Contents lists available at ScienceDirect

Pattern Recognition

journal homepage: www.elsevier.com/locate/patcog

Learning to select cuts for efficient mixed-integer programming

Zeren Huang^a, Kerong Wang^a, Furui Liu^{c,*}, Hui-Ling Zhen^{c,*}, Weinan Zhang^{a,*}, Mingxuan Yuan^c, Jianye Hao^c, Yong Yu^a, Jun Wang^b

^a Shanghai Jiao Tong University, China

^b University College London, United Kingdom

^c Noah's Ark Lab, Huawei Technologies, China

ARTICLE INFO

Article history: Received 1 April 2021 Revised 23 September 2021 Accepted 25 September 2021 Available online 4 October 2021

Keywords: Mixed-Integer programming Cutting plane Multiple instance learning Generalization ability

ABSTRACT

Cutting plane methods play a significant role in modern solvers for tackling mixed-integer programming (MIP) problems. Proper selection of cuts would remove infeasible solutions in the early stage, thus largely reducing the computational burden without hurting the solution accuracy. However, the major cut selection approaches heavily rely on heuristics, which strongly depend on the specific problem at hand and thus limit their generalization capability. In this paper, we propose a data-driven and generalizable cut selection approach, named CUT RANKING, in the settings of multiple instance learning. To measure the quality of the candidate cuts, a scoring function, which takes the instance-specific cut features as inputs, is trained and applied in cut ranking and selection. In order to evaluate our method, we conduct extensive experiments on both synthetic datasets and real-world datasets. Compared with commonly used heuristics for cut selection, the learning-based policy has shown to be more effective, and is capable of generalizing over multiple problems with different properties. CUT RANKING has been deployed in an industrial solver for large-scale MIPs. In the online A/B testing of the product planning problems with more than 10⁷ variables and constraints daily, CUT RANKING has achieved the average speedup ratio of 12.42% over the production solver without any accuracy loss of solution.

© 2021 Elsevier Ltd. All rights reserved.

1. Introduction

Combinatorial optimization (CO) is a subclass of optimization problems, where the goal is to find the optimal solution with respect to a given objective function from a finite candidate solution set. Due to its combinatorial nature (for example, integer constraints), it is usually NP hard and can mostly be formulated as mixed-integer programming (MIP) problems [1–4]. It covers a wide range of industry applications such as production planning, scheduling and manufacturing [5–8].

The difficulty of solving MIP problems lies on the non-convexity of its feasible region, which makes the general MIPs unsolvable in polynomial time. Instead of solving the MIP directly, one usually solves the corresponding LP relaxations first, and then performs rounding to generate the approximately optimal solution [9]. To facilitate such a process, a classic approach is the cutting plane method, which generates valid inequalities in the LP iterations to cut off the fractional solutions or the infeasible integer solutions, so that the convergence to optimum is accelerated [10].

* Corresponding authors.

E-mail addresses: liufurui2@huawei.com (F. Liu), zhenhuiling2@huawei.com (H.-L. Zhen), wnzhang@sjtu.edu.cn (W. Zhang). Another approach to solve MIP is the branch-and-bound algorithm [11,12], which creates branches by selecting variables to add rounding bounds to form two sub-LP problems (for a integer variable x_i with fractional value v as example, add two rounding bounds $x_i \ge \lceil v \rceil$ and $x_i \le \lfloor v \rfloor$), and then solve these sub-LP problems. In modern MIP solvers, the cutting plane technique is often combined with the branch-and-bound method to constitute the branch-and-cut framework [13], where each branch contains cutting planes as additional constraints. Heuristics are also employed in this process to cope with problems such as branching variable selection and cutting plane (cut) selection, i.e., selecting the most promising cuts to add. Such heuristics in MIP solvers are usually manually designed and heavily dependent on the problem. As a result, it shows high vulnerability with respect to the structure or the size of the MIP [14].

As a promising methodology for address above issues, many recent works [15–18] have leveraged machine learning (ML) techniques to construct efficient heuristics which is problem independent, and the majority are focusing on decision problems in the branch-and-bound algorithm, thus leaving room for machine learning methods on the direction of cut selection.

A good set of cuts is essential for the efficiency of the CO algorithms. Cuts serve the purpose to reduce the LP solution space,





which results in a smaller tree in branch-and-cut algorithm so that the number of nodes to be searched are significantly reduced. However, excessive quantity of cuts causes heavy computational workload on solving corresponding LP problems. As a consequence, deriving a good cut selection policy is of high value to the community, which has unfortunately received few attention so far. This motivates us to develop a general data-driven, machine learning based cut selection algorithm.

The purpose of this work is to construct an efficient and generalizable cut selection policy based on machine learning. The basic idea is to learn a scoring function that can measure the quality of cuts, and we formulate it as a *cut ranking* problem, in which we score each generated cut by a learned scoring function, and select a subset of cuts with the highest scores. However, such a task is non-trivial, with several remarkable technical challenges. First, in many cases, labels for individual cuts are not easy to obtain, since the impact of a single cut on MIP is relatively weak and imperceptible. Labeling good cuts individually may be infeasible. Thus, the cut selection problem naturally fits the scenarios of *multiple* instance learning (MIL), in which the collection of labels is at the bag level [19–21]. The training instances are organized into sets (also called bags), and the label is not assigned to any individual instance, but to the bag of cuts to measure the overall quality. Another important problem is the generalization ability of the learned scoring function. To enable the scoring function module to generalize to new problems, we design both static and dynamic problemspecific cut features as the inputs. Inspired by the training process of Reinforcement Learning (RL), we collect the supervised labels in an exploratory way [22]. Our proposed method is named as CUT RANKING.

To summarize, the technical contributions of our work are threefold.

- 1. We propose a novel CUT RANKING method for cut selection in the settings of multiple instance learning, which is suitable for the nature of cut selection tasks.
- 2. We study generalization ability to the cut selection policy since the designed cut features are determined by the characteristics of MIPs.
- 3. The CUT RANKING module can be applied as a subroutine in the branch-and-cut algorithm, which is generally adopted by modern MIP solvers.

The extensive experiments on various MIP problems demonstrate the superiority of CUT RANKING over previous solutions in terms of solving time and the node size of the branch-and-bound tree. Furthermore, we deploy CUT RANKING on Huawei's proprietary industrial large-scale MIP solver for production planning problems with more than 10⁷ variables and constraints daily. The A/B testings show that our solution can reduce the overall solving time by 14.98% and 12.42% on average for offline and online phases respectively without the accuracy loss of solution, which is a significant acceleration for the industrial solver.

2. Related works

The traditional approaches to tackle the MIPs mainly include: branch-and-bound [11,23,24] and cutting-plane methods [10,25]. They are widely deployed in modern solvers as the core algorithm for solving problems. However, in the era of big data, large scale problems with a lot of variables are often encountered, and those approaches suffer from very low efficiency. For scalability and speeding up the solvers, they are usually enhanced with heuristics, which are often designed by experts, based on the unique property of the problems at hand, and are not transferable or reusable when one switches to a new situation.

