

# Combining supervised learning and local search for the multicommodity capacitated fixed-charge network design problem

Charly Robinson La Rocca<sup>a,\*</sup>, Jean-François Cordeau<sup>b</sup>, Emma Frejinger<sup>a</sup>

<sup>a</sup>*Department of Computer Science and Operations Research, Université de Montréal, 2900, boul. Édouard-Montpetit, Montreal, H3T 1J4, Canada*

<sup>b</sup>*Department of Logistics and Operations Management, HEC Montréal, 3000 Côte-Sainte-Catherine Road, Montreal, H3T 2A7, Canada*

---

## Abstract

The multicommodity capacitated fixed-charge network design problem has been extensively studied in the literature due to its wide range of applications. Despite the fact that many sophisticated solution methods exist today, finding high-quality solutions to large-scale instances remains challenging. In this paper, we explore how a data-driven approach can help improve upon the state of the art. By leveraging machine learning models, we attempt to reveal patterns hidden in the data that might be difficult to capture with traditional optimization methods. For scalability, we propose a prediction method where the machine learning model is called at the level of each arc of the graph. We take advantage of off-the-shelf models trained via supervised learning to predict near-optimal solutions. Our experimental results include an algorithm design analysis that compares various integration strategies of predictions within local search algorithms. We benchmark the ML-based approach with respect to the state-of-the-art heuristic for this problem. The findings indicate that our method can outperform the leading heuristic on sets of instances sampled from a uniform distribution.

*Keywords:* Heuristics, Supervised Learning, Multicommodity Capacitated Fixed Charge Network Design, Local search, Transportation, Machine Learning.

---

---

\*Corresponding author

*Email addresses:* charly.robinson.la.rocca@umontreal.ca (Charly Robinson La Rocca), jean-francois.cordeau@hec.ca (Jean-François Cordeau), emma.frejinger@umontreal.ca (Emma Frejinger)

## 1. Introduction

The multicommodity capacitated fixed-charge network design problem (MCFND) is an influential problem in combinatorial optimization (CO). It is an NP-hard problem (Magnanti and Wong, 1984) that has been studied for decades and for which many exact algorithms (Chouman et al., 2017; Crainic and Gendron, 2021) and heuristics (Yaghini et al., 2013; Yaghini and Foroughi, 2014; Katayama, 2020) have been proposed. The research interest for the MCFND is motivated by its general structure which is compatible with applications in logistics, telecommunications, and infrastructure planning (Minoux, 1989; Quilliot, 2014). Recently, extensions to the network design problem have been developed to model congestion (Paraskevopoulos et al., 2016; Karimi-Mamaghan et al., 2020), scheduling (Hewitt and Lehu  d  , 2023; Kidd et al., 2024), multimodality (Basallo-Triana et al., 2023; Real et al., 2021) and uncertainty (An and Lo, 2016; Ghanei et al., 2023). At a high level, the objective of the MCFND is to jointly determine the structure of the network and the flows that minimize the sum of both the fixed installation costs and variable routing costs. The fixed costs are related to the activation of arcs. They can, for example, represent the decision to open a facility or a transportation route. The variable costs model the cost of sending a unit of flow on a given arc of the network. A solution to the MCFND must respect the capacity available on each arc and the demand for each commodity.

Recent advancements in heuristic methods (Gendron et al., 2018; Katayama, 2020) have enabled the discovery of high-quality solutions to large-scale instances. They often exploit continuous relaxations of MCFND, which are computationally efficient to solve (Crainic et al., 2004; Katayama et al., 2009). The performance of heuristics is often the result of a combination of many different techniques carefully designed by experts. Typically, improving the performance of a heuristic requires a deep understanding of the problem under study and significant research efforts. This limits the rate of improvement of the state of the art (SOA). Given these circumstances, we explore how a data-driven approach can help improve upon the SOA. Our method feeds a prediction from a machine learning (ML) model to a Local Search (LS) algorithm. The ML model is trained via supervised learning and labels are built using near-optimal solutions. This approach is inspired by the recent successes of ML for CO (Bengio et al., 2021).

Our work differentiates itself from the current trend in ML, which is to improve performance via increasingly larger models and datasets. This trend can be explained by scaling laws which suggest that performance strongly depends on the number of parameters, the dataset size and the amount of computing power (Kaplan et al., 2020). Many challenges arise when doing so such as increased cost, carbon emissions impact (Dhar, 2020) and reproducibility concerns (Albertoni et al., 2023). These issues are well known and many researchers are actively working on sample-efficient methods to obtain accurate predictions with small datasets (Wallingford et al., 2023; La Rocca et al., 2024). Given that one of our objectives is to remain reproducible, we choose to work with off-the-shelf ML models that are widely accessible. Our contribution lies in the design of a technique for the integration of such models within a heuristic algorithm. In particular, we propose a prediction method where the ML model is called at the level of each arc of the graph. This ensures that the output complexity does not scale exponentially with the size of the instance. This is a key feature of our method that makes it computationally efficient, even for large instances.

