



OPEN ACCESS

EDITED BY

Andreas M. Tillmann,
Technical University of Braunschweig, Germany

REVIEWED BY

Christoph Hansknecht,
Technical University of Braunschweig, Germany
Stephan Westphal,
Clausthal University of Technology, Germany

*CORRESPONDENCE

Andrea Lodi

✉ andrea.lodi@cornell.edu

RECEIVED 20 December 2022

ACCEPTED 10 July 2023

PUBLISHED 26 July 2023

CITATION

Liu D, Perreault V, Hertz A and Lodi A (2023) A machine learning framework for neighbor generation in metaheuristic search. *Front. Appl. Math. Stat.* 9:1128181. doi: 10.3389/fams.2023.1128181

COPYRIGHT

© 2023 Liu, Perreault, Hertz and Lodi. This is an open-access article distributed under the terms of the [Creative Commons Attribution License \(CC BY\)](https://creativecommons.org/licenses/by/4.0/). The use, distribution or reproduction in other forums is permitted, provided the original author(s) and the copyright owner(s) are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.

A machine learning framework for neighbor generation in metaheuristic search

Defeng Liu¹, Vincent Perreault¹, Alain Hertz¹ and Andrea Lodi^{2*}

¹Department of Mathematics and Industrial Engineering, Polytechnique Montreal, Montreal, QC, Canada,

²Jacobs Technion-Cornell Institute, Cornell Tech and Technion-IIT, New York, NY, United States

This paper presents a methodology for integrating machine learning techniques into metaheuristics for solving combinatorial optimization problems. Namely, we propose a general machine learning framework for neighbor generation in metaheuristic search. We first define an efficient neighborhood structure constructed by applying a transformation to a selected subset of variables from the current solution. Then, the key of the proposed methodology is to generate promising neighbors by selecting a proper subset of variables that contains a descent of the objective in the solution space. To learn a good *variable selection* strategy, we formulate the problem as a *classification* task that exploits structural information from the characteristics of the problem and from high-quality solutions. We validate our methodology on two metaheuristic applications: a *Tabu Search* scheme for solving a Wireless Network Optimization problem and a *Large Neighborhood Search* heuristic for solving Mixed-Integer Programs. The experimental results show that our approach is able to achieve a satisfactory trade-offs between the exploration of a larger solution space and the exploitation of high-quality solution regions on both applications.

KEYWORDS

combinatorial optimization, metaheuristics, Tabu Search, Large Neighborhood Search, machine learning, Graph Neural Networks, Mixed Integer Programming (MIP)

1. Introduction

Combinatorial Optimization (CO) is an important class of optimization problems in Operations Research (OR) and Computer Science (CS). In general, a CO problem is defined by a set of decision variables, a constrained solution space and an objective function. The goal of CO is to find optimal solutions with respect to the objective in the solution space.

Classical methods for solving CO problems can be roughly divided into three classes: exact methods, heuristics and metaheuristics. Mixed-Integer Programming (MIP) is one of the main paradigms of exact methods for modeling complex CO problems. Over the last decades, there has been increasing interest in improving the ability to solve MIPs effectively. Modern MIP solvers incorporate a variety of complex algorithmic techniques, such as primal heuristics [1], Branch and Bound (B&B) [2], cutting planes [3] and pre-processing, which results in complex and sophisticated software tools.

Exact methods are guaranteed to find optimal solutions as well as a proof of their optimality. On the one hand, due to the NP-hardness nature of many CO problems, solving them to optimality within an exact algorithm is still a very challenging task. On the other hand, in many practical applications, one is more interested in getting a good solution within a reasonable time rather than finding an optimal one. Those practical requirements have motivated the development of *specific heuristics* and *metaheuristics* (MHs). Specific heuristics

are usually designed for solving a certain type of CO problem, whereas MHs are frameworks for designing heuristics for solving general CO problems, and provide guidelines to integrate basic heuristic concepts with high-level diversification strategies.

It is worth noting that a lot of information can be produced and observed from MH processes, and therefore, a large volume of data can be collected. These data might provide valuable information about the optimization status of the process, the characteristics of the problem, the structures and properties of high-quality solutions in the solution regions being visited. However, such knowledge has not been fully exploited by traditional MH algorithms.

That can be viewed as a disadvantage when compared to application-targeted CO algorithms, in which exploiting the domain knowledge of the problem is generally preferable. Meanwhile, real-world CO problems have a rich structure. With similar instances repeatedly solved in many applications, statistical characteristics and patterns appear. This provides the opportunity for Machine Learning (ML) to extract structural properties of the problem from data and automatically produce learning-based MH strategies.

ML is a subfield of Artificial Intelligence (AI) that involves developing algorithms to learn knowledge from data and make predictions on new problems. In recent years, the application of ML techniques in CO became an emerging research area with quite a number of contributions for different purposes. On the one hand, some research has been devoted to develop heuristics for solving CO problems by ML, i.e., to perform "end-to-end learning" to directly generate good solutions for a CO problem. On the other hand, ML has been applied in combination with CO algorithms, where the learning based algorithms have the potential to achieve better performances, either because the current strategy of performing some auxiliary tasks is computationally expensive or because they are poorly understood from the mathematical viewpoint. For a detailed review of "learn to optimize," the interested reader is referred to the survey [4].

Specifically for metaheuristics, ML techniques can be used to infer patterns from data generated from MH processes. Integrating the extracted knowledge into the search strategies can lead MHs to search the solution space more efficiently and significantly improve the current performance. Recently, the application of ML techniques for MHs has attracted increasing research interest and we refer the interested reader to the following surveys [5, 6].

In this paper, we focus on the integration of ML techniques into MHs and propose a general learning-based framework for neighborhood generation in MHs. We first define an efficient neighborhood structure by applying a transformation to a selected subset of variables from the current solution. Then, the problem is to determine how to select a subset of variables that leads to a promising neighborhood of solutions containing a descent of the objective in the solution space. By conducting a classification task, our method learns good variable selection policies from both structural characteristics of the problem and high-quality solutions. We demonstrate the effectiveness of our approach on two applications: a Tabu Search scheme for solving a Wireless Network Optimization problem and a Large Neighborhood Search heuristic for solving MIPs.

2. Background

In this section, we introduce the necessary background and notation.

2.1. Combinatorial optimization

CO problems are a class of optimization problems with a set of decision variables and a defined solution space. Without loss of generality, a CO problem can be formulated into a constrained optimization problem as follows:

$$\min f(\mathbf{x}) \quad (1)$$

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

$$x_i \in \{0, 1\}, \forall i \in \mathcal{B}, \quad (3)$$

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

$$x_k \geq 0, \forall k \in \mathcal{C}, \quad (5)$$

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

Since many CO problems are NP-hard, determining optimal solutions by exact methods requires in the worst case exponential time and might be intractable, especially for large-size applications. In many practical applications where the CO problems are hard and complex, practitioners are often interested in finding good-quality solutions in an "acceptable" amount of computing time rather than solving the problem to optimality. Therefore, heuristic algorithms are developed to efficiently compute high-quality solutions.

2.2. Metaheuristics

Metaheuristics are general framework strategies for designing heuristics for solving CO problems, and provide guidelines to integrate basic heuristic schemes such as local search with high-level diversification strategies. A large part of metaheuristics are built on top of a basic *neighborhood search* (NS) scheme, where NS is a local search procedure that starts from an initial solution \mathbf{x} and iteratively search for improving solutions by exploring a series of neighborhoods.

We consider an instance $p \in \mathcal{P}$ of a CO problem, where \mathcal{P} is the set of problem instances. \mathcal{X} is the solution space, i.e., the set of feasible solutions. The neighborhood is defined as follows. Let \mathbf{x} be a solution of an instance p such that $\mathbf{x} \in \mathcal{X}$. The *neighborhood* $N(\mathbf{x})$ of solution \mathbf{x} is a subset of solutions defined from \mathbf{x} in the solution space \mathcal{X} , i.e., $N(\mathbf{x}) \subseteq \mathcal{X}$.

In general, the structure of the neighborhoods and how these neighborhoods are explored, are designed according to the characteristics of the problem at hand. A neighborhood structure is typically determined by a *transformation* operator that applies a move to the current solution in the solution space.

A transformation operator is typically represented by $\Delta: \mathcal{X} \rightarrow 2^{\mathcal{X}}$, which is a function (or a set of functions) that maps a solution \mathbf{x} to a set of solutions. If the operator Δ is defined with a set M of parameters, the operator can be defined by $\Delta: \mathcal{X} \times M \rightarrow 2^{\mathcal{X}}$. For

```

Input: an initial solution  $x$ ;
 $x^* \leftarrow x$ ;
repeat
   $N(x) \leftarrow \Delta(x)$  (or  $N(x) \leftarrow \Delta(x, m)$  with  $m \in M$ );
   $x' \leftarrow \operatorname{argmin}_{x'' \in N(x)} f(x'')$ ;
   $x \leftarrow x'$ ;
  if  $f(x) < f(x^*)$  then
     $x^* \leftarrow x$ ;
  end
until termination condition is reached;
return  $x^*$ 

```

Algorithm 1. NS-based metaheuristic.

instance, a neighborhood $N(x) \mapsto \Delta(x, m)$ can be constructed by applying the transformation operator with $m \in M$ to the current solution x .

At each NS iteration, the neighborhood $N(x)$ is generated by a $\Delta(x)$ or $\Delta(x, m)$ function. A basic template for NS-based metaheuristic is shown in Algorithm 1.

2.3. Representation learning for CO

In ML, there is a vast library of models for representing CO problems depending on the format of input data of the CO task. For example, the model could be a linear function or some non-linear Artificial Neural Network (ANN) with a set of parameters to be optimized. In Deep Learning (DL), there are many types of neural networks and architectures available for modeling CO problems. For instance, a Multilayer Perceptron (MLP) is the simplest architecture of feedforward neural networks and can be used to model problems with fixed-size input data; Graph Neural Networks (GNNs) are developed to process data that are naturally representable with graphs; Recurrent Neural Networks (RNNs) can be applied to process CO problems with sequential data, etc.