Therefore, there emerges a need for generalizable methods that are ubiquitous applicable to MIPs. Attentions are thus paid to machine learning and other data driven science due to their generalizability. Given training data, intelligent models are able to learn to solve the problem, with good performance on unseen data. Based on the role that the machine learning model takes, related literatures can be categorized into two clusters [26]. The first cluster contains methods that use ML models to replace traditional solving techniques. They can directly solve the MIP problems (such as TSPs), in which an end-to-end learning model is often used to predict the solution given the problem instance. Vinyals et al. [27] proposed the pointer network with a sequence-to-sequence architecture, and train the model through supervised learning. Bello et al. [28] introduced a reinforcement learning method to train the pointer network. Other literatures including [29,30] also use the sequence-to-sequence architecture to tackle the vehicle routing problems. More recently, Nair et al. [31] showed that the deep learning model is able to predict a good partial solution for MIP problems. The aforementioned works often use blackbox and unexplainable AI models, supported by empirical evidences but no theoretical guarantees. As a result, they sometimes show performance vulnerability in solving MIPs.

For the second cluster, in which our work can be placed, the main algorithmic framework is based on the traditional optimization algorithm, and machine learning is used to improve the heuristics. During the solution process, the ML model is repeatedly called to assist in making decisions. There have been multiple studies [15–17] about learning a branching policy in the context of branch-and-bound, in which the policy is usually trained through imitation learning to approximate a powerful heuristics named strong branching, which is effective but too slow. Khalil et al. [16] addressed the branching variable selection problem as a ranking problem.

In addition to learning a branching strategy, there are also studies on learning other core elements in the branch-and-bound framework. He et al. [32] proposed to learn a node selection policy to improve the heuristics. Khalil et al. [33] adopted a ML-based model to decide running a given heuristic or not at a branching node.

To our knowledge, the direction of deriving a ML-based cut selection policy for MIP problems have been rarely explored except Tang et al. [18], in which the authors introduced a MDP formulation for the problem of iteratively selecting cuts for a MIP, and train a reinforcement learning (RL) agent using evolutionary strategies. Our work differs from it in two aspects: first, we propose a ranking formulation, and model the learning problem in multiple instance learning settings; second, the main goal of Tang et al. [18] is to improve the efficiency of the cut selection module, i.e. to reduce the total number of cuts added, while the objective of our work is to improve the performance of the optimization algorithm, i.e. to reduce the total running time.

3. Background

Our algorithm combines MIL technique with the branch-andcut framework. In this section, we firstly introduce the background of MIP and branch-and-cut framework. Then, MIL related techniques are presented.

3.1. Mixed-integer programming background

Mixed-integer programming problem The general formulation of a MIP is as follows:

$$\arg\min_{\mathbf{x}} \{ \mathbf{z}^{\mathsf{T}} \mathbf{x} \mid \mathbf{A} \mathbf{x} \le \mathbf{b}, \mathbf{x} \in \mathbb{Z}^{p} \times \mathbb{R}^{n-p} \},$$
(1)

where **x** is the vector of decision variables, $\mathbf{z} \in \mathbb{R}^n$ is the objective coefficient vector, $\mathbf{b} \in \mathbb{R}^m$ is the right-hand side vector, $\mathbf{A} \in \mathbb{R}^{m \times n}$ is the constraint matrix.

Branch-and-cut

The branch-and-cut (BC), a combination of two classical algorithms, is widely adopted in modern MIP solvers. To address the difficulty brought by the nonconvexity on searching the optimal solution, it builds a search tree with each node corresponding to a linear programming problem. At the beginning of the algorithm, the root node corresponds to the problem with all integer constraints dropped. Then, it iteratively generates child nodes by selecting variables to branch on, that is, adding new constraints (bounds) on it. Along the paths of the tree, the space of solution is regularized by more constraints. For the problems corresponding to each node, valid cuts are added to assist searching. Taking the root node for example, assume that the added cut set $C' = \{\alpha_i^\top x \le \beta_i\}_{i=1}^{|C'|}$ is a subset of generated cut set *C*, the optimization problem becomes

$$\arg\min_{\mathbf{x}} \left\{ \mathbf{z}^{\mathsf{T}} \mathbf{x} \mid \mathbf{A} \mathbf{x} \le \mathbf{b}, \alpha^{\mathsf{T}} \mathbf{x} \le \beta, \mathbf{x} \in \mathbb{Z}^{p} \times \mathbb{R}^{n-p} \right\}.$$
(2)

The algorithm terminates when there exists no feasible solutions for one node, or we cannot obtain a better solution than the optimal one found so far. It is worth noting that the cuts are primarily added at the root node, which will often bring significant improvements.

Metric for cut quality For a given MIP, its solvability is defined as the capability of being solved, which relates to the size, the structure and other problem properties. A formal quantitative description of problem solvability is provided in Definition 1.

Definition 1 (Problem Solvability). Let *O* be the set of all the feasible optimization algorithms for solving MIPs, and *S* be the set of all the feasible MIP solvers. Assume that the computing environment is kept the same, we define the problem solvability of a MIP problem χ with respect to parameters *A*, *b*, *z* and integrality constraints as

$$PS(\chi) = \mathbb{E}_{o \sim 0, s \sim S} \left[\frac{1}{T_{\chi}} \middle| o, s \right],$$
(3)

where T_{χ} is the solving time of the MIP problem χ .

To measure the quality of selected cut subset C', we propose a metric named *problem solvability improvement* (PSI), which is calculated after adding cuts via

$$PSI = PS(\chi') - PS(\chi), \tag{4}$$

where χ' and χ represent the MIP with and without cuts, respectively.

In real practice, it is impractical to calculate *PSI* since obtaining the problem solvability is infeasible. However, when the optimization algorithm, the MIP solver and the computing environment are fixed, we can substitute *PSI* with the reduction ratio of solution time as the feedback r of selecting C':

$$r = \mathbb{E}\left[\frac{T_{\chi} - T_{\chi'}}{T_{\chi}} \middle| o, s\right],\tag{5}$$

where o is the optimization algorithm, s is the solver. A higher value of r implies a higher-quality cut subset for the MIP.

Typical cut types In cutting plane tasks, there exist various types of cuts which can be generated. Here, we list several typical types of cuts:

• **Cover Cut.** For a set of binary variables $X = \{x_1, x_2, ..., x_k\}$, a so-called knapsack constraint takes the form as

$$a_1 x_1 + a_2 x_2 + \ldots + a_k x_k \le b, \tag{6}$$

where a_1, a_2, \ldots, a_k, b are all non-negative. Let $X' = \{x'_1, x'_2, \ldots, x'_l\} \subset X$. A minimal cover cut related with the above knapsack constraint is of the form as

$$x'_1 + x'_2 + \ldots + x'_l \le l - 1.$$
 (7)

• **Gomory Cut.** The gomory cuts are generated from the rows of the simplex tableau, returned by the simplex algorithm for solving LPs. Here we use the similar notations as in Tang et al. [18]. Denote the constraint matrix and the constraint vector of the tableau as A' and b', respectively. For the *i*th row, the corresponding gomory cut can be generated by applying integer rounding as

$$(-A'_i + \lfloor A'_i \rfloor)x \le -b'_i + \lfloor b'_i \rfloor.$$
(8)

• **Clique Cut.** For a set of binary variables $X = \{x_1, x_2, ..., x_k\}$, a clique cut is of the form as

$$x_1 + x_2 + \ldots + x_k \le 1, \tag{9}$$

where at most one variable can be positive.

For a more detailed introduction of other cut types, one can refer to the surveys for cutting planes [10,25]. Note that these different types of cuts are enabled in general MIP solvers. For a given MIP instance, we can use the cut generators incorporated in the MIP solver to generate the candidate cuts.

3.2. Multiple instance learning

Multiple instance learning (MIL) concerns the problem of supervised learning where the model prediction and training are put at the level of bag of instances [19]. Each bag is composed of multiple unlabeled training instances. The goal is to predict the labels of unseen data at the bag level or at the instance level.