The open research question we attempt to answer is how to best use the ML prediction to help guide a SOA heuristic for the MCFND. We believe that our work is the first to explore this question as most of the literature on this problem focuses on exact methods and human-designed heuristics. To bridge this gap, this paper makes the following contributions:

- A framework for the integration of supervised learning and CO.
- A set of sampling routines that generate informative features for the MCFND.
- An algorithm design analysis that compares different ways to integrate the ML predictions in LS algorithms.
- A comparison of different ML models and the effect of bias adjustments.
- A study of the performance of the SOA heuristic with and without ML predictions.

The remainder of this article is organized as follows. Section 2 presents the formulation and related works on the MCFND. We also review some of the literature on supervised learning and the integration of ML in CO. This is followed in Section 3 by our methodology, which includes a general framework for hybrid ML-CO and the sampling

routines used to generate informative features. In Section 4, we quantify our performance with metrics that measure the quality of solutions and the speed at which they are found. Our test set includes the GT dataset (Hewitt et al., 2010) and large instances produced by the Canad generator (Larsen et al., 2023). For each instance, we compare the performance of the SOA method (Katayama, 2020) augmented by ML predictions against its standalone version. Finally, Section 5 concludes the paper and discusses future work.

## 2. Related Works

This section describes the problem definition and the body of work related to the MCFND, with an emphasis on heuristics developed to identify near-optimal solutions in a short amount of time. We then transition to a review of supervised learning, and the gradient boosting models that we use in our experiments. The final part of this section synthesizes recent efforts related to the integration ML and CO.

### 2.1. Problem Formulation

The MCFND can be defined on a graph  $G = (N, A)$ , where  $N$  is the set of nodes and  $A$  is the set of arcs. Each commodity  $k \in K$  has a demand  $d_k$  that must be routed from its origin  $O(k) \in N$  to its destination  $D(k) \in N$ . The amount of flow on arc  $(i, j) \in A$  for commodity  $k$  is denoted by  $x_{ij}^k$ , and it incurs a cost per unit of flow  $c_{ij}^k$ . If any amount of flow is routed on arc  $(i, j)$ , a fixed cost  $f_{ij}$  must be paid to use it. This decision is represented by the binary variable  $y_{ij}$ . The objective is to minimize the sum of both the fixed and variable costs of the arcs used to route all commodities. In terms of constraints, the total flow on arc  $(i, j)$  cannot exceed its capacity  $u_{ij}$ . Also, flow conservation must be satisfied at each node  $i \in N$ . To model this constraint, we define sets of successors  $N_i^+$  and predecessors  $N_i^-$  for each node  $i \in N$ . The arc-flow formulation (AF) of  $\mathcal{P}$  can be stated as the following mixed-integer program (MIP):

$$(\mathcal{P}) \quad \min \sum_{k \in K} \sum_{(i,j) \in A} c_{ij}^k x_{ij}^k + \sum_{(i,j) \in A} f_{ij} y_{ij} \quad (1)$$

s.t.

$$\sum_{k \in K} x_{ij}^k \leq u_{ij} y_{ij}, \quad \forall (i, j) \in A, \quad (2)$$

$$\sum_{j \in N_i^+} x_{ij}^k - \sum_{j \in N_i^-} x_{ji}^k = \begin{cases} d^k, & \text{if } i = O(k) \\ 0, & \text{if } i \neq O(k), D(k), \\ -d^k, & \text{if } i = D(k) \end{cases}, \quad i \in N, k \in K, \quad (3)$$

$$x_{ij}^k \geq 0, \quad \forall (i, j) \in A, k \in K, \quad (4)$$

$$y_{ij} \in \{0, 1\}, \quad \forall (i, j) \in A. \quad (5)$$

The objective function (1) minimizes the sum of flow routing costs and arc opening costs. Constraints (2) ensure that the total flow on each arc does not exceed its capacity and bind the decision of opening an arc to the amount of flow routed on it. Constraints (3) enforce flow conservation at each node. Finally, constraints (4) and (5) define the domains of the decision variables.

### 2.2. Continuous Relaxations for the MCFND

It is common for MCFND heuristics to exploit a continuous relaxation of problem  $\mathcal{P}$  (defined by (1) - (5)) where the integrality requirements on the binary variables  $y_{ij}$  are relaxed. The relaxed problem  $\tilde{\mathcal{P}}$  is significantly easier than the original problem  $\mathcal{P}$  and its optimal solution can be used as a tool to explore to solution space efficiently. The main hypothesis is that the solution of  $\tilde{\mathcal{P}}$  is a reasonable approximation of the optimal solution for  $\mathcal{P}$ . In particular, if the total amount of flow that goes through a given arc in  $\tilde{\mathcal{P}}$  is zero, the corresponding binary variable is likely to be zero in the optimal solution to the original problem. Inversely, if most of the capacity on the arc is used in  $\tilde{\mathcal{P}}$ , the corresponding binary variable is likely to be one in the optimal solution. This insight is the basis for scaling

approaches that iteratively solve the relaxed problem with a modified objective function (Crainic et al., 2004) or set of constraints (Katayama et al., 2009).

