to use?	Devex / steepest / quick-start steepest
Berthold and Hendel [18]	Presolve	Which scaling method to apply?	Standard / Curtis-Reid
Berthold et al. [19]	Cutting	Should cuts be applied outside the root node?	Yes / no
Turner et al. [98]	Cutting	How should cut scores be weighted?	$\mu \in \mathbb{R}_{\geq 0}^4$

The ML model is responsible for answering a single parametric question by choosing one of the available options

has been the focus of the 2023 MIP challenge [20], and again it pertains to effectively configure an MILP solver, i.e., its algorithmic decisions, by exploiting data associated with instance similarities and data collected from previous runs. For the challenge, several classes of instances were proposed, where the instances in each class differ very little, for example, only in the coefficients of the objective function. The solver that won the competition [83] was able to leverage the data of the (previous) runs, for example, the pseudocosts for making better branching decisions.

Many open questions in this vast research area still exist, making this a promising area of future work.

4 Problem representation

The standard form of Mixed Integer Linear Programs is the one presented in Eq. 1. Given $A \in \mathbb{Q}^{m \times n}$, $\mathbf{c} \in \mathbb{Q}^n$, $\mathbf{b} \in \mathbb{Q}^m$ and the partition $(\mathcal{A}, \mathcal{B}, \mathcal{C})$ of the variables, an MILP solver can start solving the instance. The ML contributions to the solving process surveyed in Sect. 3 also require information about the problem, but the data $\{A, \mathbf{c}, \mathbf{b}, (\mathcal{A}, \mathcal{B}, \mathcal{C})\}$ may be insufficient or unfit for the task at hand. In this section we review different methodologies to construct a representation X of the problem being solved. This representation is the input to a parameterized policy $f(X, \theta)$ to be trained for a specific task.

Let us start by listing the desirable properties of a representation.

1. **Permutation invariance**⁹ permuting the order of the variables and/or constraints should leave the representation unchanged.
2. **Scale invariance:** scale invariance is preferred to keep values within controlled ranges, which helps the learning process. This can be achieved with a normalization step.
3. **Size invariance:** the size of the representation should not depend on the size of the instance. This is, we require a fixed-sized description of each element that needs representation, e.g., each variable or each node.
4. **Low computational cost:** low cost of extracting, storing and processing the data.

In the following, we will make a distinction with respect to descriptors that represent general properties of the MILP and descriptors that relate to a specific variable or constraint. These descriptors may be static in nature or they may dynamically change during the solving process. We will also discuss global descriptions of the process versus local (subproblem) properties. Note that some approaches use no description of the instance, and instead learn exclusively from the performance metric (see, e.g., Chmiela et al. [26], Hendel [54]).

⁹ Note that permutation invariance is an issue beyond the ML context. The performance of MILP solvers can be affected by a change in the order of the variables or constraints. Such seemingly irrelevant changes that have an impact on the solution process are a known issue called *performance variability*: (see Lodi and Tramontani [75]).




Fig. 5 The bipartite graph representation of an MILP

4.1 The bipartite graph representation

MILPs can be represented as a bipartite graph, as shown in Fig. 5. This graph is constructed as follows: each constraint and each variable have a corresponding representative node. A constraint node is connected to a variable node if the corresponding variable has a non-zero coefficient in the corresponding constraint. Each node has an associated vector descriptor. The advantage of using a graph representation is that this data structure can be parsed by a Graph Neural Network (GNN, see Sect. 2.1 for a formal definition). This type of architecture automatically handles inputs of different sizes. The data aggregation step that ensures size invariance is one of the learnable mappings (function `comb` in Definition 3). This is, instead of manually engineering a mechanism to aggregate information, this mechanism is automatically *learned*. The use of GNNs for combinatorial optimization has experienced a rise in popularity in recent years [24] because of their ability to capture the structural properties of the instances without excessive engineering.

4.2 Representing variables individually

Some of the learning tasks discussed in Sect. 3 require a description of each variable individually. Clearly, branching rules fall under this category. This is also the case for prediction-driven heuristics and neighborhood selection policies, for which variables are mapped to values or probabilities. In this section, we will discuss three important approaches to building variable representations and how they relate to the different approaches of Sect. 3.

A straightforward approach to building variable representations is to gather a number of descriptors into a *vector representation* for each variable. Khalil et al. [63] propose a number of such descriptors, including different statistics about the set of constraints in which each variable participates. These statistics aggregate information whose length would otherwise depend on the problem size. For example, for a variable $j \in \mathcal{I}$, using the constraint coefficients $\{a_{ij}\}_{i=1}^m$ would yield a vector whose length depends on m , which is undesirable. On the contrary, using the average of these

coefficients gives a size-independent descriptor. This is a necessary step but calls to question which statistics should be included or excluded in this feature engineering step.

Alternatively, one can use the bipartite graph representation of the MILP as described in Sect. 4.1. Gasse et al. [44] were the first to use such representation to make predictions about variables, using a GNN as a mapping. The descriptors associated to the elements in the graph include structural information, such as constraint coefficients. This data is not given in the form of aggregated statistics like before. Instead, the aggregation function is part of the learnable mappings of the GNN (see Definition 3). Apart from the structural information, both Khalil et al. [63] and Gasse et al. [44] include information about the LP solution and other basic variable features such as their objective coefficient or variable type. Zarpellon et al. [106] take a different perspective, stressing the importance of historical data collected during the B&B tree. This strategy resembles SCIP's default branching rule, which considers information about past branchings, collected conflicts or cutoffs. The representation used in Zarpellon et al. [106] includes this variable information and, additionally, global information about the search tree. They argue that such a description can uncover shared structures among very diverse MILPs (see Sect. 3.2). The information collected includes statistics about the node being processed, tree composition and shape, and bound statistics, with a particular focus on unprocessed nodes.