In particular, due to the ubiquity of graph data, problems over graphs arise in numerous application domains. Moreover, given the fact that the vast majority of CO problems have a discrete nature, many of them are naturally described in graphs or can be modeled into a graph structure. For instance, a network optimization problem can be naturally modeled by a graph and a generic MIP instance can also be represented into a bipartite graph [7]. These graphs have inherent structural commonalities and patterns, and there is a potential to exploit valuable graph features to learn patterns from data. Therefore, *graph representation learning* with the application of various GNNs [8–11] has recently emerged as a popular approach for studying CO with a machine learning perspective. Without loss of generality, GNNs learn a graph embedding based on the input features and the graph structure. In a nutshell, higher-level representations of a node or an edge are obtained by kernel convolutions which leverage its local structure.

In the literature, there has been an effort to learn algorithms for solving specific CO problems and many of them apply GNNs as the representation model. On the one hand, the first attempted paradigm was "end-to-end" learning for generating a heuristic

solution by a ML model [12–17]. However, these methods are typically limited to specific CO problems in which a heuristic solution can be easily constructed, and scaling to large-size instances is an issue. On the other hand, since a wide range of constrained CO problems can be formulated into a MIP model, there has also been increasing interest in learning decision rules to improve MIP algorithms [7, 18–22]. While it is shown that this direction has a great potential to improve the state-of-the-art of MIP algorithms, convincing generalization performances, and transfer learning across instances have not been fully tackled yet.

3. Methodology

In this section, we present our framework for learning to generate high-quality solution neighbors for MH search. In Section 3.1 and 3.2, we first introduce an efficient neighborhood structure by applying a transformation to a selected subset of variables from the current solution. Then, in Section 3.3 the problem becomes to determine how to select a subset of variables that leads to a promising neighborhood, i.e., one containing a descent direction of the objective in the solution space. By conducting a classification task, our method learns promising variable selection policies from the structural characteristics of the problem and high-quality solutions.

3.1. Transformation operator and neighbor generation

The definition of the neighborhood plays a critical role in the NS-based metaheuristics. It specifies how the metaheuristic search moves in the solution space. As introduced in Section 2.2, the structure of a neighborhood is typically determined by a transformation operator, since the latter specifies how the current solution is perturbed and transformed to solution neighbors.

For instance, in a classical Traveling Salesman Problem (TSP) where a set of n cities and the distances between each pair of cities are given, the task is to find the shortest tour that visits each city exactly once. Given an arbitrary tour as the initial solution, a simple transformation operator can be defined by firstly removing k edges from the current tour and then adding k other edges to construct a new tour. This is known as the " k -OPT" neighborhood. For a generic CO problem at hand, there are many ways of defining a transformation operator. In general, two main aspects must be taken into consideration: *exploitation* (or *intensification*) and *exploration* (or *diversification*) of the solution space.

On the one hand, the smaller the perturbation induced by the operator, the closer the constructed neighborhood to the current solution. The metaheuristic search thoroughly exploits the local solution regions around the current solution. On the other hand, with more perturbation or more randomness induced by the operator, the neighbors can be defined far from the current solution. The heuristic search has a larger chance to explore more solution regions that have been less visited before. On the extreme case, when an operator completely perturbs the solution or the solution is allowed to be fully changed, the neighborhood could be expanded to the entire solution space, and the complexity

of exploring this neighborhood will be very high, as high as solving the original problem. Instead, the heuristic search becomes totally random when the operator perturbs a part of the solution randomly. Moreover, the design of the transformation operator also strongly depends on the type of problem and its representation. The effectiveness of an operator might not be the same on different types of problems.

In order to guide the metaheuristic search to promising regions of the solution space, we present a general ML framework for neighbor generation. More precisely, we design transformation operators with parameterizations that are able to construct neighborhoods containing high-quality solutions.

3.2. Variable selection for neighbor generation

In this section, we define a neighborhood structure in NS-based metaheuristics. As mentioned before, we aim at generating promising neighbors by efficient transformation operators.

In NS, improving solutions are typically found from neighbors defined by perturbing only a part of the solution. In practice, in order to control the size of the neighborhood, the ratio of perturbation is typically set to a relatively small value. Although this ratio can be increased by diversification strategies when no improving solution is found from the last search, it is still very common that the best solutions found by two consecutive local search iterations share a large part of variable values. Nevertheless, this "partial evolution" of the local optima induces a class of transformation operators for constructing structural neighborhoods, where the transformation is defined on a subset of variables.

Definition 1. A "subset" transformation operator is a function $\Delta : \mathcal{X} \times \mathcal{Y} \rightarrow 2^{\mathcal{X}}$, where \mathcal{X} denotes the solution space, \mathcal{V} is the index set of variables, and $\mathcal{Y} = \{0, 1\}^{|\mathcal{V}|}$ defines a binary decision space for selecting a subset of \mathcal{V} . Consequently, for $y \in \mathcal{Y}$, the neighborhood $N(x)$ of a solution $x \in \mathcal{X}$ is defined as $\Delta(x, y)$, and only variables x_i with $y_i = 1$ can have their value changed by transforming x into a neighbor in $N(x)$.

When the components of y are all set to 1, all the variables in \mathcal{V} will be allowed to change. As discussed before, generating neighbors over the entire variable set \mathcal{V} might result in a large local problem. On the other hand, the "partial evolution" of improving solutions in local search also indicates that it is possible to define efficient transformation operators only on a subset of variables. Moreover, common characteristics are often present in good solutions in many applications. For instance, in a classical TSP problem, two cities that are far from each other are typically disconnected in good solutions, whereas cities close to each other have a higher probability of being connected. Hence, ML techniques can be employed to learn structural information from high-quality solutions and select a subset of variables that has a high probability of defining a neighborhood that contains a descent point of the objective in the solution space.

3.3. Learning a variable selection policy for structural neighbor generation

Given an instance of a generic CO problem and a solution $x \in \mathcal{X}$, we denote its current state as $s \in \mathcal{S}$, where \mathcal{S} is the state space of the problem and typically consists of a set of selected features including x . In order to define a structural neighborhood using the subset transformation operator (defined by Definition 1), an instantiation of \mathcal{Y} is required to select a subset of variables. Hence, a variable selection policy π for selecting $y \in \mathcal{Y}$ can be defined by

$$\begin{aligned} \pi : \mathcal{S} &\longrightarrow \mathcal{Y} \\ s &\longmapsto y. \end{aligned}$$

Now the question is:

How to design a variable selection policy by which the subset transformation operator defines neighborhoods containing high-quality solutions?

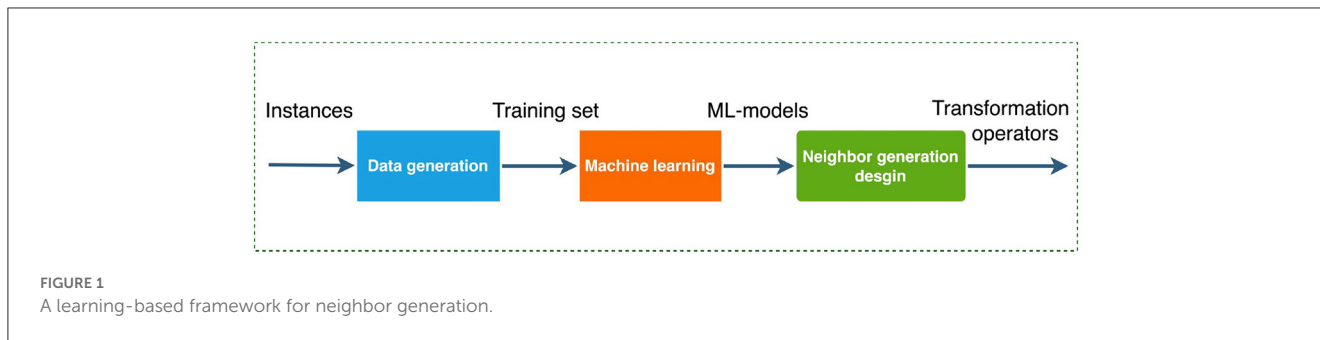
As mentioned before, the high-quality solutions found during the metaheuristic search often share common characteristics and patterns, and more importantly, many improved solutions in the neighborhood often share a part of the variables with the same values. Hence, we propose a general learning-based framework to extract these characteristics from data, and exploit the learned knowledge to guide the metaheuristic search toward compact and promising solution regions. Specifically, the framework learns a variable selection policy and the high-quality neighborhood structure will then be defined by selecting promising variables, the values of which are allowed to be changed from the current solution.

The pipeline of the framework consists of three components: data generation, machine learning and neighbor generation design. The framework is depicted in Figure 1.

In the following, we will instantiate the learning problem for variable selection as a classification task and train a variable selection policy through supervised learning. Within our method, the solution neighbors in metaheuristic search are generated based on our trained model.

3.3.1. Generation of a training data set

Our aim is to learn a policy π that maps the state s of a problem instance, to a high-quality y (referred to as *label* in the following), the binary classification decisions for selecting the best subset of variables that leads to a successful change of the current solution. For this purpose, one typically generates a training data set by applying an algorithm (referred to as *expert*) which explores the neighborhood of some states s_1, \dots, s_N , allowing to modify the values of all the variables, and then detects which are the variables whose value modification results in an improvement. Hence a label y_i is associated with each s_i ($i = 1, \dots, N$) and we thus have a training data set $\mathcal{D}_{\text{train}} = \{(s_i, y_i)\}_{i=1}^N$.



3.3.2. Training the variable selection policy

The policy we want to determine depends on parameters $\theta \in \Theta$ of a model (see below), where Θ is the set of candidate parameters. Each policy is thus denoted π_θ and, given a training data set $\mathcal{D}_{\text{train}} = \{(s_i, y_i)\}_{i=1}^N$, we aim to determine parameters θ^* such that $\pi_{\theta^*}(s_i)$ is as close as possible to y_i ($i = 1, \dots, N$). For this purpose, we solve the following classification problem:

$$\theta^* = \underset{\theta \in \Theta}{\operatorname{argmin}} \sum_{i=1}^N \mathcal{L}(\pi_\theta(s_i), y_i), \quad (6)$$

where $\mathcal{L}(\pi_\theta(s_i), y_i)$ denotes the loss function and will be defined according to the CO application (see more details in Sections 4.3.1, 4.3.2, and 5.2). In summary, given a state $s \in \mathcal{S}$, we use $y = \pi_{\theta^*}(s)$ for selecting which variables are allowed to be changed when moving to neighbors.