A popular method in this category is Slope Scaling (SS) (Crainic et al., 2004), which adjusts the objective function coefficients of each arc to reflect the effective cost of passing flow through it, by considering a combination of both fixed and variable costs. More precisely, the objective function coefficient of an arc for a given commodity becomes the sum of the variable cost and a linearization factor. This factor is the ratio of the fixed cost and the total flow on the arc. The idea is to increase the coefficient of arcs for which the fixed cost is high relative to the flow. One can also set the linearization factor to a large number to force the flow to be zero on certain arcs. Given the optimal solution  $\tilde{x}$  of  $\tilde{\mathcal{P}}$ , we can construct a feasible solution  $y^{SS}$  for the original problem by rounding up the normalized total flow on each arc. Here, bold formatting is used to denote vectors. The rounding is done according to the following equation:

$$y_{ij}^{SS} = \left\lceil \frac{\sum_{k \in K} \tilde{x}_{ij}^k}{u_{ij}} \right\rceil, \quad \forall (i, j) \in A. \quad (6)$$

Capacity Scaling (CS) (Katayama et al., 2009) is a similar approach that iteratively adjusts the capacity of each arc. The method first solves the continuous relaxation of the path-flow formulation (PF) of the MCFND to obtain a fractional solution  $\tilde{y}$ . In addition to the design variables, the PF formulation incorporates variables  $z_p^k$  to represent the flows routed through every possible path  $p$  for commodity  $k$ . These variables are defined for every path in the set  $P^k$ , which contains all feasible paths for commodity  $k$  from the origin  $O(k)$  to the destination  $D(k)$ . The PF can be stated as follows:

$$\min \sum_{(i,j) \in A} \sum_{k \in K} c_{ij}^k \sum_{p \in P^k} \delta_{ij}^p z_p^k + \sum_{(i,j) \in A} f_{ij} y_{ij} \quad (7)$$

s.t.

$$\sum_{p \in P^k} z_p^k = d^k, \quad \forall k \in K, \quad (8)$$

$$\sum_{k \in K} \sum_{p \in P^k} \delta_{ij}^p z_p^k \leq u_{ij} y_{ij}, \quad \forall (i, j) \in A, \quad (9)$$

$$\sum_{p \in P^k} \delta_{ij}^p z_p^k \leq d^k y_{ij}, \quad \forall (i, j) \in A, k \in K, \quad (10)$$

$$z_p^k \geq 0, \quad \forall p \in P^k, k \in K, \quad (11)$$

$$y_{ij} \in \{0, 1\}, \quad \forall (i, j) \in A. \quad (12)$$

The binary parameter of  $\delta_{ij}^p$  indicates whether or not the arc  $(i, j)$  is included in path  $p$ . The objective function (7) minimizes the total cost similar to the AF formulation. Constraints (8) act as flow conservation constraints for each commodity. Constraints (9) represent the capacity of each arc and (10) provide the forcing constraints. Non-negativity of path-flow variables is enforced by (11).

The CS routine iteratively solves the PF formulation (defined by (7)-(11)) and updates the capacity  $u_{ij}$  according to the following equation:

$$u'_{ij} = \lambda u_{ij} \tilde{y}_{ij} + (1 - \lambda) u_{ij}, \quad \forall (i, j) \in A, \quad (13)$$

where  $\lambda \in [0, 1]$  is the so-called capacity smoothing factor that controls the aggressiveness of the scaling. The idea is to increase the capacity of arcs that are used in the solution and decrease the capacity of arcs that are not used. The method is similar to slope scaling in that it aims to force design variables to be zero or one. If  $\tilde{y}_{ij}$  is integral for all arcs, then the solution process stops as the solution is feasible for the original problem. However, in practice, termination criteria are used to stop the process after a certain number of iterations or when the total number of fractional variables is below a given threshold.

The CS procedure in Katayama et al. (2009) integrates a column generation technique to improve its computational efficiency. This is required because the total number of paths in the PF is often huge. The method starts with a small set of paths and iteratively adds new paths to the set to improve the solution. The new paths are identified by solving

the pricing problem, which considers the reduced cost of each path flow variable. The method is effective because it allows the solution process to focus on a small set of paths that are likely to be used in the optimal solution. At each iteration of the method, the computational effort increases and, therefore, a termination criterion is used to stop the process after a certain number of iterations.