These three approaches to MILP variable representation are compared in Table 7. We report on features that describe variables individually, therefore excluding the complete list of tree features of Zarpellon et al. [106] (which can be found in their appendix). Table 7 showcases that, while some common features exist, the different representations have distinct focuses.

Most of the work surveyed in Sect. 3 uses the graph representation of Gasse et al. [44], either exactly (e.g., [51, 90, 96]) or with small modifications of the variable descriptors (e.g., [65, 74, 104]). Hybrid models also exist. In particular, Gupta et al. [50] propose extracting the Gasse et al. [44] representation at the root node and the Khalil et al. [63] representation in the rest of the nodes of the search tree. The reason is that, while the graph representation is rich, it is also computationally expensive. Gasse et al. [44] overcome this by using a GPU (graphics processing unit) to accelerate the execution of their learned function. Such computation on GPUs is common practice in the ML community, but one could argue that it is unrealistic to require the availability of a GPU. The hybrid model of Gupta et al. [50] reuses the rich but expensive representation of the root node in combination with the features of Khalil et al. [63] that are cheap and update the description at every node. This proves very effective, with their best performing model outperforming both reliability branching and the model of Gasse et al. [44] when executed without GPU acceleration.

4.3 Representing constraints individually

Analogously to variables, a description of the problem's constraints may be needed. Here, we refer both to original problem constraints and additional valid constraints that can be added as cuts. A constraint representation is necessary in two cases. First

Table 7 Three key approaches to variable descriptors

Basic	Khalil et al.	Gasse et al.	Zarpellon et al.
Objective coefficient	✓	✓	
Upper/lower bound		✓	
Type		✓	
Structural			
# of constraints the variable is in	✓	Implied	
Min and max ratios a_{ij} to b_i	✓	Implied	
Min and max ratios a_{ij} to $\sum_k a_{kj}$	✓	Implied	
Stats for constraint degrees	✓	Implied	
Stats for constraint coefficients	✓	Implied	
Stats for active constraints	✓	Implied	
LP solution			
LP basis status		✓	
LP sol value		✓	✓
LP sol at bound		✓	
LP sol fractionality	✓	✓	
LP sol reduced cost		✓	
LP sol age		✓	
Root LP sol value		✓	
Incumbent			
Incumbent value		✓	
Average incumbent value		✓	✓
Tree statistics			
Average branching depth			✓
Conflict score			✓
Conflict length score			✓
Average conflict length			✓
Pseudocost score	✓		✓
Pseudocost count stats			✓
Average inference score			✓
Average number of inferences			✓
Average cutoff score			✓
Average cutoffs of variable	✓		✓
# of implications derived			✓
# of cliques of variable			✓

Features marked as 'implied' do not appear explicitly as a variable descriptor, but can be inferred from the bipartite graph representation

and undoubtedly, whenever the task requires a decision over said constraints (e.g., which cuts to add). Second, the bipartite graph representation of MILPs discussed in the previous section (see Fig. 5) also calls for a description of the constraints, even in the case where they are to be aggregated at a later stage.

Gasse et al. [44] first proposed the bipartite graph representation using a small number of descriptors for the constraint nodes. In particular, they use the cosine similarity¹⁰ with the objective coefficients, the constraint right-hand-side, and LP information such as basis status and dual bound. This type of concise descriptions of the constraints is frequently used for the learning tasks associated with branching or primal heuristics.

In the case of cut selection, a more detailed description is preferred. Paulus et al. [84] extend the graph representation of Gasse et al. [44] with metrics that are typically considered in cut selection, such as violation, objective parallelism or sparsity. Wang et al. [100] describe each cut with a single vector of classical cut scores (see, e.g., [101]). In the case of Turner et al. [98], the classical cut scores are intrinsically taken into account in the definition of their learning task (see Sect. 3.3). For this reason, they use a graph representation of the model with a small amount of variable and constraint descriptors.

4.4 Representing a (sub-)MILP

To conclude, it is worth observing that some decision tasks are formulated at an instance or node level, therefore needing a global representation of the MILP and perhaps also the solving process. One possible approach is to aggregate variable and constraint descriptors coming from the bipartite graph representation. Liu et al. [74] use the average of the variable descriptors, while Labassi et al. [70] concatenate the average variable descriptors and the average constraint descriptors.

Other approaches include the one of Khalil et al. [64], who, in the context of scheduling of primal heuristics, build a vector representation of the current node. This representation includes comparisons to the root node and context on the node's position within the tree. Great focus is put on information coming from the LP solution, such as the objective value, average fractionalities, and statistics on the constraint activity. Berthold et al. [18, 19] build representations that are more specialized to the particular configuration task. We refer to their work for a more detailed discussion of the problem descriptors.

4.5 Outlook

In this section, we have reviewed different methodologies that build a suitable representation of MILPs that can be parsed by learning models. This includes both the selection of relevant data and data structures. At the beginning of the section, we anticipated the desired properties of such representations. All of the presented methodologies ensure permutation invariance and, to some degree, try to normalize data to establish some scale invariance without excluding important information. We also dis-

¹⁰ The cosine similarity between two vectors $\mathbf{a}$ and $\mathbf{b}$ of the same length is defined as $\mathbf{a} \cdot \mathbf{b} / \|\mathbf{a}\|_2 \|\mathbf{b}\|_2$, that is, their dot product divided by the product of their norms.

cussed the importance of the graph representation to handle input of different sizes. In the following, we elaborate on the importance of Graph Neural Networks for MILP representation, as well as the determining trade-off between model expressivity and speed.

Expressivity versus speed In the context of machine learning, model size refers to the number of parameters and operations that define the mapping function $f(\cdot, \theta)$. Larger models allow us to learn more complex relationships between input and output. However, there is a clear trade-off between model size and computational cost of execution and training. In stark contrast to other fields of application of ML, such as large language models [81] or strategic game playing [99] where the goal is super-human performance, the computational cost per execution of the ML model is decisive to whether or not it will beat its competitors, i.e., already highly efficient optimization software. It is desirable for the MILP representation to ensure low computational cost, both in terms of data extraction and processing. The work discussed in Sect. 3 shows that navigating this trade-off is an active field of research, with promising results in both small (e.g., [63]) and large models (e.g., [44]), as well as in compressing ML models without compromising the accuracy of their predictions (e.g., [50]).