For the binary classification problems, where the label is positive or negative, the standard setting of MIL is that, bags containing at least one positive instances are assigned positive labels while bags containing only negative instances are assigned with negative labels [19]. This can be relaxed to the collective assumption, which is related to problems where the label assignment is determined by more than one instances. The MIL measures the effect of a set of instances by interpreting labels of bags, and this naturally fits the scenario of our cut selection. For large-scale problems, the effect of a single cut to the solution is rather minor and imperceptible. Therefore, we use MIL related techniques as our ML model for cut selection.

4. Methodology: cut ranking

4.1. A cut ranking formulation in MIL settings

In the branch-and-cut framework, we introduce cut selection at the root node of the branching tree. Since the majority of cuts are added to the root LP in general cases, the cuts are disabled at other sub-nodes for a better evaluation of effects on the algorithm. However, we argue that the learned cut selection policy can generalize to other sub-nodes (sub-MIPs) due to that the designed cut features are of the same properties with a MIP instance.

Definition 2 (MIP Problem Property). For a MIP with parameters $M = \{A, b, z\}$, after its root LP relaxation being solved, the LP solution x_{LP}^* is accessible. We define the problem property of the MIP before cut selection as $P = \{x_{LP}^*, M\}$.

For a given MIP, its problem property *P* is defined in Definition 2. Let $C = \{c_1, c_2, ..., c_l\}$ be the candidate cut set generated at the root node, the cut selection problem is equivalent to

select an optimal cut subset C^* with respect to the problem solvability improvement (PSI) mentioned in Eq. (4):

$$C^* = \operatorname*{arg\,max}_{C'} \{ PSI \mid C' \subset C, P \}. \tag{10}$$

Due to its combinatorial structure, finding an exact solution is intractable, especially when the size of *C* becomes larger.

To tackle such a problem, we present a CUT RANKING formulation in the branch-and-cut framework. The process of cut ranking involves the training phase and the test phase. In the training phase, the learning process is modeled in MIL settings, that is, the training data are grouped into bags, and the label assignment is at the bag level. Specifically, each bag consists of several cuts sampled from the candidate cut set, and the bags are not disjoint. Denote $u_i \in \mathbb{R}^H$ as the feature vector of cut c_i , which can be derived given the problem property *P* and the cut parameters. Let $X = \{x_1, x_2, \dots, x_h\}$ be the set of feature vectors for the collected bags, where each bag feature vector $x_i \in \mathbb{R}^H$ can be constructed through a feature mapping function $\phi(\cdot)$, taking the aggregated features of cuts within the bag as the input. Let $Y = \{y_1, y_2, \dots, y_h\}$ be the labels of X. Given the training set, our goal is to train a scoring function $f_{\theta}(u)$ which can predict the score for each candidate cut *u*, with the cut feature vector as the input.

In the test phase, for the given MIP, we use the trained scoring function f_{θ} to assign scores to all the generated candidate cuts, and select the top *K*% cuts with the highest scores (*K* is a hyperparameter).

4.2. Constructing training data

The training data is composed of features *X* and labels *Y* at the bag level. To construct it, we first collect the training samples (each sample corresponds to a bag of cuts) using a certain searching strategy on multiple randomly generated instances; next we extract the designed instance-specific cut features for the cuts within each training sample; after that, we construct the bag features using the aggregated features of cuts in each bag; finally, we assign binary label to each sample through a designed labeling scheme.

4.2.1. Strategies of collecting training samples

For a given MIP, after its root LP being solved, a set of candidate cuts *C* are generated by the cut generators incorporated in the MIP solver first. Note that for a MIP instance, when the MIP solver is fixed, the generated candidate cuts are also fixed. Denote hyperparameter K% as the cut selection ratio, a subset *C'* is selected using a stochastic cut selection policy, and under the selection ratio. After adding the selected cuts, the algorithm continues until terminating and returns the solution time. To measure the quality of the selected subset *C'*, since the pre-defined problem solvability is unavailable in practice, we use the reduction ratio of solution time mentioned in Eq. (5), as an alternative to *PSI*. The reduction ratio of solution time can also be regarded as the feedback *r* of selecting a cut subset.

For each MIP instance, we repeat running the solver for multiple times, and collect a number of training samples. Since our cut selection policy is stochastic, we are able to explore different cutting results and obtain training samples with much diversity. Note that though collecting training samples leads to multiple rounds of execution of the MIP solver, the whole process is conducted in an offline way, and thus the incurred training cost is acceptable. The collected training sample can be seen as a tuple (P, C', r) consisting of the MIP property P, the selected cut subset C' (bag of cuts), and the feedback r. To improve the exploration and also the data efficiency, we collect the training samples based on two strategies, random sampling and active sampling, which are similar to the searching strategies adopted by Bello et al. [28].

Random sampling The cut selection is based on a fixed stochastic policy, which randomly selects a subset of cuts to add to the MIP. For each generated MIP training instance, the algorithm is repeatedly called for a certain number of times, in which we apply random sampling to collect the initial training samples.

Active sampling In this case, we select the cuts using a pretrained cut selection policy, which can lead to more promising training samples compared with random sampling. However, collecting samples only based on the pre-trained policy will reduce the sample diversity, which may result in learning a sub-optimal policy. To alleviate this issue, we adopt an ϵ -greedy policy [34], which is a common approach in reinforcement learning to balance the exploration and exploitation:

$$C' = \begin{cases} \text{sample from policy } \pi, & \text{with probability } 1 - \epsilon \\ \text{sample randomly,} & \text{with probability } \epsilon \end{cases}$$
(11)

where C' is the selected cut subset, π is the cut selection policy derived from the model. During the active sampling phase, the cut selection policy is still being refined using the training samples collected in this phase. The flow chart is displayed in Fig. 1. Note that active sampling is on-policy, that is, we improve the same policy which is used to collect samples. However, the whole framework of CUT RANKING is in an offline settings, that is, during the test phase, we do not continue to train our model on new MIP instances.

4.2.2. Constructing bag features from cut features

Since the bag features are constructed from the aggregated cut features, thus the first issue is the specification of cut features. To enable better generalization of the model, we design 14 problem-specific atomic features for the cut selection task. Similar to the features of branching variables as provided in Khalil et al. [16], here we list our designed features for each candidate cut. Specifically, for a given MIP instance with problem property $P = \{x_{LP}^*, M\}$, for any generated cut c_i : $\alpha_i^\top x \leq \beta_i$, its cut features are shown in Table 1.

The top two atomic features correspond to the statistical information of coefficients related to the cut, which help to capture the structural information of the cut. The other features measure the cut characteristics through different measurements as mentioned by Wesselmann and Suhl [35], which capture the association between the MIP property and the cut. Moreover, all the above atomic features can be computed quite efficiently, and thus making the time to construct features negligible in the algorithm.

Now that we are able to construct the bag features. For a generated MIP instance, we collect a number of training samples using certain sampling strategies. For each training sample with a selected cut subset C', we first compute and collect the cut features for each cut within C', and obtain the set of corresponding cut features $U_{C'} = \{u'_1, u'_2, \dots, u'_{|C'|}\}$. To prevent feature dimensions from being on different scales, we apply Z-score normalization to re-scale the feature values among all the training samples collected from the same MIP instance. Finally, we introduce a feature mapping function ϕ , which maps the aggregated cut features to the original cut feature space, and we obtain the final bag features as

$$\mathbf{x}_{C'} = \phi(u'_1, u'_2, \dots, u'_{|C'|}).$$
(12)

The mapping function can be designed to capture the association between cuts within a bag, while in this work, we define ϕ as an average function for simplicity, which calculates an average of the cut features. In other works [36,37] which are under the collective assumption in MIL settings, the similar average or weighted function is also adopted. Moreover, our empirical studies will show that such a mapping function is effective.