### 2.3. Local Search for the MCFND

In the previous section, we introduced different scaling approaches that aim to build a feasible solution for  $\mathcal{P}$  by iteratively solving a relaxed problem. However, they rarely find the optimal solution. In fact, they are most useful as a preprocessing step to warm-start another heuristic. They are designed to explore the solution space around a given solution to find a better one (Shaw, 1998). SOA heuristics for the MCFND often combine scaling and LS to find high-quality solutions. In this section, we describe two examples of LS heuristics, which are used in our experiments.

The work from Gendron et al. (2018) presents a matheuristic that combines iterative linear programming and SS to solve the MCFND. At each iteration, the solution  $\tilde{y}$  of the linear relaxation is used to build index sets  $A^0 = \{(i, j) \in A | \tilde{y}_{ij} = 0\}$  and  $A^1 = \{(i, j) \in A | \tilde{y}_{ij} = 1\}$ . In the SS method, arcs in  $A^0$  are forced to be zero by fixing the linearization factor to a large number. Inversely, the linearization factor is fixed to 0 for arcs in  $A^1$ . A feasible solution  $y^{SS}$  is then constructed using (6). In the next step, the selected variables corresponding to these index sets are fixed at their LP relaxation values in the restricted problem. This problem is then solved using a MIP solver with  $y^{SS}$  as a warm-start solution. A pseudo-cut is then added to the MIP model to remove from the feasible set all solutions that satisfy  $y_{ij} = \tilde{y}_{ij}, (i, j) \in A^0 \cup A^1$ :

$$\sum_{(i,j) \in A^0} y_{ij} + \sum_{(i,j) \in A^1} (1 - y_{ij}) \geq 1. \quad (14)$$

The method is compared with many other heuristics such as Cycle-based Tabu search (Ghamlouche et al., 2003), Path Relinking (Ghamlouche et al., 2004), Capacity Scaling (Katayama et al., 2009) and Simulated Annealing with column generation (Yaghini et al., 2013). The results show that the method is competitive on the Canad instances (Crainic et al., 2001), which constitute a widely used benchmark for the MCFND.

More recently, Katayama (2020) proposed to combine CS with a MIP neighbourhood search algorithm, which is the most effective method to produce high-quality solutions for the MCFND according to our reading of the literature. Hence, our work uses this SOA method as a benchmark to compare the performance of the ML models and, in the remainder of this manuscript, we refer to this method as LS\*. The heuristic follows a similar strategy to the one described in Gendron et al. (2018). However, there are many differences that help improve the performance of the method. First, the CS method is not only used to find a feasible solution, but also to aggressively prune the problem. The set of arcs associated with the design variables that converged to 0 during the CS phase are explicitly deleted from the problem instead of being fixed. Both the continuous and binary variables are removed from the problem, which significantly reduces the memory and time required to solve the problem. Second, LS\* adds objective cut-offs to force the MIP solver to find strictly better solutions than the incumbent. Third, the neighbourhood built around the incumbent solution is dynamically updated based on the result of the previous iteration. This neighbourhood is defined by two constraints. The first one removes the incumbent solution  $\bar{y}$  from the feasible set and forces the solver to find a new solution with at least one additional arc closed:

$$\sum_{(i,j) \in A | \bar{y}_{ij}=1} y_{ij} \leq \sum_{(i,j) \in A} \bar{y}_{ij} - 1. \quad (15)$$

The second constraint defines the size  $M$  of the neighbourhood to explore:

$$\sum_{(i,j) \in A | \bar{y}_{ij}=1} y_{ij} \geq \sum_{(i,j) \in A} \bar{y}_{ij} - M. \quad (16)$$

It states that the number of open arcs in the new solution must be at most  $M$  units smaller than the number of open arcs in the incumbent solution. If the solver fails to find a feasible solution within the time limit, the size of the neighbourhood  $M$  is decreased by a factor of 2. In the case where the solver proves that no feasible solution exists in the neighbourhood or  $M = 0$ , the algorithm stops. For more details on LS\*, we refer to the original paper (Katayama, 2020). Our work explores how ML can be used to improve the performance of LS\* by predicting near-optimal solutions via supervised learning.

## 2.4. Supervised Learning

Supervised learning is a branch of ML that aims to learn a mapping between input and output data. In the context of combinatorial optimization, the input data is a set of features that describe the problem instance and the output data is related to its solution. Often, supervised learning models for binary classification are trained by minimizing the binary cross-entropy loss

$$\mathcal{L}_B(y', y) = -\omega_1 y \log(y') - \omega_0 (1 - y) \log(1 - y'), \quad (17)$$

where  $y'$  is the predicted label and  $y$  is the true label. We use the notation  $\omega_1$  and  $\omega_0$  to denote the weights of the positive and negative classes, respectively. The weights are used to control the contribution of each class to the loss, which is useful when the classes are unbalanced. We cover here the linear binary and gradient boosting models that are used in our experiments.