The case for graph neural networks GNNs offer a powerful representation tool for MILP. They enable instance parsing with less feature engineering, as well as size and permutation invariance. Computationally, they have shown excellent performance across different tasks. Yet, we have limited understanding of the reasons behind this success. Some recent studies have uncovered some of the factors that contribute to the success of GNNs.

Chen et al. [25] study the separation and representation power of GNNs for LP. In particular, they study this in the context of three prediction tasks: predicting feasibility, boundedness, and the optimal solution vector for LP. The separation power is the ability of GNNs to distinguish different instances, i.e., their ability to output different results when given two different instances as input. Chen et al. [25] prove that given two LPs, if no GNN¹¹ can distinguish them, then both LPs have the same status in terms of feasibility and boundedness. Furthermore, they both have the same minimum- l_2 -norm optimal solution up to a permutation. Finally, they also show that the three tasks mentioned above can in fact be approximated using GNNs. Continuing this line of work, Qian et al. [88] prove that GNNs can be used to reproduce interior point methods. In particular, they show that there exists a GNN using $\mathcal{O}(m)$ message-passing operations (see Definition 3), with m the number of constraints, that can replicate any one iteration of the algorithm by Nocedal and Wright [80]. They show the same result for the more practical algorithm by Gondzio [47]. Notably, this is true when representing the LP using a modified version of the aforementioned bipartite graph representation (see Fig. 5), where a new node is added and connected to all variable and constraint nodes. This global node adds alternative routes of communication among constraint and variable nodes and is said to represent the objective function. Qian et al.

¹¹ The authors consider the family of functions defined in Eq. 8, where the combination and aggregation functions are feed-forward neural networks.

Table 8 A list of some of the most relevant instance collections

Benchmark	Composition	Size	Source	URL
MIPLIB 2017	Mixed	1065	[46]	↗
Cor@1	Mixed	364	–	↗
NN verification	Homogeneous	3692	[78]	↗
ML4CO_1	Homogeneous	10,000	[45]	↗
ML4CO_2	Homogeneous	10,000	[45]	↗
ML4CO_3	Homogeneous	118	[45]	↗

[88] also provide a computational comparison of different GNN implementations, i.e., different `comb` and `aggr` functions (see Eq. 8). This is also an active area of research, with recent works advocating for the so-called graph attention networks (e.g., [72, 93]) and other sophisticated architectures and training methods.

These results shed some light on the representation power of GNNs for MILP and strengthen the case for using them in the context of optimization problems.

5 Datasets and software

5.1 Datasets

The modelling power of MILP makes it a suitable language for a large range of applications. With the goal of measuring the performance of different algorithms, the MILP research community has curated large benchmarks, such as MIPLIB [46], that provide a heterogeneous set of instances coming from diverse applications. It should be noted that these benchmarks are considered large for MILP standards, but are orders of magnitude smaller than typical ML benchmarks. In the light of the methods surveyed in Sect. 3, which combine an ML component with classical optimization, there is a renewed need for collections of MILP instances. In this section we provide an overview of the collections that have been used in the growing body of literature. We restrict our discussion to benchmarks that are publicly available or whose generation code is easily accessible.

ML methodologies usually require vast amounts of data. For this reason, it is common to resort to instance generators, which complement the existing instance collections. Tables 8 and 9 provide a summary of both commonly used instance collections (with their size specification) and instance generators.

Apart from the size, there is the consideration of the composition of instances. When implementing a learning-augmented solver component a specification needs to be made regarding the instances of interest. In machine learning terms, we typically talk of an *instance distribution*, where the instances that are outside the scope of interest are assigned a zero probability of occurring. For some applications, it can be assumed that the representative instances have a shared combinatorial structure. The machine learning model is then expected to specialize to this structure. Con-

Table 9 A list of some of the most relevant instance generators

Benchmark	Problem type(s)	Source
Tang et al.	Max-cut	[97]
	Planning	
	Packing	
	Bin packing	
Ecole	Set cover	[87]
	Combinatorial auctions	
	Maximum independent set	
	Capacitated facility location	
MIPLearn	Bin packing	[89]
	Multi-dimensional knapsack	
	Capacitated p-median	
	Set cover	
	Set packing	
	Stable set	
	Traveling salesman	
	Unit commitment	
	Vertex cover	
GISP	Generalized independent set	[28]
FCMNF	Capacitated fixed-charge network flow	[57]

versely, some approaches are designed to detect patterns across instances of any class. Throughout Sect. 3 we have surveyed examples of both situations.

Table 10 summarizes which of the discussed approaches uses a collection of instances with mixed structures (mixed), and which use the assumption that all instances belong to the same class (homogeneous). From this we can observe that configuration decisions are more naturally framed over mixed instance collections than other tasks, like the more complex matter of branching, where some instance specialization seems valuable. Table 11 further shows an overview of the homogeneous datasets used in the work presented in Sect. 3. We can observe a pattern that highlights differences in the instances, based on which task is more challenging. Instances like GISP or FCMNF are more commonly chosen as a challenging test bed for primal heuristics, indicating that for these problems the difficulty lies in finding (optimal) solutions. On the contrary, branching is usually tested on instances where proving optimality is the main challenge, such as the ones provided by Ecole.