3.3.2.1. Modeling

Since CO problems generally have a discrete nature and many of them can be represented by a graph structure, we will apply GNN architectures [8–10] for parametrizing the variable selection policy. More precisely, if the state of a CO instance is modeled by a graph, GNNs can be implemented as the representation model for π_θ . GNNs exhibit some appealing properties for processing data in graph format. First of all, GNNs are size-and-order invariant to input data, i.e., they can process graphs of arbitrary size and topology, and the graph model is invariant to the ordering of the input elements, which brings a critical advantage compared to other neural networks. Moreover, GNNs can exploit the sparsity of the graph by localized propagation of information, making them an ideal class of models for embedding sparse CO problems [7].

The basic architecture of GNNs consists of three modules: the input module, the convolution module, and the output module. In the input module, the state s of the problem is fed into the GNN model. The input module embeds the features of s . The convolution module propagates the features of the graph components by *graph convolution layers* [7, 22] and further embeds the features.

In particular, the architecture of the graph convolution layers used in this paper applies the message passing operator, defined as

$$\mathbf{v}_i^{(h)} = f_\psi^{(h)} \left(\mathbf{v}_i^{(h-1)}, \sum_{j \in \mathcal{N}(i)} g_\phi^{(h)}(\mathbf{v}_i^{(h-1)}, \mathbf{v}_j^{(h-1)}, \mathbf{e}_{j,i}) \right), \quad (7)$$

where $\mathbf{v}_i^{(h-1)} \in \mathbb{R}^d$ denotes the feature vector of node i in layer $(h - 1)$, $\mathbf{e}_{j,i} \in \mathbb{R}^m$ denotes the feature vector of edge (j, i)

from node j to node i , $\mathcal{N}(i)$ denotes the set of nodes adjacent to i in the graph, and $f_\psi^{(h)}$ and $g_\phi^{(h)}$ denote the embedding functions parameterized by ψ and ϕ in layer h and are typically represented by neural networks (e.g., MLP architectures). The output module maps the embedding of the state into a distribution in the binary decision space for each variable and the final output is the class prediction of the variable (typically 1 for being selected into the subset, 0 for not being selected).

3.3.3. Neighbor generation design

After training the classification model for variable selection, the next step is to apply the pretrained model for neighbor generation. It is important to note that, the trained classification model itself is probabilistic, and it maps the current solution state of an instance into a probability distribution in the binary decision space for each variable. We still need to select a strategy to make binary decisions on the variables.

The most straightforward way is to apply a *greedy* strategy. The decision is made by always picking the class with a higher probability and the resulting strategy is *deterministic*. Another way is to sample decisions from the distribution, and hence the strategy is *probabilistic*. In general, the greedy strategy only selects the deterministic subset of variables and exploits the learned knowledge from the training data set. Instead, the probabilistic strategy selects from all the possible neighborhoods with a probability preference defined by the classification model, thus results in a better exploration of the less visited solution regions. There is no guarantee that one strategy is always better than the other. In practice, they can be combined. For each task, one should search for a good trade-off between exploitation of the local solution regions and exploration of the global solution space.

Finally, the template for a NS-based metaheuristic guided by the variable selection policy π is given in Algorithm 2. The neighborhood structure $\mathcal{N}(x)$ is defined by the subset transformation operator that selects a subset of variables according to policy π_{θ^*} , and modifies some of their values to lead the heuristic search to promising solution regions.

4. Application 1: tabu search in wireless network optimization

In this section, we apply our framework to the first case study, a Wireless Network Optimization (WNO) problem. We demonstrate

```

Input: an initial solution  $x$  and a variable selection policy  $\pi_{\theta^*}$ ;
 $x^* \leftarrow x$ ;
repeat
   $y \leftarrow \pi_{\theta^*}(s)$  where  $s$  is the state associated with  $x$ ;
   $N(x) \leftarrow \Delta(x, y)$ ;
   $x' \leftarrow \operatorname{argmin}_{x'' \in N(x)} f(x'')$ ;
   $x \leftarrow x'$ ;
  if  $f(x) < f(x^*)$  then
     $x^* \leftarrow x$ ;
  end
until termination condition is reached;
return  $x^*$ 
    
```

Algorithm 2. NS-based metaheuristic guided by the variable selection policy.

the effectiveness of our learning-based framework by generating solution neighbors in a Tabu Search scheme.

4.1. The tactical WNO problem

When telecommunications are necessary but standard networks are unavailable, as in the case of disaster relief operations, small temporary wireless networks, called tactical networks, are set up. These typically connect between 10 and 50 nodes (the key locations that must communicate) in a single network. The design of such networks can be optimized such that the network’s weakest link is maximized. This, in turn, guarantees that all nodes can receive important information in a timely manner.

Tactical wireless network design is a complex non-linear combinatorial optimization problem that includes three sub-problems: the design of the topology (P_0), the configuration of the network (P_1), and the configuration of the antennas (P_2). These sub-problems are nested such that P_1 is defined given a topology, and P_2 is defined given a topology and a network configuration, namely

$$\underbrace{\max_{\text{topology } t \in T}}_{P_0} \quad \underbrace{\max_{\text{network configuration}}}_{P_1} \quad \underbrace{\max_{\text{antenna configurations}}}_{P_2} \quad o.$$

where T is the space of valid topologies and o is the objective function of the problem. The problem has been proposed by an industrial partner and the modeling of the full tactical wireless network design optimization problem can be found in Perreault [23].

Given a set of nodes V , the topology $t \in T$ can be any undirected tree (V, E) that describes how information travels in the network between every pair of nodes. Each edge in the topology represents a direct connection between two antennas in the network. A network configuration selects a root node (also known as a master hub or gateway) and assigns waveforms and channels/frequencies to the edges. The waveforms are the communication protocols that depend on the local structure of the edges and the channels and frequencies characterize their radio signals. Together with the antenna configurations, they determine which edges interfere with each other. An antenna configuration

requires to define an angular alignment with respect to the azimuth as well as a set of activated beams, in the case of multi-beam antennas. Given all these properties, the radio signals can be physically modeled by also taking into account the path losses and fade margins of the terrain between every pair of nodes. The data transmission speed (called direct throughput) TP_{uv} for all the edges $[u, v] \in E$ can then be computed and, depending on the traffic scenario X and its distribution of congestion n_{uv}^X in the edges, their effective throughputs TP_{uv}/n_{uv}^X can also be computed. There are three traffic scenarios: scenario A is a single communication between any two nodes, scenario B is all nodes communicating with the root node, and scenario C is all nodes communicating with each other. For all traffic scenarios, the congestion of any directed edge (u, v) can be computed as a function of the number of descendants desc_v , defined by the root node selected in the network configuration.

To evaluate a single topology $t \in T$, problem P_1 , which is itself a complex combinatorial optimization problem, must be solved, and solving it requires solving P_2 many times as well. In turn, problem P_2 can be efficiently solved by a simple geometrically-based heuristic. Solving P_1 to optimality every time that a topology needs to be evaluated is too costly, especially in a NS context. It can be approximated in such a way that it can be solved efficiently by exhaustive enumeration. We denote the resulting approximated objective function by $\bar{f}(t)$. However, this approximation does not necessarily provide a feasible network because it does not define a proper frequency assignment. Alternatively, P_1 can also be tackled directly by considering greedy frequency assignments which do provide feasible networks, although this takes considerably more time. We denote this final objective function by $f(t)$.

Given such procedures for solving P_1 and P_2 , P_0 , given by

$$\max_{t \in T} f(t),$$

can be solved with a NS-based MH in the space of topologies, where local directions of descent are computed with $\bar{f}(t)$.

4.2. Topology tabu search

4.2.1. Neighborhoods

The neighborhood structure dictates how the local search can move in the space of solutions. Our neighborhood N is the edge-swap neighborhood. To construct the neighbor topologies, for each tree edge $e \in E$ in the current topology, we remove it, which disconnects the topology into two connected components, and we consider every possible way of reconnecting these two connected components with another edge. Hence, $\Delta(t)$ is the set of all topologies that can be obtained from t by an edge-swap move.

An example of neighbor topologies is depicted in Figure 2, where the edges in blue were “swapped”.

4.2.2. Tabu search

Tabu Search (TS) is a MH that explores the solution space by temporarily not allowing to move in the direction from which it came. The only case in which this rule does not apply is when moving in such a direction improves the incumbent solution. This

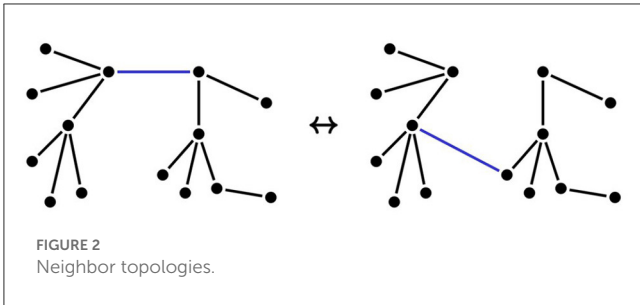


FIGURE 2 Neighbor topologies.

```

Input: an initial topology  $t$ ;
 $t^* \leftarrow t; \bar{t}^* \leftarrow t; L_{drop} \leftarrow \emptyset; L_{add} \leftarrow \emptyset;$ 
repeat
   $N(t) \leftarrow \Delta(t);$ 
   $t' \leftarrow \operatorname{argmax}_{t'' \in N(t)} \bar{f}(t'')$ 
  s.t.  $t''$  is not tabu or  $\bar{f}(t'') > \bar{f}(\bar{t}^*);$ 
  if  $\bar{f}(t') > \bar{f}(\bar{t}^*)$  then
     $\bar{t}^* \leftarrow t';$ 
  end
  Make move  $t \leftarrow t'$  and update tabu lists  $L_{drop}$  and  $L_{add}$ ;
  if  $f(t) > f(t^*)$  then
     $t^* \leftarrow t;$ 
  end
until termination condition is reached;
return  $t^*$ 
    
```

Algorithm 3. P_0 Topology Tabu Search.

is done by using tabu lists that record the opposites of the last few moves, which then become tabu for the next few moves.