Linear binary classifiers are widely appreciated for their simplicity, and efficiency in handling linearly separable datasets. They are defined by a set of weights and a bias term, which are multiplied by the input features and added to the result, respectively. The sign of the result determines the predicted label. Once trained with  $\mathcal{L}_B$ , they produce a hyperplane that separates the input features into two classes (Hastie et al., 2001b). The model's performance is constrained by the fact that it can only represent linear relationships, which limits its ability to capture complex patterns in the data. This can be mitigated by introducing non-linear transformations of the input features, which is the basis for kernel methods (Hastie et al., 2001a).

Gradient Boosted Machines (GBM) represent a class ML algorithms that combine multiple weak learners to form a strong predictive model. The essence of GBMs lies in iteratively correcting the residuals of the ensemble, thereby improving the model's accuracy with each step. The foundational work for GBM (Friedman, 2001) introduced the concept of boosting and demonstrated its efficacy in minimizing various loss functions in a stepwise fashion. Empirical results show that boosted tree methods tend to perform exceptionally well across a wide range of tasks when compared to other alternatives such as support vector machines (SVM) or neural networks (Caruana and Niculescu-Mizil, 2006; Borisov et al., 2022). Building upon this foundation, the development of XGBoost (Chen and Guestrin, 2016) and, more recently, Light GBM (LGBM) (Ke et al., 2017) marked a significant advancement in the field. These implementations of GBMs include various algorithmic enhancements, such as regularized learning objectives, parallelized tree construction, and efficient memory management. Given the tabular nature of our dataset that relates to the structure of the MCFND, we chose to focus on GBMs as our primary ML model. In particular, we use LGBM, a widely recognized and accessible library, which contributes to the reproducibility of our work.

## 2.5. ML for CO

There is a vast literature on the use of ML techniques to solve CO problems (Dai et al., 2017; Bengio et al., 2021; Kotary et al., 2021). Within MIP solvers, ML predictions can support selection strategies for nodes (Khalil et al., 2022), cuts (Turner et al., 2023) or branches (Khalil et al., 2016; Lodi and Zarpellon, 2017). In this section, we focus on the use of ML to inform a heuristic that reduces the search space of the problem with the goal of finding high-quality solutions in a short amount of time. A common strategy is fix-and-optimize (Hanafi and Todosijevic, 2017), where ML is used to identify which subset of variables to fix. La Rocca et al. (2024) propose an entropy minimization technique to predict, in a one-shot fashion, the optimal solution of generic MIP instances with SOS1 constraints. Given the tripartite graph structure of a MIP instance, Graph Neural Networks (GNN) (Hamilton, 2020; Liu et al., 2022) are often used to capture patterns in the data and predict solution values for binary variables. The work of Ding et al. (2020) demonstrates the effectiveness of this strategy on eight different classes of CO problems. The authors show that a local branching cut (Fischetti and Lodi, 2003) of the form

$$\sum_{(i,j) \in A} y_{ij}(1 - y'_{ij}) + (1 - y_{ij})y'_{ij} \leq \beta, \quad (18)$$

which restricts the search space within a distance  $\beta$  of the predicted solution  $y'$ , can improve the primal gap when compared to SCIP's aggressive heuristic setting. They also compare the precision of their GNN model with XGBoost and show that they have similar performance when both models have access to the LP relaxation solution. This suggests that GBMs can be a reasonable alternative to GNNs even if the data has a graph structure.

Their methodology is applicable to a wide range of problems and, for that reason, their experimental results are compared to a generic MIP solver. For many CO problems, specialized algorithms are available that exploit the structure of the problem to find high-quality solutions. The travelling salesman problem (TSP) is a well-known example of a CO problem that has been studied extensively and for which specialized solvers exist such as Concorde (Applegate et al., 2006).

To show that ML can be competitive with SOA algorithms, Sun et al. (2021) propose a reduction strategy for TSP instances that uses an SVM to predict which decision variables to remove. They introduce the idea of stochastic sampling of feasible tours to generate statistical features for the SVM. This is effective because one can generate a feasible route of the TSP in polynomial time. We mention this specific work because we use a similar strategy to generate so-called *sampling features*, which we introduce in Section 3. Using the sampling data, the authors construct statistical measures that aggregate the information for each arc. Their first measure is computed using a weighted sum, where weights are inversely proportional to the rank of the solution. Mathematically, the rank-based feature  $\phi_{ij}^r$  for arc  $(i, j)$  is defined as follows:

$$\phi_{ij}^r = \sum_{k=1}^m \frac{y_{ij}^k}{r_k}, \quad (19)$$

where  $r_k$  is the rank of the solution  $k$  based on the objective value, and  $m$  is the total number of samples. They also developed a measure that uses the Pearson correlation coefficient between each variable and objective values across samples. Their strategy is motivated by the hypothesis that arcs that frequently appear in high-quality solutions are more likely to be part of the optimal solution. Their results show that the SVM model is competitive when compared to other generic problem reduction methods such as Construct, Merge, Solve and Adapt (CSMA) (Blum et al., 2016).