Fig. 1. The flow chart of active sampling.

Table 1

Cut features' descriptions and counts.

Feature	Description	Count
stats. for cut coeffs.	the mean, max, min, stdev. of cut coefficients α_i	4
stats. for obj. coeffs.	the mean, max, min, stdev. of objective coefficients of cut variables	4
support	the proportion of non-zero coefficients in α_i	1
integral support	the proportion of non-zero coefficients w.r.t integer-constrained variables in α_i	1
normalized violation	$\max\{0, \frac{\alpha_i^\top x_{ip}^* - \beta_i}{ \beta_i }\}$, which measures the cut violation of the present LP solution	1
distance	$\frac{ \alpha_i^{-\gamma}x_p-\beta_i }{ \alpha_i }$, which measures the Euclidean distance between the present LP solution and the hyperplane $\alpha_i^{\top}x = \beta_i$ determined	1
	by the cut	
parallelism	$\frac{z^2 a_{\mu}}{ z a_{\mu} }$, which measures the parallelism between the objective function and the cut	1
expected improvement	$\frac{ z a_i^{o_i} x_{ip}^{o_i} - \beta_i }{ \alpha_i }$, which is an estimation for objective improvement with the cut	1

4.2.3. Assigning ranking labels to training samples

The labels for our cut selection task are defined to be binary, taking the value of 1 or 0. The positive labels are assigned to highquality cut subsets, which are preferable for the selection policy. To measure the quality of a sampled bag, we use the pre-defined reduction ratio of solution time as the feedback signal, and present a labeling scheme based on the ranking of bags. Note that deriving a precise labeling scheme requires collecting all the possible bags for a MIP instance, which is infeasible in general cases. Therefore, we adopt an approximation method, that is to sample a number of bags for each MIP instance, and assign the corresponding bag labels according to their rankings. Empirically, we find that the training is stable under such a labeling scheme.

For training bags $X = \{x_1, x_2, ..., x_{h'}\}$ collected from the same MIP instance, let $R = \{r_1, r_2, ..., r_{h'}\}$ be the set of corresponding feedbacks. For each sample x_i , we define its label as

$$y_j = \begin{cases} 1, & r_j \text{ ranks in the } \lambda\% \text{ highest feedbacks} \\ 0, & \text{otherwise} \end{cases}$$
(13)

where $\lambda \in (0, 100)$ is a tunable hyper-parameter, which controls the percentage of positive samples.

The above ranking based binary labeling scheme is suitable for cut selection since the main goal is to distinguish the high-quality cuts between the poor-quality ones without a full ranking of all candidate cuts.

4.3. Learning a scoring function

In this subsection, we first introduce the basic architecture of the scoring function, and then present how to train the model.

4.3.1. Scoring function architecture

We parameterize the scoring function as a neural network $f_{\theta}(x)$ with parameters θ . In the training phase, for each training bag, the model takes its bag features x as inputs, and outputs a probability distribution with two dimensions, which correspond to P(y = 1|x) and P(y = 0|x), respectively.

The architecture of the model consists of two parts. First, it extracts the bag embeddings x^e through a multi-layer perceptron (MLP), and then feed the embeddings to a softmax layer to output the probability distribution as

$$x^{e} = MLP(x; \theta)$$

$$P(y|x) = \text{softmax}(x^{e}).$$
(14)

Since the same form of bag features and the cut features, we are able to apply the scoring function at the cut level in the test phase.

4.3.2. Training the scoring function

We define the probability P(y = 1|x) output by the model as the score for the input bag or the cut. The goal is to train the model to output higher scores for positive samples, or high-quality cuts. Given a set of training data $\{(x_i^j, y_i^j)_{i=1}^{h_j}\}_{j=1}^N$ collected from *N* problem instances, we optimize the model parameters θ to minimize the loss function comprising the cross-entropy loss and the regularization loss as

$$L(\theta) = L_{ce}(\theta) + \gamma \Omega(\theta), \qquad (15)$$

where γ is a hyper-parameter for regularization penalty, and L_{ce} is the cross-entropy loss as

$$L_{ce}(\theta) = -\sum_{j=1}^{N} \sum_{i=1}^{h_j} y_i^j \log p_{\theta}(y_i^j = 1 | x) + (1 - y_i^j) \log p_{\theta}(y_i^j = 0 | x).$$
(16)

To avoid the overfitting of the model, we adopt a common L2 regularization into the model, which places restrictions on the model parameters as

$$\Omega(\theta) = \|\theta\|_2^2. \tag{17}$$

Finally, the major steps of our CUT RANKING method are summarized in Algorithm 1. The algorithm consists of three phases:

Algorithm 1 CUT RANKING in MIP settings.

1: Data Collection Phase:

- 2: Randomly generate a set of training instances D_{train} .
- 3: For each instance in D_{train} :
- 4: Sample training bags using strategies mentioned in Section 4.2.1;
- 5: Construct the bag features via Section 4.2.2;
- 6: Assign labels to the bags via Section 4.2.3.
- 7: Training Phase:
- 8: Initialize the scoring function f_{θ} and the learning rate μ .
- 9: Repeat
- 10: For each batch in training data:
- 11: Calculate the loss $L(\theta)$ in Equation 15;
- 12: Optimize model parameters: $\theta \leftarrow \theta \mu \nabla_{\theta} L(\theta)$;
- 13: until θ converge.

14: Test Phase:

- 15: For each instance in the test set D_{test} :
- 16: Obtain the scores for each candidate cut using f_{θ} ;
- 17: Select the top *K*% cuts with the highest scores.

data collection phase, training phase and the test phase. In the data collection phase, we collect and construct the training bags and labels; in the training phase, we train our scoring function in a supervised fashion; in the test phase, we apply the scoring function to each candidate cut, and select a cut subset with the highest scores.

5. Experiments

5.1. Experimental setup

5.1.1. Benchmarks

Our benchmarks consist of synthetic MIP problems and the real-world production planning problems. The synthetic MIPs consist of four classical classes of problems: Set Cover, Knapsack, Planning and General MIP. For the ease of data collection, we find the problem instances which are solvable within 25 seconds. For each class of problems, we randomly generate 100 training instances and 30 test instances using different random seeds.

The large-scale real-world daily production planning problems are divided into two phases, the *Offline Phase* during January 2021, and the *Online Phase* during March 2021. These two phases correspond to the offline datasets and online datasets for our experiments.

5.1.2. Metrics

The cut selection policy will affect the branch-and-cut algorithm from two aspects, the solution time and the number of nodes visited. Compared with the cases without cuts, we define two metrics for cut selection based on the above algorithm feedbacks: the reduction ratio of solution time; and the reduction ratio of the number of nodes visited.

Notably, the solution time is not directly determined by the number of nodes visited, since we need to also consider the solving time of each node relaxation. Therefore, the primary metric to qualify the cut selection policy is the reduction ratio of solution time. For each conducted experiment, we show the mean and standard deviation of results on the test set, and highlight the best average results.