For our edge-swap neighborhood, two tabu lists are necessary: a list L_{drop} for forbidding edges to be dropped and a list L_{add} for forbidding edges to be added. Let $t' \in N(t)$ be a neighbor of the current topology t , obtained by replacing an edge e with an edge e' : we consider t' as tabu if $e \in L_{drop}$ or (non exclusive) $e' \in L_{add}$. When a move is actually made from the current topology, the newly dropped edge is added to L_{add} (it cannot be added back for the next few moves) and the newly added edge is added to L_{drop} (it can not be dropped for the next few moves).

The lengths of the tabu lists determine for how many moves the dropped/added edges have tabu status. As observed in Perreault [23], reasonable lengths for L_{drop} and L_{add} are $\lfloor (\sqrt{|V|} - 1)/2 \rfloor$ and $\lfloor \sqrt{|V|(|V| - 1)/2} \rfloor$, respectively, where $\lfloor x \rfloor$ means x rounded to the nearest integer. The pseudocode for the resulting *topology Tabu Search* metaheuristic is given in Algorithm 3.

4.3. Learning to generate edge-swap neighbors for TS

As a NS-based MH, TS can be roughly described as a NS scheme plus a "tabu" strategy for preventing cycling by keeping a short-term memory of visited solutions stored in the tabu list. As discussed above, the topology design for the wireless network optimization problem can be solved by applying a topology TS

algorithm (Algorithm 3). In this TS scheme, a neighborhood $N(t)$ of the current solution t is constructed by applying an edge-swap move operator. Specifically, the move operator consists of two steps: an edge will be dropped from the current topology (by enforcing the value of the dropped edge variable from 1 to 0) and another edge will be added to complete a new topology (by enforcing the value of the added edge variable from 0 to 1). As a result, the edge-swap neighborhood consists of all the edge-swap moves.

In Algorithm 3, the neighborhood search at each TS iteration explores the entire edge-swap neighborhood by enumerating all possible moves. Although the algorithm is guaranteed to find a locally optimal solution by exploiting the entire edge-swap neighborhood, it is generally slow in terms of computing time since a complex subproblem (P_1) has to be solved to evaluate each "edge-swap" move, and another subproblem (P_2) has to be solved to evaluate each solution in P_1 .

A potential improvement of the enumeration strategy is to apply a random strategy to sample from droppable edge variables and addable edge variables, thus evaluating only a subset of possible moves. The random strategy generally explores a larger solution space than the enumeration strategy because, if the overall amount of time for the algorithm is fixed, it is able to do more iterations. However, random sampling might not be efficient enough for guiding the search toward high-quality neighborhoods. To achieve a better trade-off between exploitation and exploration, we propose to exploit structural characteristics of improving moves, and learn good variable selection policies for selecting "drop" edge variables and "add" edge variables to generate size-reduced, but high-quality edge-swap neighborhoods. The scheme for variable selection consists of two components: a classifier for selecting the "drop" edge variables and another classifier for selecting the "add" edge variables. The scheme is depicted in Figure 3 and a detailed pseudocode will be given in Algorithm 4.

4.3.1. Learning to drop edges

Let \mathcal{S} denote the state space of the topologies of a WNO instance, and let $\mathcal{Y}^d = \{0, 1\}^{|\mathcal{I}|}$ denote the set of possible binary decisions for the droppable edge variables, where \mathcal{I} is the index set of droppable edges in the current topology. We aim to learn a "drop-edge" policy that maps each state $s \in \mathcal{S}$ to a label $\pi_{\theta^*}^d(s) = y^d \in \mathcal{Y}^d$. The generation of a training dataset is explained in Section 4.4.1. Roughly speaking, we consider a set J of WNO instances, and apply T_j iterations of Algorithm 3 to each instance $j \in J$. At each iteration of each instance, we determine the edges that can be dropped and replaced by another edge to produce a better topology. This gives a label for every visited state. Let $\{s_i^{(j)}\}_{i=1}^{T_j}$ be the states of the j^{th} instance, and let $\{y_i^{d(j)}\}_{i=1}^{T_j}$ be the corresponding labels. As explained in Section 3.3.2, the policy $\pi_{\theta^*}^d$ is obtained by solving the following problem :

$$\theta^* = \operatorname{argmin}_{\theta \in \Theta} \sum_{j=1}^{|J|} \sum_{i=1}^{T_j} \mathcal{L} \left(\pi_{\theta}^d(s_i^{(j)}), y_i^{d(j)} \right), \tag{8}$$

where \mathcal{L} is the loss function.

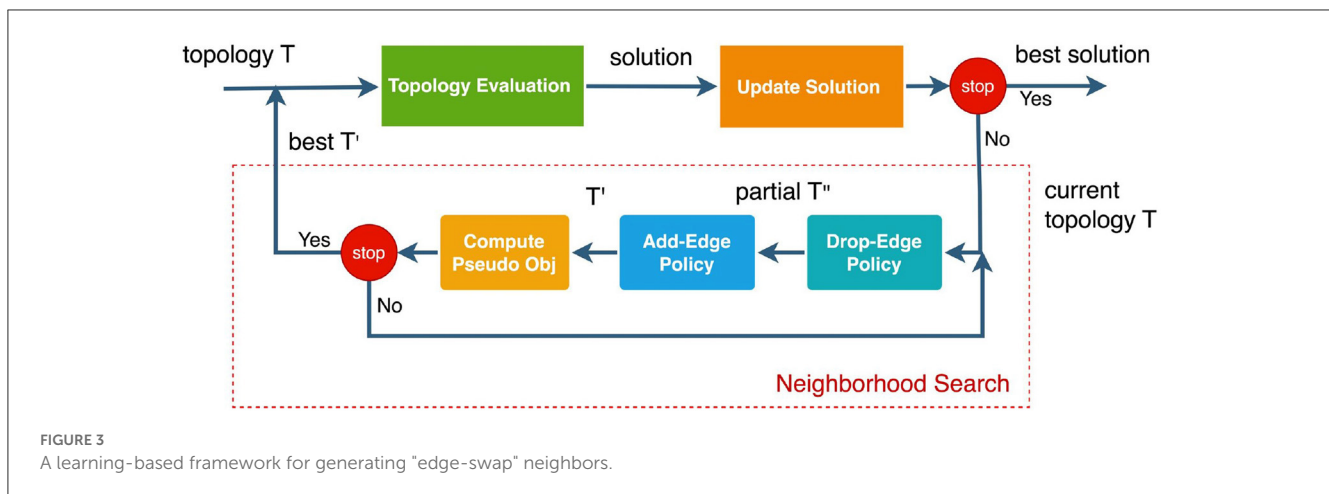


FIGURE 3 A learning-based framework for generating "edge-swap" neighbors.

4.3.1.1. Feature design

Since the topologies naturally have a graph structure, we represent each state s by a graph. More precisely, given a WNO instance with n nodes, we consider d features for each node and e features for each edge. The matrix of node features is denoted by $V \in \mathbb{R}^{n \times d}$ and the tensor of edge features is denoted by $E \in \mathbb{R}^{n \times n \times e}$. The features used for learning which edges have the best potential to lead to an improvement when dropped are listed in Table 1¹.

4.3.1.2. GNN model

Because of the intrinsic graph structure of the topologies, it is natural to apply GNN to model the drop-edge classifier. The chosen GNN architecture has 3 modules: the input module, the convolution module and the output module. The output module embeds the hidden features extracted from the convolution module and maps the embedding of each edge in the current topology into a two-neuron output.

4.3.1.3. Loss function

From the data generation process, we observed that class distribution is unbalanced and only $< 50\%$ of droppable edges lead to an improving move. Since the objective of the learning is to select as many "improving" edges to be dropped as possible, it is reasonable to put more effort in improving the predictions on the minority class. Therefore, we apply a *weighted* Cross Entropy (WCE) loss to train the model. Specifically, we add a penalty factor λ for the "improving" class and $1 - \lambda$ for the "non-improving" class. Formally, the WCE loss is defined as

$$\mathcal{L}_{wce}(\hat{y}, y) = \frac{1}{|\mathcal{I}|} \sum_{i=1}^{|\mathcal{I}|} (-\lambda y_i \log \hat{y}_i - (1 - \lambda)(1 - y_i) \log(1 - \hat{y}_i)), \tag{9}$$

where \hat{y} denotes the prediction of probability given by the policy model, y denotes a label, and \mathcal{I} is the index set of droppable edges in the current topology. The penalty factor $\lambda \in [0.5, 1]$ imposes a larger loss for the "improving" class during training.

¹ More detailed definitions of the listed features can be found in Perreault [23].

4.3.2. Learning to add edges

After dropping an edge, the next step is to select an edge to be added from all addable edges. Let \mathcal{S} denote the state space of all graphs obtained from a topology by dropping an edge, and let $\mathcal{Y}^a = \{0, 1\}^{|\mathcal{I}|}$ be the set of possible binary decisions for addable edges, where \mathcal{I} is the index set of addable edges. We aim to learn an "add-edge" policy $\pi_{\theta^*}^a$ that maps each state $s \in \mathcal{S}$ to a label $y^a \in \mathcal{Y}^a$.

The generation of a training dataset is similar to what was done for dropping edges. Algorithm 3 is applied to a set J of WNO instances for a few iterations in order to determine which edges have the potential to improve a topology once one of their edges has been dropped. More details are given in Section 4.4.1. Let $\{s_i^{(j)}\}_{i=1}^{T_j}$ be the states of the j^{th} instance, and let $\{y_i^{a(j)}\}_{i=1}^{T_j}$ be the corresponding labels. The policy $\pi_{\theta^*}^a$ is obtained by solving the following problem :

$$\theta^* = \operatorname{argmin}_{\theta \in \Theta} \sum_{j=1}^{|J|} \sum_{i=1}^{T_j} \mathcal{L}(\pi_{\theta}^a(s_i^{(j)}), y_i^{a(j)}), \tag{10}$$

4.3.2.1. Feature design

As was done for dropping edges, we model a state s by a graph. Given a WNO instance with n nodes, d features for each node, and e features for each edge, we denote by $V \in \mathbb{R}^{n \times d}$ the matrix of node features and by $E \in \mathbb{R}^{n \times n \times e}$ the tensor of edge features. The features used for learning which edges have the best potential to lead to an improvement when added are listed in Table 2.