In this paper, our main objective is to study how ML predictions can be used to accelerate a SOA heuristic for the MCFND. From our assessment of the literature, this has not been explored before. We believe that our work is relevant because the MCFND has a wide range of applications, which means that our results can be relevant to many other related problems. Furthermore, we aim to propose a methodology that is sample efficient and easily reproducible. In the next sections, we describe the methodology followed by the experimental setup used to evaluate its performance.

### 3. Methodology

We follow a similar approach that of Sun et al. (2021), where an SVM is trained to predict which arcs should be excluded from TSP instances. The core idea is to solve a set of training instances to near-optimality to collect labels for the supervised learning model. Our methodological contributions come from the adaptations we made to the original approach to make it suitable for the MCFND, and from the fact that we integrate ML predictions into a SOA heuristic.

An important characteristic of our method is related to the decision to make predictions locally on each individual arc instead of globally on the entire graph (Liu et al., 2022). Given an instance with  $|A|$  arcs, the output space size is thus  $2^{|A|}$  because each arc has two possible outcomes. The exponential growth makes global predictions challenging to scale to large instances, while maintaining the ML approach tractable, efficient and accurate. Local predictions mean we can use a simple binary classifier as our supervised learning model. The main trade-off is that we lose the ability to capture information from the entire graph. This can be mitigated via feature engineering where we extract graph and sampling features to inject global information into the model.

**Graph features.** Each arc in the MCFND has three basic characteristics: the capacity, the variable cost and the fixed cost. The first set of features is derived by normalizing characteristics to their maximum value within the instance. The second set contains node graph features, which represent the local topology of the graph. We use the degree, the type (source or sink), and aggregated metrics (mean and max) of the basic arc features in the neighbourhood for source and destination nodes. Note that the exact design choice for graph features is not critical to the performance of the heuristic. As we show in Section 4, local graph features are weak predictors of the optimal class.

**Sampling features.** The idea to sample solutions to generate features is inspired by Sun et al. (2021), where random TSP routes are generated to collect data for the supervised learning model. They construct two measures that aggregate the information using statistical tools. For our purposes, we include statistical features in a similar fashion.

For each variable, we compute the average value and how frequently it is closed or open. However, we decided to not solely depend on aggregated metrics because, unlike the TSP, we cannot generate at will a consistently large sample size. Our preliminary experiments reveal that, for some instances, our sampling routine could only find a few dozen feasible solutions. For that reason, we decided to include as-is the 10 best feasible solutions found during the sampling process as features. The data is passed to the model without further processing to limit our dependence to human designed features.

**Sampling routines.** Given that, for the MCFND, there is not a well established way to build feasible solutions in polynomial time akin to the TSP, we propose sampling routines to generate solutions. They take advantage of a continuous relaxation of the problem, which is computationally efficient to solve. They are designed to generate a large and diverse set of solutions in a short amount of time. In particular, we propose Randomized Slope Scaling (RSS) and Local Search Fast Sampling (LSFS), which use SS and CS, respectively. We utilize two different sampling routines to diversify and reduce the bias in the set of samples.

The RSS routine is based on Gendron et al. (2018), where a MIP solver and SS are called sequentially. We found that this heuristic is not competitive with  $LS^*$  in terms of solution quality. However, it is effective as a sampling routine because it can produce feasible solutions relatively quickly. We modified the original implementation in a few different ways to improve its ability to explore the solution space. First, we pass a random value for both the `Seed` and `Heuristics` parameters of Gurobi at each iteration. Second, we set the objective coefficient to a large number for 1% of the variables chosen at random. Third, the time limit per iteration is set to 20 seconds instead of 100.

The LSFS routine uses the same implementation as  $LS^*$  with a few hyperparameter changes. The time limit for the MIP solver is also set to 20 seconds and we reduce the number of columns generated to limit the computational time associated with CS.

**Training.** The training process is straightforward because we use off-the-shelf ML libraries that have high-level interfaces available. In a supervised setting, we need to provide the ML model  $\psi$  with a set of features and their corresponding labels. The featurizer  $\phi$  is a function that takes the set of solutions  $\mathcal{X}$  and returns a set of features according to the feature engineering process described above. The set of binary labels  $\mathcal{Y}$  are built by merging the top three best known solutions instead of just the best one, i.e., all arcs that are open in at least one of the solutions are labelled as open. Both the features and the labels are passed to the library-specific fit routine to generate a trained ML model  $\hat{\psi}$  as follows:

$$\hat{\psi} \leftarrow \text{fit}(\psi, \phi(\mathcal{X}), \mathcal{Y}). \quad (20)$$