5.2 Software

In connection to instance generators, there has been increasing interest in developing libraries that help the process of data generation, training and testing in the context of ML-augmented MILP solving. Some examples of these are the library Ecole [\[87\]](#), or

Table 10 Classification of the literature based on whether they use a mixed or an homogeneous instance collection (or both)

	Mixed	Homogeneous
<i>Primal heuristics</i>		
Ding et al. [34]		✓
Nair et al. [78]	✓	✓
Khalil et al. [65]		✓
Song et al. [95]		✓
Wu et al. [104]		✓
Sonnerat et al. [96]	✓	✓
Liu et al. [74]	✓	✓
Huang et al. [58]		✓
Khalil et al. [64]	✓	✓
Chmiela et al. [26]		✓
Hendel [54]	✓	
<i>Branching</i>		
Khalil et al. [63]	✓	
Gasse et al. [44]		✓
Gupta et al. [50]		✓
Etheve et al. [37]		✓
Nair et al. [78]	✓	✓
Zarpellon et al. [106]	✓	
Gupta et al. [51]		✓
Scavuzzo et al. [90]		✓
<i>Cut selection</i>		
Tang et al. [97]		✓
Paulus et al. [84]		✓
Wang et al. [100]	✓	✓
Turner et al. [98]	✓	
Li et al. [71]	✓	✓
<i>Node selection</i>		
He et al. [53]		✓
Yilmaz and Yorke-Smith [105]		✓
Labassi et al. [70]		✓
<i>Configuration decisions</i>		
Kruber et al. [69]	✓	
Hendel et al. [55]	✓	
Berthold and Hendel [18]	✓	
Berthold et al. [19]	✓	
Turner et al. [98]	✓	

Table 11 Common homogeneous instance collections and where they are used

	Ecole	NNv	GISP	FCMNF	Tang et al	Other
<i>Primal heuristics</i>						
Ding et al. [34]						✓
Nair et al. [78]		✓				✓
Khalil et al. [65]			✓	✓		
Song et al. [95]						✓
Wu et al. [104]						✓
Sonnerat et al. [96]		✓				✓
Liu et al. [74]	✓		✓			
Huang et al. [58]	✓					✓
Khalil et al. [64]			✓			
Chmiela et al. [26]			✓			
<i>Branching</i>						
Gasse et al. [44]	✓					
Gupta et al. [50]	✓					
Etheve et al. [37]						✓
Nair et al. [78]		✓				✓
Gupta et al. [51]	✓					
Scavuzzo et al. [90]	✓					
<i>Cut selection</i>						
Tang et al. [97]					✓	
Paulus et al. [84]		✓			✓	
Wang et al. [100]	✓					✓
Li et al. [71]	✓	✓			✓	✓
<i>Node selection</i>						
He et al. [53]						✓
Yilmaz and Yorke-Smith [105]	✓					
Labassi et al. [70]			✓	✓		✓

the more recent MIPLearn [89]. Their goal is to provide a standardized platform for the research community for fast prototyping and testing by removing the barrier of challenging software implementation. These libraries provide ways to easily implement learning tasks, such as branching or warm-starting. For an up-to-date specification of the features they provide we refer to their documentation.

We end the section by remarking that developing software of this type is far from trivial and present some structural challenges. Depending on the problem representation (see Sect. 4), one needs to provide an increasingly tight interface with the MILP solver in order to collect the required data. This is both in the case of offline and (even more so) online learning. For example, in the case of Ecole, the development in strict correlation with SCIP has been instrumental. Clearly, in the case one wants to use a

commercial solver, a sufficiently tight interface and integration might be impossible to reach, thus limiting the type of methods one can implement. We extend the discussion on the challenges in the following section.

6 Conclusions, perspective and challenges

The work covered in this article testifies to the growing interest in the integration of ML methodologies within MILP solvers. This is an emerging technology that has already fostered remarkable success within its short history,¹² and is likely to play a key role in future algorithmic developments. Beyond the discussion of the literature, we have highlighted some methodological trends and characterised the common grounds with respect to instance representation, learning algorithms and benchmarking.

Meaningful steps forward have been taken in answering the more pressing research questions. For example, the literature shows that some learning tasks seem to be formulated more naturally than others over heterogeneous instance collections. In other cases, an argument can be made in favour of the applicability of specializing to a certain structure, which makes the learning task easier. Studies like the ones in Zarpellon et al. [106] and Fischetti et al. [41] indicate that instance representations that describe the global solution process allow to more easily recognize patterns across different combinatorial structures. This is especially promising because a key challenge for the integration of ML-augmented methods into MILP solvers is their generalization properties, at least as long as solvers are conceived as one-configuration-fits-all software.

It is also interesting to note the various efforts to define efficient success metrics for the different learning tasks, such as imitation targets or reward functions. Solving instances to optimality is to be avoided because of the computational effort, but performance proxies that substitute solving time must be carefully chosen.

Already substantial progress has also been made in creating the right environment for easily implementing and testing ML models for MILP solving. Existing software infrastructure includes, for example, curated instance generators, code that simplifies the solver interface and standardized testing procedures. This further helps in evaluating and comparing the different methodologies. Other efforts to bring the research community together are competitions, like ML4CO [45], which encourage progress in well-defined tasks as well as fair comparisons among the proposed methods. A significant challenge resides on the software versus hardware side: many learning methods, e.g., those relying on neural networks, especially benefit from the use of GPUs, while MILP technology is inherently CPU based. The CPU versus GPU interaction is currently a relevant obstacle for ML-augmented MILP.

Overall, a key trend seems to be building more dynamic solving strategies. MILP solvers generate a plethora of statistics during execution that often go unused. This is fertile ground for learning algorithms, which can unlock more dynamic solvers that automatically adapt the solving strategy based on prescriptions derived from such solving statistics. This poses exciting new questions and challenges.

¹² This is the case of methods that have been included in commercial (e.g., [18, 22]) or non-commercial (e.g., [8, 56]) MILP solvers.

Finally, we remark that the phase the field is entering now is that of a more systematic transfer of the proofs of concept discussed in this survey to the MILP software. The successful transfers reviewed in this article—and recalled at the beginning of this section—give solid evidence but do not define yet a mature path for such a transfer to happen because of the challenges above. This is the next, in our opinion achievable, step to make.