4.3.2.2. GNN model

As for the drop-edge classifier, we also apply GNN to model the add-edge classifier for selecting the edge to be added. The GNN architecture is the same as for the drop-edge classifier.

4.3.2.3. Loss function

The class distribution for addable edges is even more unbalanced than for the droppable case. Actually, only $< 10\%$ of addable edges lead to an improving move. Therefore, we also applied the WCE loss to train the model with an emphasis on the minority class.

4.4. Numerical experiments

This section contains the experimental results for the WNO application. After presenting the data collection, in Sections 4.4.2 and 4.4.3, we discuss the experimental setting and the evaluation metrics, respectively. Finally, the results are reported and discussed in Section 4.4.4.

4.4.1. Data collection

4.4.1.1. Problem instance generation

For the numerical experiments, the instances were generated in the following way. First, independent coordinates for the explicit nodes $v \in V$ are iteratively generated using

$$(x_v, y_v) = (\sqrt{U_0} \cos(2\pi U_1), \sqrt{U_2} \sin(2\pi U_3)) \quad (11)$$

where $U_0, U_1, U_2, U_3 \sim \mathcal{U}(0, 1)$ are independent and identically distributed (IID). These coordinates are then scaled to match the average distance ratio of 10 km. This coordinate generation is repeated until the minimum distance is above 2 km and the maximum distance is below 150 km. Once this is achieved, for each possible edge $[u, v]$ with $u, v \in V$, a random uniform variable is sampled for the path loss and another is sampled for the fade margin, both according to empirical cumulative distribution functions derived from North-American datasets.

We generate small instances of 10 nodes, as well as larger instances of 30 nodes. For these instances, the TS MH is initialized with a minimum spanning tree with a specific distance based on the path losses and fade margins of the terrain between every pair of nodes [23].

4.4.1.2. Training data generation

To collect data for training the drop-edge and add-edge classifiers, we call Algorithm 3 to evaluate all possible moves in the edge-swap neighborhood, and compute labels according to the learning task. In particular, given a WNO instance, and an initialization of the topology, we execute Algorithm 3 as follows: for each TS iteration, we first call the TS algorithm to evaluate all possible edge-swap moves. For each droppable edge, we evaluate all the addable edges. If the best resulting edge-swap move within the neighborhood of the droppable edge leads to a better approximated pseudo-objective, then the droppable edge is labeled as "improving;" otherwise, it is labeled as "non-improving." Given the dropped edge, the label of each addable edge is also decided by the quality of the resulting edge-swap move. If the move leads to a better approximated pseudo-objective, then the addable edge within the neighborhood of the corresponding dropped edge is labeled as "improving;" otherwise, it is labeled as "non-improving."

4.4.1.3. TS guided by GNNs

The TS algorithm with GNN classifiers (TS-GNN) is obtained by using the learned policies $\pi_{\theta^*}^d$ and $\pi_{\theta^*}^a$ to generate edge-swap neighbors. Hence, the pseudocode of TS-GNN is the same as Algorithm 3, except that the neighborhood $N(t)$ of t results from the subset transformation operator that performs edge-swaps according to $\pi_{\theta^*}^d$ and $\pi_{\theta^*}^a$. A detailed pseudocode is given in Algorithm 4.

Input: an initial topology t , a variable selection policy $\pi_{\theta^*}^d$ for dropping edges, and a variable selection policy $\pi_{\theta^*}^a$ for adding edges;

$t^* \leftarrow t; \bar{t}^* \leftarrow t; L_{drop} \leftarrow \emptyset; L_{add} \leftarrow \emptyset;$

repeat

$N(t) \leftarrow \emptyset$

$y^d \leftarrow \pi_{\theta^*}^d(\mathbf{s})$ where \mathbf{s} is the state associated with t ;

for every e_i in t such that $y_i^d = 1$ **do**

set s_i equal to the state obtained from \mathbf{s} by dropping e_i from topology t ;

$y^a \leftarrow \pi_{\theta^*}^a(s_i)$;

for each reconnecting edge e_j such that $y_j^a = 1$ **do**

set $t_{i,j}$ equal to the topology obtained from t by dropping edge e_i and adding edge e_j ;

$N(t) \leftarrow N(t) \cup \{t_{i,j}\}$;

end

end

$t' \leftarrow \operatorname{argmax}_{t'' \in N(t)} \bar{f}(t'')$

s.t. t'' is not tabu or $\bar{f}(t'') > \bar{f}(\bar{t}^*)$;

if $\bar{f}(t') > \bar{f}(t^*)$ **then**

| $\bar{t}^* \leftarrow t'$;

end

Make move $t \leftarrow t'$ and update tabu lists L_{drop} and L_{add} ;

if $f(t) > f(t^*)$ **then**

| $t^* \leftarrow t$;

end

until termination condition is reached;

return t^*

Algorithm 4. P_0 Topology Tabu Search with GNN classifiers.

4.4.2. Experimental setting

4.4.2.1. Training

We train the drop-edge policy and add-edge policy on the collected training dataset separately. The training dataset is generated from 30 small instances of 10 nodes. The collected data is split into training (70%), validation (10%), and test (20%) sets.

4.4.2.2. Evaluation

We evaluate the compared algorithms listed in Section 4.4.4 on two evaluation sets. The first evaluation set contains 50 small instances of 10 nodes. In addition, to evaluate the generalization performance of our approach on larger instances, we also evaluate the model trained on small instances on a large evaluation set with 5 instances of 30 nodes.

4.4.2.3. Experimental environment

Our code is written in Python 3.9 and we use Pytorch 1.60 [24], Pytorch Geometric 1.7.0 [25] for implementing and training the GNNs.

4.4.3. Evaluation metrics

We use the *primal integral* [26] to measure the performance of the compared algorithms. The *primal integral* was originally proposed to measure the performance of primal heuristics for solving mixed-integer programs. The metric takes into account

TABLE 1 Description of the features in the "droppable" graph (N, E).

| Tensor | Feature | Description |
|--------|--------------|---|
| V | coordinates | X and Y coordinates in km |
| | desc_v | Normalized desc _v = ((desc _v - 1)/(V - 1)) |
| E | pass_loss | Path loss in dB |
| | fade_margin | Fade margin in dB |
| | waveform | Binary (point-to-point or point-to-multipoint) |
| | channel | Frequency channel |
| | n_beams | Number of beams used by the antenna of the predecessor |
| | tp_uv_a | Throughput scenario A |
| | tp_uv_b | Throughput scenario B |
| | tp_uv_c | Throughput scenario C |
| | tp_uv_m | Throughput by mixed scenarios of A, B, C |
| | edge indices | Edge indices in current topology |

TABLE 2 Description of the features in the "addable" graph (N, E).

| Tensor | Feature | Description |
|--------|--------------|--|
| V | coordinates | X and Y coordinates in km |
| | desc_v | Normalized desc _v = ((desc _v - 1)/(V - 1)) |
| E | edge_type | Binary indicator of addable edges |
| | pass_loss | Path loss in dB |
| | fade_margin | Fade margin in dB |
| | edge indices | Edge indices in current topology |

both the quality of solutions and the computing time spent to find those solutions during the solving process. To define the primal integral, we first consider a *scaled primal gap function* $p(t)$ as a function of time, defined as

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

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

$$\bar{\gamma}(\tilde{\mathbf{x}}) = \frac{|f(\tilde{\mathbf{x}}_{\text{opt}}) - f(\tilde{\mathbf{x}})|}{|f(\tilde{\mathbf{x}}_{\text{opt}}) - f(\tilde{\mathbf{x}}_{\text{init}})|},$$

where $f(\tilde{\mathbf{x}})$ denotes the objective value given solution $\tilde{\mathbf{x}}$, $\tilde{\mathbf{x}}_{\text{opt}}$ is either the optimal solution or the best one known for the instance and $\tilde{\mathbf{x}}_{\text{init}}$ is the initial solution. The original unscaled *primal gap* is defined as

$$\gamma(\tilde{\mathbf{x}}) = \frac{|f(\tilde{\mathbf{x}}_{\text{opt}}) - f(\tilde{\mathbf{x}})|}{|f(\tilde{\mathbf{x}}_{\text{opt}})|}.$$

Let $t_{\text{max}} > 0$ be the time limit for executing the heuristic. The primal integral measure is then defined as