**Inference.** Algorithm 1 describes the Supervised Local Search (SLS) algorithm, which is a high-level description of our inference pipeline. It is designed to be flexible, so it can be adapted to different CO problems. It takes as input the problem  $\mathcal{P}$ , the trained ML model  $\hat{\psi}$ , the local search algorithm  $\mathcal{A}_{LS}$  and the set of sampling routines  $\mathbb{A}_{SPL}$ . In the first phase, the algorithm parallelizes the sampling process using the routines in  $\mathbb{A}_{SPL}$ . The solutions collected are accumulated in a set  $\mathcal{X}$ , which is then used to get an incumbent solution  $\bar{\mathbf{y}}$  that can warm-start the solution process. We get an estimate of near-optimal solutions  $\mathbf{y}'$  as the output of the ML model with  $\phi(\mathcal{X})$  as its input. The local search algorithm  $\mathcal{A}_{LS}$  is then called to attempt to find a high-quality solution  $\mathbf{y}^{LS}$  of the problem  $\mathcal{P}$ . Note that Algorithm 1 does not explicitly describe how the ML predictions are integrated in the solution process. This is an implementation detail of the algorithm  $\mathcal{A}_{LS}$ , and Section 4.3 provides experimental results for concrete strategies.

**Problem reduction.** The problem reduction step is a key part of the  $LS^*$  algorithm and one of the main reasons why it is so effective. At every CS iteration, each arc with a fractional value below a certain threshold is removed from the problem. This significantly reduces the memory and time required to solve the problem. However, it introduces a challenge related to the integration of guiding solutions. In particular, we found that both the incumbent  $\bar{\mathbf{y}}$  and the ML prediction  $\mathbf{y}'$  are rarely feasible in the reduced problem. Furthermore,  $LS^*$  already incorporates local cuts (15) and (16), which are more restrictive than (18). For these reasons, the cut (18) does not improve the performance of  $LS^*$ . To circumvent the misalignment between the reduced problem and the guiding solutions, we decided to replace the original CS-based reduction by a routine that uses ML to predict which arcs to remove, akin to Sun et al. (2021). This is motivated by the fact that the CS routine can take up to 1 hour to execute. Also, we can make the required adjustment to make sure the incumbent and the ML prediction is feasible in the reduced problem.



---

**Algorithm 1** Supervised Local Search (SLS) algorithm

---

```
procedure SLS( $\mathcal{P}, \hat{\psi}, \mathcal{A}_{LS}, \mathbb{A}_{SPL}$ )  
   $X = \{\}$  ▷ Initialize the set of solutions  
  for (In parallel)  $\alpha \in \mathbb{A}_{SPL}$  do  
     $X \leftarrow \text{sample}(\mathcal{P}, \alpha)$  ▷ Generate features using the sampling routine  
  end for  
   $\bar{y} \leftarrow \text{get\_best}(X)$  ▷ Get best solution from the set  
   $y' \leftarrow \hat{\psi}(\phi(X))$  ▷ Forward pass of the ML model  
   $y^{LS} \leftarrow \text{solve}(\mathcal{P}, \mathcal{A}_{LS}, \bar{y}, y')$  ▷ Solve using the local search algorithm  
  return  $y^{LS}$   
end procedure
```

---

## 4. Experimental Results

This section presents the results of our experiments. We first describe the experimental setup and the different datasets used to test our method. We then present some intermediate results on the extended runs and the sampling routines, which give us the necessary information to train the ML model. We include results on the tuning procedures for LS\* and our ML models. This is followed by an algorithm design analysis where we test different ways to integrate the ML prediction in an LS algorithm. Finally, we explain the main result of our work, which is the comparison of the performance of LS\* with and without ML predictions.

**Metrics.** The metrics we use to evaluate the performance of the heuristics are the primal gap (PG) and the primal integral (PI). The primal gap is the relative difference between the objective function values of the best solution found by the heuristic and of the best known solution. The primal integral is calculated using the area under the curve of the primal gap over time. The primal gap is a measure of the quality of the solutions found, and the primal integral is a measure of the speed at which the heuristic finds high-quality solutions. Each table contains the mean, the geometric shifted mean and quantiles of the corresponding metric.

**Setup.** Our software stack is written in Julia and we use the JuMP package (Lubin et al., 2023) to interface with the Gurobi Optimizer version 10.1 (Gurobi Optimization, LLC, 2024). We run each scenario on a single thread of an Intel Gold 6148 Skylake (2.4 GHz) processor and 35 GiB of memory. We use the original implementation of LS\*, which was written in C++ by Katayama (2020).