Acknowledgements The authors are indebted to two anonymous referees for helpful comments on the original version of the article. The third author would like to warmly thank the Government of Canada for establishing the Canada Excellence Research Chair (CERC) program. It generously supported the “Data Science for Real-time Decision-making” CERC at Polytechnique Montréal that the author had the privilege to lead (2015–2022) and whose fantastic team has been instrumental to shape his knowledge of the topic.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article’s Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article’s Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

1. Achterberg, T.: Constraint integer programming. Ph.D. thesis, Technischen Universität Berlin (2007)
2. Achterberg, T., Wunderling, R.: Mixed integer programming: analyzing 12 years of progress. In: *Facets of Combinatorial Optimization: Festschrift for Martin Grötschel*, Springer, pp. 449–481 (2013)
3. Achterberg, T., Koch, T., Martin, A.: Branching rules revisited. *Oper. Res. Lett.* **33**(1), 42–54 (2005)
4. Achterberg, T., Berthold, T., Heinz, S., Koch, T., Wolter, K.: *Constraint Integer Programming: Techniques and Applications*. ZIB-Report 08-43, Zuse Institute Berlin (2008)
5. Achterberg, T., Bixby, R.E., Gu, Z., Rothberg, E., Weninger, D.: Presolve reductions in mixed integer programming. *INFORMS J. Comput.* **32**(2), 473–506 (2020)
6. Alvarez, A.M., Wehenkel, L., Louveaux, Q.: Online learning for strong branching approximation in branch-and-bound, <https://orbi.uliege.be/handle/2268/192361>, working paper (2016)
7. Alvarez, A.M., Louveaux, Q., Wehenkel, L.: A machine learning-based approximation of strong branching. *INFORMS J. Comput.* **29**(1), 185–195 (2017)
8. Anderson, D., Hendel, G., Le Bodic, P., Viernickel, M.: Clairvoyant restarts in branch-and-bound search using online tree-size estimation. *Proc. AAAI Conf. Artif. Intell.* **33**, 1427–1434 (2019)
9. Andreello, G., Caprara, A., Fischetti, M.: Embedding cuts in a branch and cut framework: a computational study with $\{0, \frac{1}{2}\}$ -cuts. *INFORMS J. Comput.* **19**, 229–238 (2007)
10. Applegate, D., Bixby, R., Chvátal, V., Cook, W.: Finding cuts in the TSP (a preliminary report). Technical Report 95-05, DIMACS (1995). <https://api.semanticscholar.org/CorpusID:972108>
11. Balas, E., Ceria, S., Cornuéjols, G., Natraj, N.: Gomory cuts revisited. *Oper. Res. Lett.* **19**(1), 1–9 (1996)
12. Balcan, M.F., Dick, T., Sandholm, T., Vitercik, E.: Learning to branch. In: *International Conference on Machine Learning*, PMLR, pp. 344–353 (2018)
13. Balcan, M.F.F., Prasad, S., Sandholm, T., Vitercik, E.: Sample complexity of tree search configuration: cutting planes and beyond. *Adv. Neural. Inf. Process. Syst.* **34**, 4015–4027 (2021)
14. Barnhart, C., Laporte, G.: (eds) *Transportation, Handbooks in Operations Research and Management Science*, vol 14. North-Holland (2007)
15. Bengio, Y., Lodi, A., Prouvost, A.: Machine learning for combinatorial optimization: a methodological tour d’horizon. *Eur. J. Oper. Res.* **290**(2), 405–421 (2021)
16. Berthold, T.: Measuring the impact of primal heuristics. *Oper. Res. Lett.* **41**(6), 611–614 (2013)