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

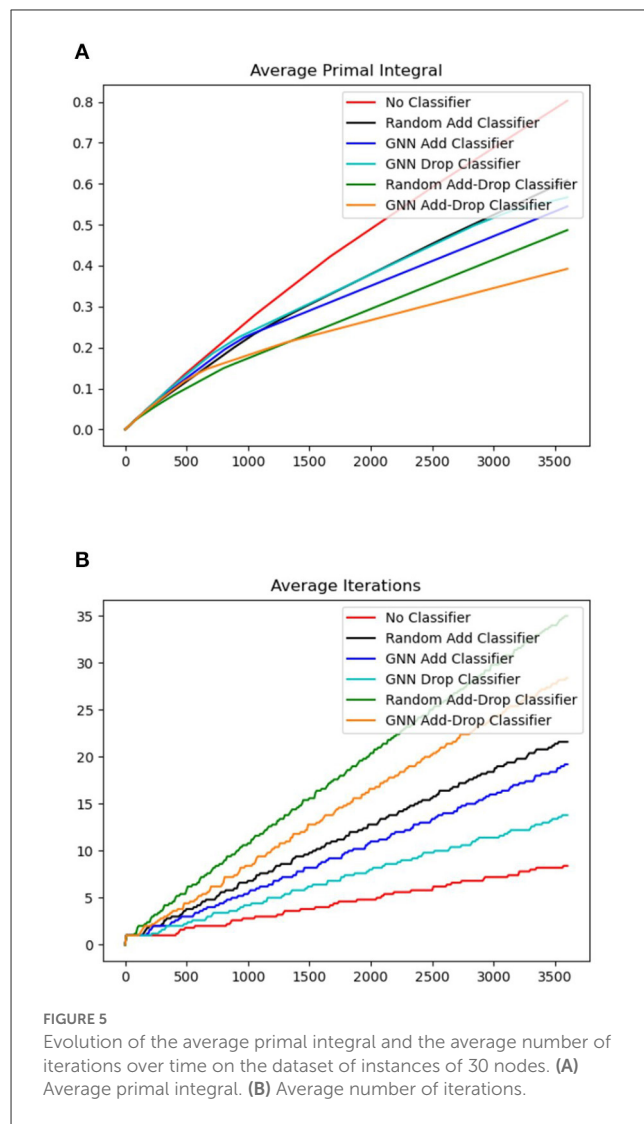
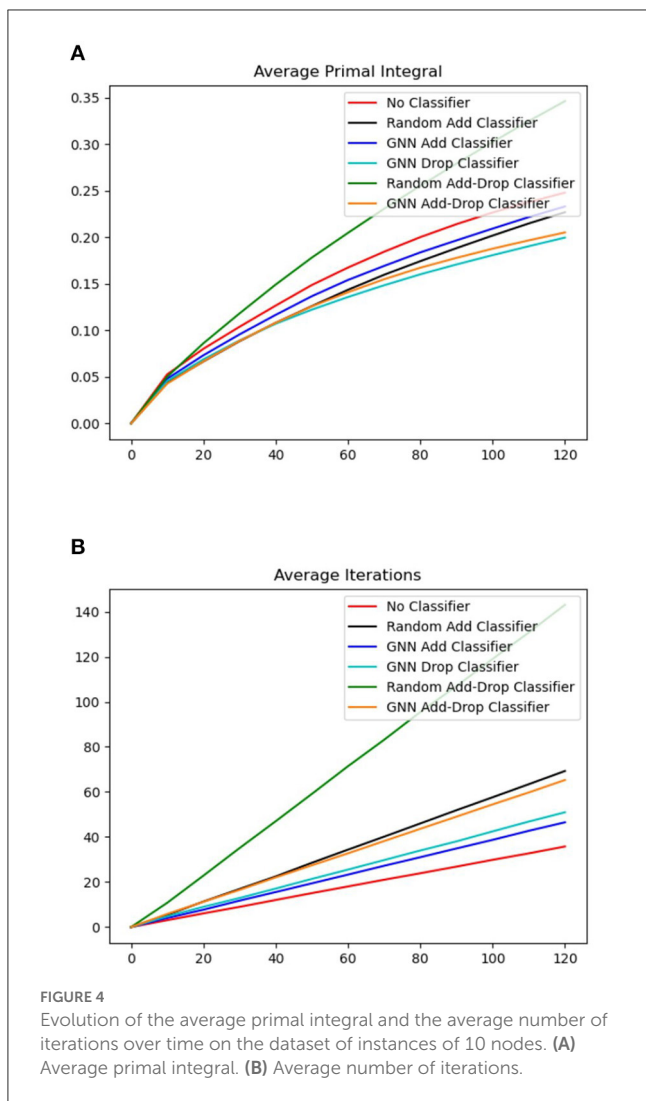
4.4.4. Results

In order to validate our approach, we compare multiple versions of our *TS-GNN* algorithm with the baselines. Specifically, we compare the following algorithms:

- TS with *No-Classifier*, the baseline topology TS algorithm with an enumeration strategy that evaluates all the possible edge-swap neighbors;
- TS with *Random-Add-Classifier*, the baseline topology TS algorithm with a sampling strategy that randomly selects edges to be added for generating edge-swap neighbors;
- TS with *Random-Add-Drop-Classifier*, the baseline topology TS algorithm with a sampling strategy that randomly selects both the edges to be dropped and the edges to be added for generating edge-swap neighbors;
- TS with *GNN-Add-Classifier*, the TS algorithm plus the GNN classifier for selecting the edges to be added for generating edge-swap neighbors;
- TS with *GNN-Drop-Classifier*, the TS algorithm plus the GNN classifier for selecting the edges to be dropped for generating edge-swap neighbors;
- TS with *GNN-Add-Drop-Classifier*, the TS algorithm plus the GNN classifiers both for selecting the edges to be dropped and for selecting the edges to be added.

We evaluate the performance of all the compared algorithms on both small instances with 10 nodes and larger instances with 30 nodes. For each test set, we compute the average primal integral defined in Section 4.4.3 as well as the average number of iterations. The primal integral is our main performance metric for evaluating metaheuristics and it measures the speed of convergence of the objective over the entire search time (the smaller, the better). In addition, the number of iterations counts the number of neighborhoods explored by the algorithm. With the same running time, it reflects how much exploration is guaranteed by modifications made to the basic TS scheme. Indeed, larger number of iterations for the same amount of time indicates a faster moving between neighborhoods and is generally preferable, although the outcome of the search still depends on the quality of the neighborhoods. The evaluation results of the compared algorithms are shown in Figures 4, 5 for instances with 10 and 30 nodes, respectively. Figure 4 plots the evolution of primal integral of the compared algorithms over the running time and Figure 5 plots the evolution of number of TS iterations over time. Table 3 reports more statistical details about the final primal integral when the compared TS algorithms reach the running time limit.

From these results, we observe that all our ML-based TS algorithms with GNN classifiers outperform the original TS baseline with no classifier. The *No-Classifier* baseline employs an enumeration strategy that evaluates all the possible moves in the edge-swap neighborhood. Although it is guaranteed to find the best move at each TS iteration, its more extensive exploitation results in the lowest number of iterations over all the compared algorithms. For small instances of 10 nodes, the performance of *No-Classifier* is still decent, as the solution space of these instances is small for exploration. However, it gives the worst results compared to all other algorithms on larger instances



because the low number of iterations (slow exploitation) and the resulting lack of exploration becomes the main bottleneck of the algorithm.

The algorithms with random classifiers, on the contrary, are more efficient in terms of number of iterations. The *Random-Add-Classifier* algorithm explores a random subset of addable edges in the edge-swap neighborhood and the *Random Add-Drop Classifier* generates an even smaller neighborhood by sampling from both droppable and addable edges. On small instances with 10 nodes, *Random-Add-Classifier* achieves a better performance with a lower primal integral although the *Random-Add-Drop-Classifier* algorithm has a larger number of iterations. This is because, for small instances, the quality of the solutions in the sampled neighborhoods is more important than the number of iterations. Whereas for larger instances of 30 nodes, *Random-Add-Drop-Classifier* performs better than *Random-Add-Classifier* since the sampling efficiency becomes more important for exploring larger edge-swap neighborhoods.

On the one hand, the *No-Classifier* baseline can fully exploit each edge-swap neighborhood (*exploitation*) but has a low efficiency in terms of number of iterations. The baselines with random classifiers are more effective in increasing the number of iterations by reducing the size of each neighborhood (*exploration*), however, the quality of the neighborhood is restricted by its sampling efficiency. On the other hand, our complete ML-based algorithm with *GNN-Drop-Add-Classifier* achieves a smaller primal integral than all the baseline algorithms with a reasonably large number of iterations. Moreover, as the size of instances increases, the *GNN-Drop-Add-Classifier* algorithm becomes more competitive and significantly outperforms all the compared algorithms. These results demonstrate that our approach offers a good trade-off in both exploration and exploitation of the solution space. Since the GNN classifiers are only trained with data generated from small instances, the results also show that our method generalizes well on larger instances.

TABLE 3 Statistics on final primal integral [mean, standard deviation (SD), maximum, and minimum] over instances from the evaluation set of 10 and 30 nodes.

| Algorithm | Ten nodes | | | | Thirty nodes | | | |
|----------------------------|--------------|-------|-------|-------|--------------|-------|-------|-------|
| | Mean | SD | Max | Min | Mean | SD | Max | Min |
| No classifier | 0.248 | 0.218 | 0.829 | 0.028 | 0.856 | 0.075 | 1.000 | 0.788 |
| Random add classifier | 0.227 | 0.228 | 0.735 | 0.017 | 0.574 | 0.179 | 0.727 | 0.222 |
| GNN add classifier | 0.233 | 0.226 | 0.865 | 0.023 | 0.701 | 0.22 | 1.000 | 0.332 |
| Random Add-drop classifier | 0.346 | 0.257 | 0.845 | 0.011 | 0.534 | 0.316 | 0.964 | 0.063 |
| GNN Add-drop classifier | 0.205 | 0.206 | 0.740 | 0.012 | 0.475 | 0.284 | 0.966 | 0.242 |

The bold value in the mean column indicates the best primal integral value over all the compared algorithms.

5. Application 2: large neighborhood search in MIP

In this section, we explain how to apply our methodology to another application: a Large Neighborhood Search (LNS) heuristic for solving MIPs.

5.1. Large neighborhood search

LNS is a refinement heuristic, i.e., given an initial solution, it is applied to improve the solution by exploring a "large" neighborhood. There are several ways of describing a LNS scheme. We adopt the following simple one based of three building blocks:

- *destroy* function d : fixes the values of a subset of variables to the current solution x and "destroy" the rest. The output of this function is a sub-MIP with a neighborhood $N(x)$. The size of the LNS neighborhood, i.e. the number of variables to be destroyed, will be selected as a hyperparameter;
- *repair* function r : rebuilds the destroyed solution², typically by solving a sub-MIP defined by $N(x)$;
- *accept* function a : decides whether the new solution should be accepted or rejected.

Given as an input a feasible solution \bar{x} , LNS searches for the best feasible solution \tilde{x} in the neighborhood of \bar{x} (the size of the neighborhood is a parameter) and then updates $\bar{x} \leftarrow \tilde{x}$. This is repeated until a termination condition is met. The pseudocode of the scheme is given in Algorithm 5.