5.1.3. Baselines

For synthetic datasets, we examine our CUT RANKING module against five widely used manually-designed heuristics for cut selection, including:

- RANDOM: select cuts according to a stochastic policy.
- VIOLATION: select the cuts with larger violation.
- NORMALIZED VIOLATION: select the cuts with larger normalized violation.
- DISTANCE: select the cuts with larger Euclidean distance from the root LP solution x_{IP}^* .
- PARALLELISM: select the cuts which are more parallel to the objective function.

For real-world datasets, we compare CUT RANKING with the fine-tuned manually heuristics which are adopted in Huawei's proprietary solver.

Note that RL2C [18] is not included in the baselines since their formulation is based on sequential decision making, and the main goal of RL2C is to reduce the total number of added cuts to solve the MIP to optimality. The main algorithmic framework of RL2C is the cutting plane method, which iteratively adds cuts to the initial LP relaxation. Specifically, in RL2C, a new cut is selected and added from the candidate cut set at each step and the solver is required to execute the LP instantly to obtain the change of objective value $\mathbf{z}^{\mathsf{T}}\mathbf{x}$. However, our defined task is essentially a one-step decision problem, and the main goal of CUT RANKING is to improve the solution time of MIP in the algorithmic framework of branch-and-cut. Therefore, RL2C is highly incompatible with our settings, and its RL formulation is impractical for large-scale MIPs since it may lead to much more sampling and computational costs. Moreover, the policy architecture of RL2C is based on LSTM, and the network inputs include all the constraints and available candidate cuts, which also makes it infeasible to apply in large-scale scenarios.

5.1.4. Implementation of algorithms

The algorithmic framework for synthetic MIPs is branch-andcut. We implement the vanilla branch-and-bound algorithm, and use the open-source solver Python-MIP [38] for solving LPs and generating cuts¹ For the real-world datasets, the optimization is based on a proprietary industrial solver of Huawei Company. We enable cut selection before the rounding procedure.

5.1.5. Hyper-parameters

Policy architecture The implemented policy network of CUT RANKING is a 4-layer fully-connected neural network, including an input layer, two hidden layers with 30 and 15 hidden units respectively and tanh activation, and an output layer. As mentioned in Section 4.3.1, the input layer accepts the cut (or bag) features with dimension 14 as the network inputs. The output layer outputs a probability distribution with two dimensions, and we define the positive probability as the score for the input cut (or bag). *Hyperparameters of algorithms*

For both synthetic datasets and real-world datasets, we set the hyper-parameter *K* to 30, λ to 50, and γ to 0.1 after hyperparameter tuning. For synthetic datasets, the number of sampled bags for each MIP instance is 100, and the total number of collected training samples is 10⁴. For real-world datasets, to improve the sample efficiency, the exploring policy starts from a welltuned heuristic combined with ϵ -greedy. The training samples are

¹ The code for reproducing the experiments on synthetic datasets is available at https://github.com/logWin322/Learn2SelectCuts.

Table 2

Evaluation results of cut selection policies in terms of	the reduction ratio of solving time (higher is better), and the reduction ratio of visited nodes (higher is be	etter).
--	---	---	---------

Method	Set cover		Knapsack		Planning	Planning		General MIP	
	Time	Nodes	Time	Nodes	Time	Nodes	Time	Nodes	
Random	0.09 ± 0.45	0.27 ± 0.48	0.14 ± 0.27	0.62 ± 0.67	0.07 ± 0.14	0.40 ± 0.30	-0.10 ± 0.11	-0.03 ± 0.12	
VIOLATION	0.17 ± 0.19	0.35 ± 0.18	0.21 ± 0.33	0.50 ± 0.88	-0.01 ± 0.24	0.36 ± 0.35	0.11 ± 0.19	0.25 ± 0.20	
NORM-VIOLATION	0.16 ± 0.22	0.34 ± 0.22	0.25 ± 0.24	$\textbf{0.70} \pm 0.38$	0.17 ± 0.20	0.48 ± 0.30	0.13 ± 0.39	0.20 ± 0.35	
DISTANCE	0.12 ± 0.36	0.30 ± 0.40	0.10 ± 0.19	0.41 ± 0.44	0.01 ± 0.15	0.35 ± 0.33	0.10 ± 0.28	0.25 ± 0.20	
Parallelism	0.06 ± 0.34	0.23 ± 0.37	0.16 ± 0.18	0.58 ± 0.34	-0.06 ± 0.17	0.34 ± 0.36	0.03 ± 0.20	0.26 ± 0.20	
Cut Ranking	$\textbf{0.21} \pm 0.16$	$\textbf{0.49} \pm 0.16$	$\textbf{0.27}\pm0.20$	0.69 ± 0.32	$\textbf{0.18} \pm 0.17$	$\textbf{0.48} \pm 0.28$	$\textbf{0.32} \pm 0.25$	$\textbf{0.38} \pm 0.20$	

collected from the daily production planning problems within a month.

5.2. Experiments on synthetic MIP datasets

5.2.1. Experiment I: the quality of selected cuts

To check the effectiveness of our proposed ranking-based cut selection policy, we conduct comparative experiments on four classes of MIP problems. For Set Cover, the number of elements and sets are both set to 200, and the resulting problem size is 200×200 ; for Knapsack, the number of items is 700, and the resulting problem size is 701×700 ; for Planning, the number of factories and demands is 20 and 50, respectively, and the resulting problem size is 140×1420 ; for General MIP, owing to its complex problem size varies from class to class, the mean solution time (without cuts) is close, thus the difficulty of MIP instances for each class is at the same level. For each class of instances, the number of generated candidate cuts is roughly 20.

As shown in Table 2, in terms of the problem solving time, our proposed CUT RANKING policy has achieved higher average reduction ratio of solution time over other baseline policies on all the problems, which leads to less solution time. Moreover, the CUT RANKING policy has also shown to be more stable on multiple instances since the standard deviation is relatively smaller compared to the mean. For the Knapsack and Planning problems, the average performance of the heuristic NORMALIZED VIOLATION is comparable to us, while with a large variance, which indicates that such a heuristic often suffers from performance fluctuations. The results are similar for other human-designed heuristics, compared to the CUT RANKING policy, they have shown larger performance variance, and may slow down the solving process in many cases.

Considering the impact of cut selection on the size of the branch-and-cut search tree, for the CUT RANKING policy, the number of nodes visited has decreased more significantly than other baselines over the Set Cover and General MIP problem instances. For the Knapsack and Planning problems, our proposed policy has also shown to be competitive, and also achieve smaller variance.

Overall, these results indicate that the CUT RANKING policy has improved the optimization algorithm more significantly compared to the human-designed heuristics. Moreover, the CUT RANKING policy is capable of speeding up the solving time for all the generated test cases, while the effects of other cut selection heuristics suffer from instability.

5.2.2. Experiment II: study of generalization ability

To test if our proposed policy has the generalization ability over problems with different structures or scales, we conduct three experiments to try to answer the following questions:

- Can the CUT RANKING policy generalize to the same class of problems with different sizes?
- Can the CUT RANKING policy generalize to the same class of problems with different coefficient ranges?
- Can the CUT RANKING policy generalize to the problems with different structures?

Problem size Table 3 presents the results on Knapsack problems with different problem sizes. The learning-based policy is trained on the problem instances with 700 items, and tested against other heuristics on instances with 600, 800, 900 and 1000 items. As can be seen from the table, the CUT RANKING policy has shown a higher reduction ratio of solution time over other baselines on test instances with 600, 800 and 1000 items. For problem instances with 800 items, although the heuristic VIOLATION and PARALLELISM have achieved a slightly higher averaged reduction ratio, our policy has a much lower variance, thus is more preferable for cut selection.