17. Berthold, T., Csizmadia, Z.: The confined primal integral: a measure to benchmark heuristic MINLP solvers against global MINLP solvers. *Math. Program.* **188**(2), 523–537 (2021)
18. Berthold, T., Hendel, G.: Learning to scale mixed-integer programs. *Proc. AAAI Conf. Artif. Intell.* **35**, 3661–3668 (2021)
19. Berthold, T., Francobaldi, M., Hendel, G.: Learning to use local cuts. (2022). *arXiv preprint [arXiv:2206.11618](https://arxiv.org/abs/2206.11618)*
20. Bolusani, S., Besançon, M., Gleixner, A., Berthold, T., d'Ambrosio, C., Muñoz, G., Paat, J., Thomopoulos, D.: The MIP Workshop 2023 Computational Competition on Reoptimization. (2023). *arXiv preprint [arXiv:2311.14834](https://arxiv.org/abs/2311.14834)*
21. Bonami, P., Lodi, A., Zarpellon, G.: Learning a classification of mixed-integer quadratic programming problems. In: *Proceedings of the International Conference on Integration of Constraint Programming, Artificial Intelligence, and Operations Research (CPAIOR)*, Springer, vol. 15, pp. 595–604 (2018)
22. Bonami, P., Lodi, A., Zarpellon, G.: A classifier to decide on the linearization of mixed-integer quadratic problems in CPLEX. *Oper. Res.* **70**(6), 3303–3320 (2022)
23. Bubeck, S., Cesa-Bianchi, N., et al.: Regret analysis of stochastic and nonstochastic multi-armed bandit problems. *Found. Trends Mach. Learn.* **5**(1), 1–122 (2012)
24. Cappart, Q., Chételat, D., Khalil, E.B., Lodi, A., Morris, C., Veličković, P.: Combinatorial optimization and reasoning with graph neural networks. *J. Mach. Learn. Res.* **24**(130), 1–61 (2023)
25. Chen, Z., Liu, J., Wang, X., Yin, W.: On representing linear programs by graph neural networks. In: *Proceedings of the International Conference on Learning Representations*, vol. 11 (2023)
26. Chmiela, A., Khalil, E., Gleixner, A., Lodi, A., Pokutta, S.: Learning to schedule heuristics in branch and bound. *Adv. Neural. Inf. Process. Syst.* **34**, 24235–24246 (2021)
27. Chmiela, A., Gleixner, A., Lichocki, P., Pokutta, S.: Online learning for scheduling MIP heuristics. In: *Proceedings of the International Conference on Integration of Constraint Programming, Artificial Intelligence, and Operations Research (CPAIOR)*, Springer, pp. 114–123 (2023)
28. Colombi, M., Mansini, R., Savelsbergh, M.: The generalized independent set problem: polyhedral analysis and solution approaches. *Eur. J. Oper. Res.* **260**(1), 41–55 (2017)
29. Conforti, M., Cornuéjols, G., Zambelli, G.: *Integer Programming Models*, Graduate Texts in Mathematics, vol. 271, Springer, pp. 45–84 (2014)
30. Cornuéjols, G., et al.: Revival of the Gomory cuts in the 1990's. *Ann. Oper. Res.* **149**(1), 63–66 (2007)
31. Dey, S.S., Molinaro, M.: Theoretical challenges towards cutting-plane selection. *Math. Program.* **170**, 237–266 (2018)
32. Deza, A., Khalil, E.B.: Machine learning for cutting planes in integer programming: a survey. In: *International Joint Conference on Artificial Intelligence*, ijcai.org, vol. 32, pp. 6592–6600 (2023)
33. Dilkina, B., Gomes, C.P., Malitsky, Y., Sabharwal, A., Sellmann, M.: Backdoors to combinatorial optimization: feasibility and optimality. In: *Proceedings of the International Conference on Integration of Constraint Programming, Artificial Intelligence, and Operations Research (CPAIOR)*, Springer, vol. 6, pp. 56–70 (2009)
34. 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. *Proc. AAAI Conf. Artif. Intell.* **34**, 1452–1459 (2020)
35. Duchi, J., Hazan, E., Singer, Y.: Adaptive subgradient methods for online learning and stochastic optimization. *J. Mach. Learn. Res.* **12**(7) (2011)
36. Dulac-Arnold, G., Levine, N., Mankowitz, D.J., Li, J., Paduraru, C., Goyal, S., Hester, T.: Challenges of real-world reinforcement learning: definitions, benchmarks and analysis. *Mach. Learn.* **110**(9), 2419–2468 (2021)
37. Etheve, M., Alès, Z., Bissuel, C., Juan, O., Kedad-Sidhoum, S.: Reinforcement learning for variable selection in a branch and bound algorithm. In: *Proceedings of the International Conference on Integration of Constraint Programming, Artificial Intelligence, and Operations Research (CPAIOR)*, Springer, vol. 17, pp. 176–185 (2020)
38. FICO: FICO Xpress Optimizer (2023). <https://www.fico.com/en/products/fico-xpress-optimization>
39. Fischetti, M., Lodi, A.: Local branching. *Math. Program.* **98**, 23–47 (2003)
40. Fischetti, M., Lodi, A.: *Heuristics in Mixed Integer Programming*. Wiley Encyclopedia of Operations Research and Management Science (2010)
41. Fischetti, M., Lodi, A., Zarpellon, G.: Learning MILP resolution outcomes before reaching time-limit. In: *Proceedings of the International Conference on Integration of Constraint Programming, Artificial Intelligence, and Operations Research (CPAIOR)*, Springer, pp. 275–291 (2019)