5.2. Learning to generate neighbors for LNS

In a LNS heuristic, the solution neighbors are obtained by the *destroy* function that selects a subset of integer variables to be "destroyed," thus freed for neighbor generation. The remaining integer variables will be fixed to their values in the current solution. The classic LNS algorithm applies a randomized sampling strategy or hand-crafted rules for defining the destroy function (see, e.g.,

```

Input: an initial solution  $x$ ;
 $x^* \leftarrow x$ ;
repeat
   $x' \leftarrow r(d(x))$ ;
  if  $a(x', x)$  then
     $x \leftarrow x'$ ;
  end
  if  $f(x) < f(x^*)$  then
     $x^* \leftarrow x$ ;
  end
until termination condition is reached;
return  $x^*$ 
    
```

Algorithm 5. The LNS heuristic.

[27]). However, our observation shows that only a small subset of variables have the potential to improve the current solution by changing their values and such a subset is strongly dependent to the structure of the problem. Hence, in order to optimize the performance of LNS, we aim at learning new variable selection strategies to select a subset of variables to be freed. In particular, we investigate the dependencies between the state of the problem, defined by a set of both static and dynamic features collected from the LNS procedure (e.g., context of the problem, incumbent solution) and the binary decisions about integer variables to be either freed or fixed.

We now show how to use a GNN model for learning a destroy policy. The LNS neighborhoods will then be defined based on our pretrained GNN classifier.

Let \mathcal{S} denote the state space of a MIP instance, and let $\mathcal{Y} = \{0, 1\}^{|\mathcal{I}|}$ denote all candidate subsets of variables to be destroyed, where \mathcal{I} is the index set of integer variables. We aim to learn a destroy policy π_{θ^*} that maps each state s to a label $y \in \mathcal{Y}$.

To generate a training dataset, we consider a set J of instances, and apply T_j iterations of the local branching heuristic [28] to each instance $j \in J$. We thus determine improving subsets of variables which give us labels. More details about this process can be found in Section 5.3.1. Let $\{s_i^{(j)}\}_{i=1}^{T_j}$ be the states of the j^{th} instance, and let $\{y_i^{(j)}\}_{i=1}^{T_j}$ be the corresponding labels. As explained in Section 3.3.2, the policy π_{θ^*} can be obtained by solving the following classification problem:

$$\theta^* = \operatorname{argmin}_{\theta \in \Theta} \sum_{j=1}^{|J|} \sum_{i=1}^{T_j} \mathcal{L}(\pi_{\theta}(s_i^{(j)}), y_i^{(j)}). \tag{12}$$

² In some cases, the repaired solution can be worse than the destroyed solution.

TABLE 4 Description of the features in the bipartite graph $s = (V, C, E)$.

| Tensor | Feature | Description |
|--------|---------|--|
| V | sol_val | Solution value. |
| C | bias | Right-hand side value of the constraint. |
| E | coef | Constraint coefficient. |

5.2.1. Feature design

We represent each state s of a MIP by a bipartite graph (V, C, E) [7]. More precisely, assume that the considered MIP has n variables, m constraints, d features for the variables and q features for the constraints. The variables of the MIP with their feature matrix $V \in \mathbb{R}^{n \times d}$ are represented on one side of the graph. On the other side are nodes corresponding to the constraints with $C \in \mathbb{R}^{m \times q}$ being their feature matrix. A variable node i and a constraint node j are connected by an edge (i, j) if variable i appears in constraint j of the MIP. Finally, $E \in \mathbb{R}^{n \times m \times e}$ denotes the tensor of edge features, with e being the number of features for each edge. The features in the bipartite graph are listed in Table 4.

5.2.2. GNN model

Given that the state of a MIP instance can be represented as a bipartite graph, we propose to use GNN to parameterize the model for the destroy policy. Our GNN architecture again consists of 3 modules: the input module, the convolution module, and the output module. For a bipartite graph, a convolution layer is decomposed into two half-layers: one half-layer propagates messages from variable nodes to constraint nodes through edges, and the other one propagates messages from constraint nodes to variable nodes. We refer the reader to Gasse et al. [7] for more details. The output module embeds the features extracted from the convolution module for the prediction of each variable, which maps the graph representation embedding of each variable into a two-neuron output.

5.2.3. Loss function

The class distribution is highly unbalanced. In the training dataset, we observe that only $< 10\%$ variables belong to the "destroy" class, i.e., the variables in this class have the potential to improve the current solution by changing their values. In order to adapt to the imbalanced distribution, we applied the WCE loss and focal loss [29] to train the model.

5.2.4. LNS guided by GNNs

Our refined LNS heuristic with GNN classifier (LNS-GNN) is obtained by using the destroy policy π_{θ^*} in Algorithm 5. Hence, given a solution x and its associated state s , we set y equal to $\pi_{\theta^*}(s)$, we then destroy some variables x_i with $y_i = 1$ to get $N(x)$, and we finally obtain a repaired neighbor x' of x by solving the sub-MIP defined by $N(x)$.

5.3. Numerical experiments

In this section, we present the experimental results for the ML-based LNS. As for the WNO application, we first present the data collection, then, we discuss the experimental setting in Section 5.3.2 and the results in Section 5.3.3. The evaluation metrics remain the same as presented in Section 4.4.3.

5.3.1. Data collection

5.3.1.1. MIP benchmark

To train, evaluate and compare algorithms, we consider 126 MIP instances taken from the MIPLIB [30] dataset. For each instance, an initial feasible solution is required to start the LNS heuristic. We use an intermediate solution found by SCIP [31], typically the best solution obtained by SCIP at the end of the root node computation, i.e., before branching.

5.3.1.2. Training data generation

To collect data for the classification task, we use the Local Branching (LB) algorithm of Fischetti and Lodi [28]. More precisely, given a MIP instance and an initial incumbent \bar{x} , we call the LB algorithm to explore as many neighbors as possible, with a time limit of 600 s, and by limiting to 25% the number of destroyed variables. The best solution x^* found by LB is compared with \bar{x} and the variables with changed values are labeled as improving variables for x . These improving variables are then destroyed from x , and the partial solution is repaired by calling a MIP solver. The resulting solution becomes the new incumbent to which we repeat the same process. This is done for a set J of instances, as explained in Section 5.2.

5.3.2. Experimental setting

5.3.2.1. Training

In the classification task for producing a good destroy policy, the GNN model learns from the features of the MIP formulation and its incumbent solution. We train our GNN classifier with 29 of the 126 considered MIP instances.

5.3.2.2. Evaluation

The algorithms are compared and evaluated on the remaining 97 binary MIP instances.

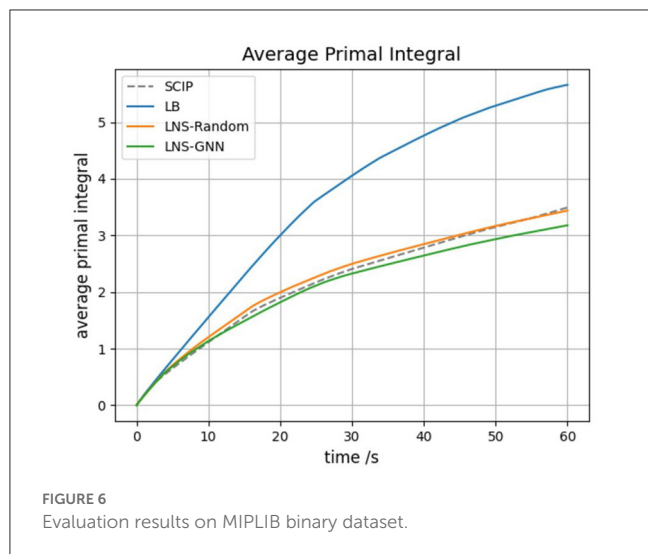
5.3.2.3. Experiment environment

Our code is written in Python 3.7 and we use Pytorch 1.60 [24], Pytorch Geometric 1.7.0 [25], PySCIPOpt 3.1.1 [32], and SCIP 7.01 [31] for developing our models and solving MIPs.

5.3.3. Results

We compare our *LNS-GNN* algorithm with the following algorithms:

- *SCIP*, the default SCIP solver;
- *LNS-Random*, the baseline LNS algorithm with a sampling strategy that randomly selects variables to destroy;
- *LB*, the local branching algorithm [28].



For all these algorithms, we have limited to 40 the number of variables that can be destroyed to generate neighbors.

For measuring the heuristic performance of the compared algorithms, we also compute the average primal integral defined in Section 4.4.3. We ran the listed algorithms for 60 s on each instance. The results are shown in Figure 6.

From the results in Figure 6, we can see that the primal integral of the *LB* baseline is the largest over the entire solving time, which indicates that exploring all the possible neighbors within the same Hamming distance to the current solution is computationally expensive. Moreover, the *LNS-Random* baseline performs better than *LB* baseline by applying a random sampling strategy for selecting the subset of variables to be freed. The default *SCIP* algorithm achieves similar primal integral as *LNS-Random*, which indicates that adding a random LNS strategy to the MIP solver would not significantly improve the overall heuristic behavior of MIP solving. Moreover, from Table 5, we observe that all the LNS algorithms achieve a lower final primal gap than the default *SCIP* algorithm, showing that building a LNS-based primal heuristic on top of a MIP solver generally helps to find a better solution faster. Our *LNS-GNN* algorithm presents the best heuristic behavior in terms of the primal integral, showing that the pretrained GNN model is able to produce structural neighborhoods that contain improving solutions. These results demonstrate that our approach achieves a better trade-off between exploitation and exploration of the promising solution regions. Although, potentially, any of the baseline algorithms could be tuned to obtain better results, this is true for *LNS-GNN* too, and overall we believe that these results show clear promise.