The CUT RANKING policy also results in a smaller branching tree compared to most baseline heuristics. From these two aspects, we conclude that our proposed policy trained on problem instances of a certain scale can be applied to the same class of problems of different scales.

Coefficient ranges The parameters of Knapsack problems include *the maximal number, the maximal value* and *the maximal weight* for one type of item, which restricts the range of each randomly generated coefficient. We generate four sets of Knapsack instances with parameters set to 10, 20, 50 and 100. We train our cut selection module on the instances with parameters set to 10, and test on other three sets of instances.

The results are set out in Table 4, from which we can observe that the CUT RANKING policy clearly outperforms other baselines on problem instances with coefficients range between 0 and 20. For problems with coefficients range between 0 and 50, the CUT RANK-ING policy is still superior to the baselines, while the performance

Table 3

Evaluation results of cut selection policies on Knapsack problems with different scales. We train our learning module on 100 randomly generated Knapsack instances with 700 items.

Method	600 items		800 items		900 items		1000 items	
	Time	Nodes	Time	Nodes	Time	Nodes	Time	Nodes
Random Violation Norm-Violation Distance	$\begin{array}{c} 0.20 \pm 0.29 \\ 0.27 \pm 0.31 \\ 0.18 \pm 0.31 \\ 0.21 \pm 0.28 \end{array}$	$\begin{array}{c} 0.69 \pm 0.28 \\ 0.75 \pm 0.30 \\ \textbf{0.75} \pm 0.31 \\ 0.62 \pm 0.35 \end{array}$	$\begin{array}{c} 0.20 \pm 0.24 \\ 0.21 \pm 0.25 \\ 0.30 \pm 0.26 \\ 0.29 \pm 0.27 \end{array}$	$\begin{array}{c} 0.64 \pm 0.31 \\ 0.67 \pm 0.32 \\ 0.74 \pm 0.33 \\ 0.67 \pm 0.39 \end{array}$	$\begin{array}{c} 0.26 \pm 0.41 \\ 0.31 \pm 0.27 \\ 0.24 \pm 0.28 \\ 0.22 \pm 0.39 \end{array}$	$\begin{array}{c} 0.65 \pm 0.70 \\ \textbf{0.77} \pm 0.27 \\ 0.71 \pm 0.27 \\ 0.65 \pm 0.60 \end{array}$	$\begin{array}{c} 0.21 \pm 0.20 \\ 0.23 \pm 0.24 \\ 0.17 \pm 0.24 \\ 0.22 \pm 0.22 \end{array}$	$\begin{array}{c} 0.55 \pm 0.40 \\ 0.50 \pm 0.54 \\ 0.50 \pm 0.53 \\ 0.57 \pm 0.45 \end{array}$
Parallelism Cut Ranking	$\begin{array}{c} 0.10\pm0.24\\ \textbf{0.28}\pm0.15\end{array}$	$\begin{array}{c} 0.58 \pm 0.40 \\ 0.74 \pm 0.22 \end{array}$	$\begin{array}{c} 0.27\pm0.26\\ \textbf{0.33}\pm0.24\end{array}$	$\begin{array}{c} 0.66\pm0.38\\ \textbf{0.80}\pm0.31\end{array}$	$\begin{array}{c} \textbf{0.33} \pm 0.40 \\ 0.31 \pm 0.16 \end{array}$	$\begin{array}{c} 0.74 \pm 0.68 \\ 0.71 \pm 0.18 \end{array}$	$\begin{array}{c} 0.19 \pm 0.20 \\ \textbf{0.25} \pm 0.18 \end{array}$	$\begin{array}{c} 0.53\pm0.41\\ \textbf{0.61}\pm0.38\end{array}$

Table 4

Evaluation results of cut selection policies on Knapsack problems with the same size (700 items) but different ranges of coefficients. We train our learning module on Knapsack instances with coefficients range between (0,10].

Method	$0 < coeff. \leq 20$	$0 < coeff. \leq 20$			$0 < coeff. \leq 100$		
	Time	Nodes	Time	Nodes	Time	Nodes	
Random	0.42 ± 0.54	0.63 ± 0.87	0.36 ± 0.35	0.70 ± 0.58	0.14 ± 0.27	0.62 ± 0.67	
VIOLATION	0.43 ± 0.51	0.64 ± 0.80	0.35 ± 0.40	$\textbf{0.71} \pm 0.44$	0.21 ± 0.33	0.50 ± 0.88	
Norm-Violation	0.41 ± 0.52	0.62 ± 0.80	0.37 ± 0.27	0.67 ± 0.40	0.25 ± 0.24	0.70 ± 0.38	
DISTANCE	0.30 ± 0.32	0.56 ± 0.38	0.32 ± 0.26	0.68 ± 0.30	0.10 ± 0.19	0.41 ± 0.44	
Parallelism	0.33 ± 0.34	0.58 ± 0.41	0.33 ± 0.28	0.65 ± 0.34	0.16 ± 0.18	0.58 ± 0.34	
Cut Ranking	$\textbf{0.52} \pm 0.30$	0.86 ± 0.23	$\textbf{0.39} \pm 0.27$	0.68 ± 0.39	0.23 ± 0.19	0.65 ± 0.30	

Table 5

Evaluation results of cut selection policies on different classes of problems. We trained our learning module on Knapsack instances, and test on other classes of problem instances which at the same difficulty level.

Method	Set cover Time Nodes		Planning		General MIP	General MIP	
			Time Nodes Time N		Nodes	Time	Nodes
Random	0.09 ± 0.45	0.27 ± 0.48	0.07 ± 0.14	0.40 ± 0.30	-0.10 ± 0.11	-0.03 ± 0.12	
VIOLATION	0.17 ± 0.19	0.35 ± 0.18	-0.01 ± 0.24	0.36 ± 0.35	0.11 ± 0.19	0.25 ± 0.20	
Norm-Violation	0.16 ± 0.22	0.34 ± 0.22	$\textbf{0.17} \pm 0.20$	$\textbf{0.48}\pm0.30$	0.13 ± 0.39	0.20 ± 0.35	
DISTANCE	0.12 ± 0.36	0.30 ± 0.40	0.01 ± 0.15	0.35 ± 0.33	0.10 ± 0.28	0.25 ± 0.20	
Parallelism	0.06 ± 0.34	0.23 ± 0.37	-0.06 ± 0.17	0.34 ± 0.36	0.03 ± 0.20	$\textbf{0.26} \pm 0.20$	
Cut Ranking	$\textbf{0.22} \pm 0.18$	$\textbf{0.52}\pm0.17$	0.16 ± 0.11	0.33 ± 0.26	$\textbf{0.15}\pm0.23$	0.23 ± 0.18	

gap between them has decreased much. Such a phenomenon is more striking for problems with larger ranges of coefficients, as demonstrated in the rightmost two columns of Table 4, the CUT RANKING policy does not have clear advantages over the heuristic NORMALIZED VIOLATION.

Taken together, the results reveal that our CUT RANKING policy can generalize to the problems with different coefficient ranges. However, problems with a large range of coefficients will limit the generalization ability of the learned policy.

Problem structure To explore the generalization ability of the CUT RANKING policy on problems with different structures, we test the policy trained on Knapsack instances on Set Cover, Planning and General MIP problems. The generated problem size in this experiment is kept the same as experiment I.