42. Gamrath, G., Schubert, C.: Measuring the impact of branching rules for mixed-integer programming. In: Operations Research Proceedings 2017: Selected Papers of the Annual International Conference of the German Operations Research Society (GOR), Freie Universität Berlin, Germany, September 6–8, 2017, Springer, pp. 165–170 (2018)
43. Gamrath, G., Anderson, D., Bestuzheva, K., Chen, W.K., Eifler, L., Gasse, M., Gemander, P., Gleixner, A., Gottwald, L., Halbig, K., et al.: The SCIP Optimization Suite 7.0 (2020)
44. Gasse, M., Chételat, D., Ferroni, N., Charlin, L., Lodi, A.: Exact combinatorial optimization with graph convolutional neural networks. *Adv. Neural Inf. Process. Syst.* **32** (2019)
45. Gasse, M., Bowly, S., Cappart, Q., Charfreitag, J., Charlin, L., Chételat, D., Chmiela, A., Dumouchelle, J., Gleixner, A., Kazachkov, A.M. et al.: The machine learning for combinatorial optimization competition (ML4CO): results and insights. In: NeurIPS 2021 Competitions and Demonstrations Track, PMLR, pp. 220–231 (2022)
46. Gleixner, A., Hendel, G., Gamrath, G., Achterberg, T., Bastubbe, M., Berthold, T., Christophel, P., Jäck, K., Koch, T., Linderoth, J., et al.: MIPLIB 2017: data-driven compilation of the 6th mixed-integer programming library. *Math. Program. Comput.* **13**(3), 443–490 (2021)
47. Gondzio, J.: Interior point methods 25 years later. *Eur. J. Oper. Res.* **218**(3), 587–601 (2012)
48. Goodfellow, I., Bengio, Y., Courville, A.: Deep Learning. MIT Press, Cambridge (2016)
49. Greenacre, M., Groenen, P.J., Hastie, T., d’Enza, A.I., Markos, A., Tuzhilina, E.: Principal component analysis. *Nature Rev. Methods Primers* **2**(1), 100 (2022)
50. Gupta, P., Gasse, M., Khalil, E., Mudigonda, P., Lodi, A., Bengio, Y.: Hybrid models for learning to branch. *Adv. Neural. Inf. Process. Syst.* **33**, 18087–18097 (2020)
51. Gupta, P., Khalil, E.B., Chételat, D., Gasse, M., Lodi, A., Bengio, Y., Kumar, M.P.: Lookback for learning to branch. *Trans. Mach. Learn. Res.* <https://openreview.net/forum?id=EQpGkw5rvL>, expert Certification (2022)
52. Gurobi Optimization, LLC: Gurobi Optimizer Reference Manual (2023). <https://www.gurobi.com>
53. He, H., Daume III, H., Eisner, J.M.: Learning to search in branch and bound algorithms. *Adv. Neural Inf. Process. Syst.* **27** (2014)
54. Hendel, G.: Adaptive large neighborhood search for mixed integer programming. *Math. Program. Comput.* **14**(2), 185–221 (2022)
55. Hendel, G., Miltenberger, M., Witzig, J.: Adaptive algorithmic behavior for solving mixed integer programs using bandit algorithms. In: Operations Research Proceedings 2018, Springer, Berlin pp. 513–519 (2019)
56. Hendel, G., Anderson, D., Le Bodic, P., Pfetsch, M.E.: Estimating the size of branch-and-bound trees. *INFORMS J. Comput.* **34**(2), 934–952 (2022)
57. Hewitt, M., Nemhauser, G.L., Savelsbergh, M.W.: Combining exact and heuristic approaches for the capacitated fixed-charge network flow problem. *INFORMS J. Comput.* **22**(2), 314–325 (2010)
58. Huang, T., Ferber, A.M., Tian, Y., Dilkina, B., Steiner, B.: Searching large neighborhoods for integer linear programs with contrastive learning. In: International Conference on Machine Learning, PMLR, pp. 13869–13890 (2023)
59. Hutter, F., Hoos, H.H., Leyton-Brown, K., Stützle, T.: ParamILS: an automatic algorithm configuration framework. *J. Artif. Intell. Res.* **36**, 267–306 (2009)
60. IBM (2023) IBM ILOG CPLEX Optimizer. <https://www.ibm.com/products/ilog-cplex-optimization-studio/cplex-optimizer>
61. Karamanov, M., Cornuéjols, G.: Branching on general disjunctions. *Math. Program.* **128**, 403–436 (2011)
62. Khalil, E.: Towards tighter integration of machine learning and discrete optimization. Ph.D. thesis, Georgia Institute of Technology (2019)
63. Khalil, E., Le Bodic, P., Song, L., Nemhauser, G., Dilkina, B.: Learning to branch in mixed integer programming. In: Proceedings of the AAAI Conference on Artificial Intelligence, vol. **30**(1) (2016)
64. Khalil, E.B., Dilkina, B., Nemhauser, G.L., Ahmed, S., Shao, Y.: Learning to run heuristics in tree search. In: International Joint Conference on Artificial Intelligence, pp. 659–666 (2017)
65. Khalil, E.B., Morris, C., Lodi, A.: MIP-GNN: A data-driven framework for guiding combinatorial solvers. *AAAI* (2022)
66. Khalil, E.B., Vaezipour, P., Dilkina, B.: Finding backdoors to integer programs: a monte carlo tree search framework. *Proc. AAAI Conf. Artif. Intell.* **36**, 3786–3795 (2022)
67. Kingma, D.P., Ba, J.: Adam: a method for stochastic optimization (2014). arXiv preprint [arXiv:1412.6980](https://arxiv.org/abs/1412.6980)

68. Kirk, R., Zhang, A., Grefenstette, E., Rocktäschel, T.: A survey of zero-shot generalisation in deep reinforcement learning. *J. Artif. Intell. Res.* **76**, 201–264 (2023)
69. Kruber, M., Lübbecke, M.E., Parmentier, A.: Learning when to use a decomposition. In: *Proceedings of the International Conference on Integration of Constraint Programming, Artificial Intelligence, and Operations Research (CPAIOR)*, Springer, vol. 14, pp. 202–210 (2017)
70. Labassi, A.G., Chételat, D., Lodi, A.: Learning to compare nodes in branch and bound with graph neural networks. *Adv. Neural Inf. Process. Syst.* (2022)
71. Li, S., Ouyang, W., Paulus, M., Wu, C.: Learning to configure separators in branch-and-cut. *Adv. Neural Inf. Process. Syst.* **36**, 60021–60034 (2023)
72. Lin, J., Zhu, J., Wang, H., Zhang, T.: Learning to branch with tree-aware branching transformers. *Knowl.-Based Syst.* **252**, 109455 (2022)
73. Linderoth, J.T., Savelsbergh, M.W.P.: A computational study of search strategies for mixed integer programming. *INFORMS J. Comput.* **11**, 173–187 (1999)
74. Liu, D., Fischetti, M., Lodi, A.: Learning to search in local branching. *Proc. AAAI Conf. Artif. Intell.* **36**, 3796–3803 (2022)
75. Lodi, A., Tramontani, A.: Performance variability in mixed-integer programming. In: *Theory Driven by Influential Applications*, INFORMS, pp. 1–12 (2013)
76. MIPLIB: MIPLIB 2017—The Mixed Integer Programming Library. <https://miplib.zib.de/> (2017)
77. Mitchell, T.M.: *Machine Learning*. McGraw Hill, New York (2017)
78. Nair, V., Bartunov, S., Gimeno, F., von Glehn, I., Lichocki, P., Lobov, I., O'Donoghue, B., Sonnerat, N., Tjandraatmadja, C., Wang, P. et al.: Solving mixed integer programs using neural networks (2020). arXiv preprint [arXiv:2012.13349](https://arxiv.org/abs/2012.13349)
79. Nemhauser, G.L., Wolsey, L.A.: *Integer and Combinatorial Optimization*. Wiley Interscience Series in Discrete Mathematics and Optimization. Wiley, London (1999)
80. Nocedal, J., Wright, S.J.: *Linear programming: Interior-point methods*. In: *Numerical Optimization*, Springer, chap 14 (2006)
81. OpenAI: Introducing ChatGPT (2022). <https://openai.com/blog/chatgpt>, [Online; accessed 04-April-2023]
82. Padberg, M., Rinaldi, G.: A branch-and-cut algorithm for the resolution of large-scale symmetric traveling salesman problems. *SIAM Rev.* **33**(1), 60–100 (1991)
83. Patel, K.K.: Progressively strengthening and tuning MIP solvers for reoptimization. (2023). arXiv preprint [arXiv:2308.08986](https://arxiv.org/abs/2308.08986)
84. Paulus, M.B., Zarpellon, G., Krause, A., Charlin, L., Maddison, C.: Learning to cut by looking ahead: cutting plane selection via imitation learning. In: *International Conference on Machine Learning*, PMLR, pp. 17584–17600 (2022)
85. Pochet, Y., Wolsey, L.A.: *Production Planning by Mixed Integer Programming*. Springer Series in Operations Research and Financial Engineering, Springer (2006)
86. Pomerleau, D.A.: Efficient training of artificial neural networks for autonomous navigation. *Neural Comput.* **3**(1), 88–97 (1991)
87. Prouvost, A., Dumouchelle, J., Scavuzzo, L., Gasse, M., Chételat, D., Lodi, A.: Ecole: A gym-like library for machine learning in combinatorial optimization solvers (2020). arXiv preprint [arXiv:2011.06069](https://arxiv.org/abs/2011.06069)
88. Qian, C., Chételat, D., Morris, C.: Exploring the power of graph neural networks in solving linear optimization problems (2023). arXiv preprint [arXiv:2310.10603](https://arxiv.org/abs/2310.10603)
89. Santos Xavier, A., Qiu, F., Gu, X., Becu, B., Dey, S.S.: MIPLearn: an extensible framework for learning-enhanced optimization (2023). <https://doi.org/10.5281/zenodo.8018181>
90. Scavuzzo, L., Chen, F.Y., Chételat, D., Gasse, M., Lodi, A., Yorke-Smith, N., Aardal, K.: Learning to branch with tree MDPs. *Adv. Neural Inf. Process. Syst.* (2022)
91. Schrijver, A.: *Theory of Linear and Integer Programming*. Wiley Interscience Series in Discrete Mathematics and Optimization. Wiley, London (1998)
92. Sewak, M.: *Deep Reinforcement Learning*. Springer, Berlin (2019)
93. Seyfi, M., Banitalebi-Dehkordi, A., Zhou, Z., Zhang, Y.: Exact combinatorial optimization with temporo-attentional graph neural networks. In: *Proceedings of the Joint European Conference on Machine Learning and Knowledge Discovery in Databases (ECML PKDD)*, Springer, pp. 268–283 (2023)