6. Conclusions and future work

In this work, we presented a methodology for integrating machine learning techniques into metaheuristics for solving combinatorial optimization problems. Namely, we proposed a general ML framework for neighbor generation in metaheuristic

TABLE 5 Primal integral (mean \pm standard deviation) for MIPLIB problems with a time limit of 60s for each instance.

| Algo. | Final primal integral | Final primal gap |
|------------|-------------------------------------|-------------------------------------|
| SCIP | 3.493 \pm 7.242 | 3.766 \pm 13.058 |
| LB | 5.661 \pm 10.216 | 2.783 \pm 10.883 |
| LNS-Random | 3.438 \pm 7.233 | 2.724 \pm 11.217 |
| LNS-GNN | 3.179 \pm 6.017 | 2.536 \pm 7.942 |

The bold value in each column of the table indicates the best achieved value of the metric over all the compared algorithms.

search. We firstly defined a neighborhood structure constructed by applying a transformation operator to a selected subset of variables from the current solution. Then, the key of the proposed methodology is to learn a variable selection policy by which the subset transformation operator is able to produce neighborhoods containing high-quality solutions. We formulated the variable selection problem as a classification problem that exploits structural information from the characteristics of the problem and high-quality solutions.

We demonstrated our methodology on two applications. The first problem we addressed occurs in the context of Wireless Network Optimization, where a Tabu Search metaheuristic is used for the topology design sub-problem. In a predefined topology neighborhood structure, we trained classification models to select sized-reduced, but high-quality neighborhoods. This allowed the metaheuristic search to execute more iterations within the same amount of time. As each iteration requires to solve a series of complex combinatorial sub-problems, more iterations entail a greater exploration of the solution space. In addition, to demonstrate the broader applicability of our approach, we also applied our framework to the Large Neighborhood Search metaheuristic for solving MIPs. The experimental results of the two applications have shown that our approach is able to learn a satisfactory trade-off between the exploration of a larger solution space and the exploitation of promising solution regions in metaheuristic search.

Although deep neural networks such as GNNs have been wildly applied to represent combinatorial optimization problems, the current GNNs might not be expressive enough to capture all the crucial patterns from data [11]. For future research, it would be interesting to develop more expressive ML models that exhibit better transferability and scalability across broader classes of problems.

Data availability statement

Publicly available datasets were analyzed in this study. This data can be found at: <https://github.com/pandat8>.

Author contributions

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

Funding

This work was supported by Canada Excellence Research Chair in Data Science for Real-Time Decision-Making at Polytechnique Montreal.

Conflict of interest

The authors declare that the research was conducted in the absence of any commercial or financial relationships

that could be construed as a potential conflict of interest.

Publisher's note

All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers. Any product that may be evaluated in this article, or claim that may be made by its manufacturer, is not guaranteed or endorsed by the publisher.

References

- Blum C, Roli A. Metaheuristics in combinatorial optimization: Overview and conceptual comparison. *ACM Computing Surveys*. (2003) 35:268–308. doi: 10.1145/937503.937505
- Land AH, Doig AG. An automatic method for solving discrete programming problems. In: Jünger M, M, Leibling TM, Naddef D, Nemhauser GL, Pilleblank WR, Reinelt G, Rinaldi G, Wolsey LA, editors. *50 Years of Integer Programming 1958–2008*. Berlin: Springer (2010). p. 105–32.
- Bixby R, Rothberg E. Progress in computational mixed integer programming—a look back from the other side of the tipping point. *Ann Operat Res*. (2007) 149:37–41. doi: 10.1007/s10479-006-0091-y
- Bengio Y, Lodi A, Prouvost A. Machine learning for combinatorial optimization: a methodological tour d'horizon. *Eur J Operat Res*. (2021) 290:405–21. doi: 10.1016/j.ejor.2020.07.06
- Talbi EG. Machine learning into metaheuristics: a survey and taxonomy. *ACM Comput Surveys*. (2021) 54:1–32. doi: 10.1145/3459664
- Karimi-Mamaghan M, Mohammadi M, Meyer P, Karimi-Mamaghan AM, Talbi EG. Machine learning at the service of meta-heuristics for solving combinatorial optimization problems: a state-of-the-art. *Eur J Operat Res*. (2022) 296:393–422. doi: 10.1016/j.ejor.2021.04.032
- Gasse M, Chetelat D, Ferroni N, Charlin L, Lodi A. Exact combinatorial optimization with graph convolutional neural networks. In: Wallach H, Larochelle H, Beygelzimer A, Alche-Bue F, Fox E, Garnett R, editors. *Advances in Neural Information Processing Systems*, Vol. 32. Curran Associates, Inc (2019). Available online at: https://proceedings.neurips.cc/paper_files/paper/2019/file/d14c2267d848abeb81fd590f371d39bd-Paper.pdf
- Gori M, Monfardini G, Scarselli F. A new model for learning in graph domains. In: *Proceedings. 2005 IEEE International Joint Conference on Neural Networks*, vol. 2. Montreal, QC: IEEE (2005). p. 729–34.
- Scarselli F, Gori M, Tsoi AC, Hagenbuchner M, Monfardini G. The graph neural network model. *IEEE Trans Neural Netw*. (2008) 20:61–80. doi: 10.1109/TNN.2008.2005605
- Hamilton WL, Ying R, Leskovec J. Representation learning on graphs: methods and applications. *arXiv preprint arXiv:170905584*. (2017). doi: 10.48550/arXiv.1709.05584
- Cappart Q, Chételat D, Khalil E, Lodi A, Morris C, Veličković P. Combinatorial optimization and reasoning with graph neural networks. *arXiv preprint arXiv:210209544*. (2021). doi: 10.48550/arXiv.2102.09544
- Khalil E, Dai H, Zhang Y, Dilkina B, Song L. Learning combinatorial optimization algorithms over graphs. In: Guyon I, Von Luxburg U, Bengio S, Wallach H, Fergus R, Vishwanathan S, Garnett R, editors. *Advances in Neural Information Processing Systems*, Vol. 30. Curran Associates, Inc (2017). Available online at: https://proceedings.neurips.cc/paper_files/paper/2017/file/d9896106ca98d3d05b8cbdf4fd8b13a1-Paper.pdf
- Nazari M, Oroojlooy A, Snyder L, Takác M. Reinforcement learning for solving the vehicle routing problem. *Adv Neural Inform Process Syst*. (2018) 31:4240. doi: 10.48550/arXiv.1802.04240
- Zhang C, Song W, Cao Z, Zhang J, Tan PS, Chi X. Learning to dispatch for job shop scheduling via deep reinforcement learning. *Adv Neural Inform Process Syst*. (2020) 33:1621–32. doi: 10.48550/arXiv.2010.12367
- Bello I, Pham H, Le QV, Norouzi M, Bengio S. Neural combinatorial optimization with reinforcement learning. *arXiv preprint arXiv:161109940*. (2016). doi: 10.48550/arXiv.1611.09940
- Gao L, Chen M, Chen Q, Luo G, Zhu N, Liu Z. Learn to design the heuristics for vehicle routing problem. *arXiv preprint arXiv:200208539*. (2020). doi: 10.48550/arXiv.2002.08539
- Liu D, Lodi A, Tanneau M. Learning chordal extensions. *J Glob Opt*. (2021) 81:3–22. doi: 10.48550/arXiv.1910.07600
- He H, Daume III H, Eisner JM. Learning to search in branch and bound algorithms. *Adv Neural Inform Process Syst*. (2014) 27:3293–301. doi: 10.5555/2969033.2969194
- 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*. (2016). p. 30. Available online at: <https://ojs.aaai.org/index.php/AAAI/article/view/10080> (accessed February 21, 2016).
- Khalil E, Dai H, Zhang Y, Dilkina B, Song L. Learning combinatorial optimization algorithms over graphs. In: *Advances in Neural Information Processing Systems*. (2017). p. 30.
- Balcan M-F, Dick T, Sandholm T, Vitercik E. Learning to branch. In: Dy J, Krause A, editors. *Proceedings of the 35th International Conference on Machine Learning*, Vol. 80. PMLR (2018). p. 344–53. Available online at: <http://proceedings.mlr.press/v80/balcan18a/balcan18a.pdf>
- Liu D, Fischetti M, Lodi A. Learning to Search in Local Branching. *Proc AAAI Conf Artif Intell*. (2022) 36:3796–803. doi: 10.1609/aaai.v36i4.20294
- Perreault V. *Tactical Wireless Network Design for Challenging Environments*. Montreal, QC: Polytechnique Montreal. (2022).
- Pasze A, Gross S, Massa F, Lerer A, Bradbury J, Chanan G, et al. Pytorch: an imperative style, high-performance deep learning library. *Adv Neural Inform Process Syst*. (2019) 32:8026–37. doi: 10.48550/arXiv.1912.01703
- Fey M, Lenssen JE. Fast graph representation learning with PyTorch Geometric. *arXiv preprint arXiv:190302428*. (2019). doi: 10.48550/arXiv.1903.02428
- Berthold T. Measuring the impact of primal heuristics. *Operat Res Lett*. (2013) 41:611–4. doi: 10.1016/j.orl.2013.08.007
- Danna E, Rothberg E, Le Pape C. Exploring relaxation induced neighborhoods to improve MIP solutions. *Math Progr*. (2005) 102:71–90. doi: 10.1007/s10107-004-0518-7
- Fischetti M, Lodi A. Local branching. *Math Progr*. (2003) 98:23–47. doi: 10.1007/s10107-003-0395-5
- Lin T-Y, Goyal P, Girshick R, He K, Dollár P. Focal Loss for dense object detection. *IEEE Trans Pattern Anal Mach Intell*. (2020) 42: 318–27. doi: 10.1109/TPAMI.2018.2858826
- Gleixner A, Hendel G, Gamrath G, Achterberg T, Bastubbe M, Berthold T, et al. *MIPLIB 2017: data-driven compilation of the 6th mixed-integer programming library*. *Math Progr Comput*. (2021) 2021:1–48. doi: 10.1007/s12532-020-00194-3
- Gamrath G, Anderson D, Bestuzheva K, Chen WK, Eifler L, Gasse M, et al. *The SCIP Optimization Suite 7.0*. Berlin: Zuse Institute. (2020). p. 10–20. Available online at: <http://nbn-resolving.de/urn:nbn:de:0297-zib-78023> (accessed March 30, 2020).
- Maher S, Miltenberger M, Pedroso JB, Rehfeldt D, Schwarz R, Serrano F. *PySCIPOpt: mathematical programming in python with the SCIP optimization suite*. In: *Mathematical Software – ICMS*. (2016). Cham: Springer International Publishing (2016). p. 301–7.