From the results shown in Table 5, the CUT RANKING policy outperforms other baselines on Set Cover problems, which shows greater average improvement with a lower variance. For Planning problems, the heuristic NORMALIZED VIOLATION has achieved slightly better average performance compared to the learned policy, however, the variance is much larger. For General MIPs, the CUT RANKING policy also outperforms other baselines, nevertheless, it fails to reduce the solving time of some test instances. In view of the more complex problem structure for General MIPs, it is more difficult for the learned policy to generalize to these problem instances.

In summary, these results show that the CUT RANKING policy has certain generalization ability on problems with different structures. The results also indicate that the General MIPs and the Knapsack problems may have a relatively greater distinction in problem structures, which may lead to increased divergence between the distribution of training and test dataset.

5.3. Experiments on large-scale real-world tasks

To further evaluate the quality of the proposed CUT RANKING policy, we embed the cut selection module in an industrial largescale MIP solver developed by Huawei Company, and conduct both offline and online experiments on the real-world production planning problems with more than 10⁷ variables and constraints daily. To our knowledge, this is the first study to apply machine learning into cut selection for large-scale MIPs with more than 10⁷ variables and constraints.

MIP statistics The daily production-planning problems have similar problem structures, while problem properties change day by day. For a better description of problem characteristics, we record the number of variables and constraints, the density of the constraint matrix, the mean and standard deviation of objective coefficients, and the number of integer variables.

Table 6 shows the mean MIP statistics of production planning datasets. The upper six sub-figures of Figs. 2 and 3 demonstrate the visualization of problem statistics during January 2021 and March 2021 for offline and online datasets, respectively. Moreover, the number of generated candidate cuts for each problem instance is between 1200 and 1400.

Comparison experiments Our learning module is trained on the offline collected samples, and applied to both offline and online datasets. We compare our CUT RANKING policy against the production planning solver with a manually heuristic for cut selection. As shown in Figs. 2 and 3, our CUT RANKING policy has led to less solution time on both offline and online datasets without the accuracy loss of solution, and the average speedup ratio has reached 14.98% and 12.42%, respectively. Interestingly, for the problem instance of day 7 in offline datasets, the solver with the manually heuristic is unable to return a solution within a limited time; for the problem instance of day 5 in online datasets, the solver with the manually heuristic costs much more time to obtain the solution compared to the solver with the learning module. These results have further demonstrated the importance of deriving a proper cut selection policy. Our proposed ranking-based cut selection policy has shown to be more robust and efficient compared to the baseline, and is also generalizable to the daily problems with different properties.

Strategy captured by CUT RANKING CUT RANKING will learn to find the informative cut features or the informative feature combinations for the given class of MIP problems. We analyze CUT RANK-ING on the real-world production planning problems since they have particular problem structures. As we have tested, the heuristic based on PARALLELISM performs well on the problem instances, which indicates that PARALLELISM is one of the informative cut features for the production planning problems. We find that CUT 8000





Fig. 2. The daily MIP statistics and the evaluation results for offline datasets.



Fig. 3. The daily MIP statistics and the evaluation results for online datasets.

Table 6

Tl

ıe	mean	MIP	statistics	of	the	large-scale	production	planning	tasks.



Fig. 4. Box plot of informative cut features for production planning problems.

RANKING has learned to select cuts with larger PARALLELISM as well. Besides, CUT RANKING also prefers cuts with larger mean value of objective coefficients of cut variables. For a better demonstration, we visualize the informative features found by CUT RANKING using the box plot, which can be seen in Fig. 4.

6. Conclusion

In this paper, we presented CUT RANKING method for cut selection in the context of branch-and-cut algorithm for MIPs. To tackle the infeasibility of acquiring the supervised label for a single cut the task, we proposed to model the learning process in the settings of multiple instance learning. Moreover, we designed several problem-specific features for cuts, and provided a scheme to construct bag features and labels for training. The experimental results on synthetic MIPs have demonstrated that our learned ranking-based cut selection policy is more competitive compared to other manually heuristics, and also with generalization ability on problems with different scales, coefficient ranges or structures. For the real-world production planning tasks, our CUT RANKING method has also significantly improved the efficiency Huawei's industrial MIP solver without the accuracy loss of solution, achieving a speedup ratio of 14.98% and 12.42% in offline and online A/B testings, respectively. The empirical findings of this work provide a deeper insight into the generalization ability of machine learning techniques for cut selection, and reveal that the machine learning module can be incorporated in the solver to improve the solution process even for large-scale MIPs.

Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

Weinan Zhang is supported by "New Generation of AI 2030" Major Project (2018AAA0100900) and National Natural Science Foundation of China (62076161). The work is also sponsored by Huawei Innovation Research Program.

References

- R.E. Bixby, A brief history of linear and mixed-integer programming computation, Doc. Math. (2012) (2012) 107–121.
- [2] A. Richards, J. How, Mixed-integer programming for control, in: Proceedings of the 2005, American Control Conference, 2005., IEEE, 2005, pp. 2676–2683.
- [3] T. Achterberg, R. Wunderling, Mixed integer programming: analyzing 12 years of progress, in: Facets of Combinatorial Optimization, Springer, 2013, pp. 449–481.
- [4] R.E. Bixby, M. Fenelon, Z. Gu, E. Rothberg, R. Wunderling, Mixed-integer programming: a progress report, in: The Sharpest Cut: The Impact of Manfred Padberg and His Work, SIAM, 2004, pp. 309–325.
- [5] T. Wu, K. Akartunali, J. Song, L. Shi, Mixed integer programming in production planning with backlogging and setup carryover: modeling and algorithms, Discrete Event Dyn. Syst. 23 (2) (2013) 211–239.
- [6] T. Schouwenaars, B. De Moor, E. Feron, J. How, Mixed integer programming for multi-vehicle path planning, in: 2001 European Control Conference (ECC), IEEE, 2001, pp. 2603–2608.
- [7] A.B. Keha, K. Khowala, J.W. Fowler, Mixed integer programming formulations for single machine scheduling problems, Comput. Ind. Eng. 56 (1) (2009) 357–367.
- [8] A.R. Amaral, A mixed-integer programming formulation for the double row layout of machines in manufacturing systems, Int. J. Prod. Res. 57 (1) (2019) 34–47.
- [9] L.A. Wolsey, Mixed integer programming, Wiley Encycl. Comput. Sci. Eng. (2007) 1–10.
- [10] H. Marchand, A. Martin, R. Weismantel, L. Wolsey, Cutting planes in integer and mixed integer programming, Discrete Appl. Math. 123 (1-3) (2002) 397-446.
- [11] E.L. Lawler, D.E. Wood, Branch-and-bound methods: a survey, Oper. Res. 14 (4) (1966) 699–719.
- [12] M. Ris, J. Barrera, D.C. Martins Jr, U-curve: a branch-and-bound optimization algorithm for u-shaped cost functions on boolean lattices applied to the feature selection problem, Pattern Recognit. 43 (3) (2010) 557–568.
- [13] J.E. Mitchell, Branch-and-cut algorithms for combinatorial optimization problems, Handb. Appl. Optim. 1 (2002) 65–77.
- [14] T. Achterberg, T. Berthold, G. Hendel, Rounding and propagation heuristics for mixed integer programming, in: Operations Research Proceedings 2011, Springer, 2012, pp. 71–76.
- [15] M.-F. Balcan, T. Dick, T. Sandholm, E. Vitercik, Learning to branch, in: International Conference on Machine Learning, PMLR, 2018, pp. 344–353.
- [16] E.B. Khalil, P.L. Bodic, L. Song, G. Nemhauser, B. Dilkina, Learning to branch in mixed integer programming, in: Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence, 2016, pp. 724–731.
- [17] M. Gasse, D. Chételat, N. Ferroni, L. Charlin, A. Lodi, Exact combinatorial optimization with graph convolutional neural networks, in: Proceedings of the 33rd International Conference on Neural Information Processing Systems-Volume 2, 2019, pp. 15580–15592.