94. Shaw, P.: Using constraint programming and local search methods to solve vehicle routing problems. In: International Conference on Principles and Practice of Constraint Programming, Springer, pp. 417–431 (1998)
95. Song, J., Lanka, R., Yue, Y., Dilkina, B.: A general large neighborhood search framework for solving integer linear programs. *Adv. Neural. Inf. Process. Syst.* **33**, 20012–20023 (2020)
96. Sonnerat, N., Wang, P., Ktena, I., Bartunov, S., Nair, V.: Learning a large neighborhood search algorithm for mixed integer programs (2021). arXiv preprint [arXiv:2107.10201](https://arxiv.org/abs/2107.10201)
97. Tang, Y., Agrawal, S., Faenza, Y.: Reinforcement learning for integer programming: learning to cut. In: International Conference on Machine Learning, PMLR, pp. 9367–9376 (2020)
98. Turner, M., Koch, T., Serrano, F., Winkler, M.: Adaptive cut selection in mixed-integer linear programming. *Open J. Math. Optim.* **4**, 5 (2023)
99. Vinyals, O., Babuschkin, I., Chung, J., Mathieu, M., Jaderberg, M., Czarnecki, W., Dudzik, A., Huang, A., Georgiev, P., Powell, R., Ewalds, T., Horgan, D., Kroiss, M., Danihelka, I., Agapiou, J., Oh, J., Dalibard, V., Choi, D., Sifre, L., Sulsky, Y., Vezhnevets, S., Molloy, J., Cai, T., Budden, D., Paine, T., Gulcehre, C., Wang, Z., Pfaff, T., Pohlen, T., Yogatama, D., Cohen, J., McKinney, K., Smith, O., Schaul, T., Lillicrap, T., Apps, C., Kavukcuoglu, K., Hassabis, D., Silver, D.: AlphaStar: Mastering the Real-Time Strategy Game StarCraft II. (2019) <https://deepmind.com/blog/alphastar-mastering-real-time-strategy-game-starcraft-ii/>
100. Wang, Z., Li, X., Wang, J., Kuang, Y., Yuan, M., Zeng, J., Zhang, Y., Wu, F.: Learning cut selection for mixed-integer linear programming via hierarchical sequence model. In: International Conference on Learning Representations (2023) <https://openreview.net/forum?id=Zob4P9bRNcK>
101. Wesselmann, F., Stuhl, U.: Implementing cutting plane management and selection techniques. In: Technical Report, University of Paderborn (2012) <https://optimization-online.org/?p=12261>
102. Wirtz, M., Hahn, M., Schreiber, T., Müller, D.: Design optimization of multi-energy systems using mixed-integer linear programming: which model complexity and level of detail is sufficient? *Energy Convers. Manage.* **240**, 114249 (2021)
103. Wolsey, L.A.: *Integer Programming*. Wiley, London (2020)
104. Wu, Y., Song, W., Cao, Z., Zhang, J.: Learning large neighborhood search policy for integer programming. *Adv. Neural. Inf. Process. Syst.* **34**, 30075–30087 (2021)
105. Yilmaz, K., Yorke-Smith, N.: A study of learning search approximation in mixed integer branch and bound: node selection in SCIP. *AI* **2**(2), 150–178 (2021)
106. Zarpellon, G., Jo, J., Lodi, A., Bengio, Y.: Parameterizing branch-and-bound search trees to learn branching policies. *Proc. AAAI Conf. Artif. Intell.* **35**, 3931–3939 (2021)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.