- [18] Y. Tang, S. Agrawal, Y. Faenza, Reinforcement learning for integer programming: learning to cut, in: International Conference on Machine Learning, PMLR, 2020, pp. 9367–9376.
- [19] M.-A. Carbonneau, V. Cheplygina, E. Granger, G. Gagnon, Multiple instance learning: a survey of problem characteristics and applications, Pattern Recognit. 77 (2018) 329–353.
- [20] J. Foulds, E. Frank, A review of multi-instance learning assumptions, Knowl. Eng. Rev. 25 (1) (2010) 1–25.
- [21] V. Cheplygina, D.M. Tax, M. Loog, Multiple instance learning with bag dissimilarities, Pattern Recognit. 48 (1) (2015) 264–275.
 [22] K. Arulkumaran, M.P. Deisenroth, M. Brundage, A.A. Bharath, Deep reinforce-
- [22] K. Arulkumaran, M.P. Deisenroth, M. Brundage, A.A. Bharath, Deep reinforcement learning: a brief survey, IEEE Signal Process. Mag. 34 (6) (2017) 26– 38.
- [23] J. Clausen, Branch and Bound Algorithms-Principles and Examples, Department of Computer Science, University of Copenhagen, 1999, pp. 1–30.
- [24] D.R. Morrison, S.H. Jacobson, J.J. Sauppe, E.C. Sewell, Branch-and-bound algorithms: a survey of recent advances in searching, branching, and pruning, Discrete Optim. 19 (2016) 79–102.
- [25] G. Cornuéjols, Valid inequalities for mixed integer linear programs, Math. Program. 112 (1) (2008) 3–44.
- [26] Y. Bengio, A. Lodi, A. Prouvost, Machine learning for combinatorial optimization: a methodological tour d'horizon, Eur. J. Oper. Res. 290 (2) (2021) 405–421.
- [27] O. Vinyals, M. Fortunato, N. Jaitly, Pointer networks, in: Proceedings of the 28th International Conference on Neural Information Processing Systems-Volume 2, 2015, pp. 2692–2700.
- [28] I. Bello, H. Pham, Q.V. Le, M. Norouzi, S. Bengio, Neural combinatorial optimization with reinforcement learning, in: The Workshop Track of the 5th International Conference on Learning Representations, 2017, pp. 1–5.
- [29] W. Kool, H. Van Hoof, M. Welling, Attention, learn to solve routing problems!, in: Proceedings of the 7th International Conference on Learning Representations, 2019, pp. 1–25.
- [30] M. Nazari, A. Oroojlooy, M. Takáč, L.V. Snyder, Reinforcement learning for solving the vehicle routing problem, in: Proceedings of the 32nd International Conference on Neural Information Processing Systems, 2018, pp. 9861–9871.
- [31] V. Nair, S. Bartunov, F. Gimeno, I. von Glehn, P. Lichocki, I. Lobov, B. O'Donoghue, N. Sonnerat, C. Tjandraatmadja, P. Wang, et al., Solving mixed integer programs using neural networks, arXiv preprint arXiv:2012.13349 (2020).
- [32] H. He, H. Daumé III, J. Eisner, Learning to search in branch-and-bound algorithms, in: Proceedings of the 27th International Conference on Neural Information Processing Systems-Volume 2, 2014, pp. 3293–3301.
- [33] E.B. Khalil, B. Dilkina, G.L. Nemhauser, S. Ahmed, Y. Shao, Learning to run heuristics in tree search, in: Proceedings of the 26th International Joint Conference on Artificial Intelligence, 2017, pp. 659–666.
- [34] R.S. Sutton, A.G. Barto, Reinforcement Learning: An Introduction, MIT Press, 2018.
- [35] F. Wesselmann, U.H. Suhl, Implementing Cutting Plane Management and Selection Techniques, Technical Report, University of Paderborn, 2012.
- [36] Z. Wang, V. Radosavljevic, B. Han, Z. Obradovic, S. Vucetic, Aerosol optical depth prediction from satellite observations by multiple instance regression, in: Proceedings of the 2008 SIAM International Conference on Data Mining, SIAM, 2008, pp. 165–176.
- [37] N. Pappas, A. Popescu-Belis, Explaining the stars: weighted multiple-instance learning for aspect-based sentiment analysis, in: Proceedings of the 2014 Conference on Empirical Methods In Natural Language Processing (EMNLP), 2014, pp. 455–466.

[38] M.J. Saltzman, Coin-or: an open-source library for optimization, in: Programming Languages and Systems in Computational Economics and Finance, Springer, 2002, pp. 3–32.

Zeren Huang is a master student in computer science at the Shanghai Jiao Tong University, Shanghai, China. His current research interest includes machine learning, deep learning and reinforcement learning

Kerong Wang is a junior undergraduate student in School of Electronic Information and Electrical Engineering, Shanghai Jiao Tong University, Shanghai, China. His current research interest includes machine learning and reinforcement learning

Furui Liu is a senior researcher at Huawei Noah's Ark Lab. He obtained the Ph.D. from The Chinese University of Hong Kong in 2019, and the Bachelor in Engineering from Zhejiang University. He was a visiting student at Max Planck Institute for Intelligent Systems, and Microsoft Research Asia. His research interests include causal inference, deep learning and applications in large scale problems

Hui-Ling Zhen received the Ph.D. degree on applied mathematics from Beijing University of Posts and Telecommunications. She is a Senior Engineer with Huawei Noah's Ark Lab, Hong Kong. Her current research interests include large-scale optimization, constraint programming, as well as their applications in supply chain management and chip design

Weinan Zhang is now a tenure-track associate professor at Shanghai Jiao Tong University. His research interests include reinforcement learning, deep learning and data science with various real-world applications. Weinan earned his Ph.D. from University College London in 2016 and B.E. Shanghai Jiao Tong University in 2011

Mingxuan Yuan received the Ph.D. degree from Hong Kong University of Science and Technology. He is a Principal Researcher with Huawei Noah's Ark Lab. He has led several projects in telecommunication data mining and supply chain optimization. His research interests include spatiotemporal data analytics and enterprise operation optimization models

Jianye Hao is now an associate professor at Tianjin University and Director of Decision Making and Reasoning Lab of Noah's Ark Lab, Huawei. He was a postdoctoral research fellow in the Computer Science and Artificial Intelligence Laboratory at MIT (2013–2015). His research interests include deep reinforcement learning, and multiagent systems.

Yong Yu is a professor in the Department of Computer Science, Shanghai Jiao Tong University. His research interests include information systems, web search, data mining, and machine learning. He has published more than 200 papers and served as a PC member of several conferences including WWW, RecSys, and a dozen of other related conferences (e.g., NIPS, ICML, SIGIR, ISWC, etc.) in these fields

Jun Wang is a professor of information systems and data science at University College London. He has published more than 100 research papers in leading journals and conference proceedings including the ACM Transactions on Information Systems, the IEEE Transactions on Knowledge and Data Engineering, WWW, CIKM, SIGIR, KDD, SIGMM, AAAI, etc. He received the Doctoral Consortium award in ACM SIGIR'06, and the Best Paper Prize in ECIR'09, ECIR'12, ADKDD'14, and SIGIR'17