**Datasets.** We experiment with sets of instances with different characteristics. The first set contains the 24 GT instances (Hewitt et al., 2010). They are respectively subdivided into GTFL and GTFT, which stand for instances with loose and tight capacities. We chose to work with these instances because they are challenging and they are the ones used in the original work (Katayama, 2020). The second set contains instances built by a modernized version of the generator used to produce the Canad dataset (Larsen et al., 2023). Among others, the generator has parameters to control the size of the instance (number of arcs and commodities), and its topology. We set the instance size to be similar to GT (between 2,000 and 3,000 arcs, and around 200 commodities) to minimize the potential performance impact on LS\* due to calibration issues. The sets A and B use the grid topology, whereas C and D use the circular one. The sets A and C are exclusively used for training purposes and the sets B and D are used to test the method. Instances in test sets are about 30% larger than those in the training sets, to ensure that the method can generalize. Each generated set contains 30 instances, for a total of 144 instances when we include the GT instances.

### 4.1. Extended Runs and Sampling Routines

This section covers how to collect the required data to train the ML model. Features and labels are generated using the sampling routines and extended runs, respectively. Table 1 contains the optimality gaps given by the Gurobi default configuration (grb) after 12-hour runs. It gives us an estimate of the relative difficulty of each dataset. We notice how GT instances are meaningfully more challenging than the generated ones. Dataset A instances are the easiest with an average optimality gap of 1.3%, and dataset GTFL instances are the hardest with an average gap of 27.37%. The circular topology is marginally more difficult than the grid one. Test instances are slightly more challenging because they are larger.

Table 2 includes primal gaps for Gurobi and LS\* on each dataset. The main takeaway is the relative effectiveness of LS\* when specifically applied to GT instances. Gurobi has an average gap of 24.84% on GTFL, and LS\* has an average of 0%. This means the SOA heuristic systematically finds the best solution. The story is different for easier instances where Gurobi tends to outperform LS\*.

To generate the features, we run each sampling routine for 5 minutes and collect all the solutions found during the solution process. The results related to sampling routines on GT instances are summarized in Tables 3 and 4. The first table presents the quality of the best solution found by each routine, and the second contains the solution count. Gurobi (GRBP-5) is included as a reference. The LSFS routine is executed with three different column generation values: 3, 10 and 50. We note that LSFS-3 is the most effective routine in terms of solution quality with an average gap of 13.45%. An increase in the number of columns generated does not help because of the limited time budget. The RSS routine’s main advantage is its ability to generate a large number of solutions. It can generate more than 10 times more solutions than LSFS in the same period of time. We found similar results on the generated instances.

Table 1: Optimality gap (%) using Gurobi for 12-hour runs

Dataset	Quantiles			Mean	Geomean
	0.1	0.5	0.9		
GTFT	16.58	26.36	39.92	27.37	26.88
GTFL	12.04	30.31	44.69	29.10	27.15
A	0.88	1.35	1.62	1.30	2.28
B	1.85	2.46	2.90	2.40	3.37
C	1.01	1.57	2.20	1.64	2.60
D	2.61	3.85	4.84	3.82	4.74

Table 2: Primal gap (%) for 12-hour runs with baseline algorithms

Scenario	Quantiles			Mean	Geomean
	0.1	0.5	0.9		
grb-GTFT	0.18	13.34	31.05	14.70	9.09
grb-GTFL	2.51	26.33	44.84	24.84	18.45
grb-A	0.00	0.00	0.11	0.04	1.04
grb-B	0.00	0.35	0.62	0.35	1.32
grb-C	0.00	0.00	0.13	0.02	1.02
grb-D	0.16	0.64	1.08	0.67	1.63
LS*-GTFT	0.00	0.00	0.91	0.23	1.15
LS*-GTFL	0.00	0.00	0.00	0.00	1.00
LS*-A	0.00	0.05	0.61	0.19	1.16
LS*-B	0.00	0.00	0.51	0.12	1.10
LS*-C	0.00	0.09	0.53	0.21	1.18
LS*-D	0.00	0.00	0.00	0.01	1.01

#### 4.2. Tuning LS\* and the ML Model

The original implementation of LS\* was tuned on GT instances, so we do not change its default parameters for that dataset. We do, however, tune the capacity smoothing factor  $\lambda$  on the newly generated instances. To do so, we solve each instance in the training sets for 30 minutes with different values of  $\lambda$ . The results for dataset A are presented in Table 5. According to the average primal gap, the best value for  $\lambda$  is 0.05. It is also the best setting for dataset C. We chose to tune  $\lambda$  because it is the parameter that is explicitly calibrated in the original work. Since generated instances have a similar size than GT instances, we expect that the remaining parameters should be adequate.

Table 3: Primal gap (%) for sampling algorithms on dataset GT

Scenario	Quantiles			Mean	Geomean
	0.1	0.5	0.9		
LSFS-3	0.00	1.90	35.95	13.45	4.85
LSFS-10	0.00	15.17	93.13	41.60	17.84
LSFS-50	0.00	92.10	120.22	83.89	31.66
RSS	2.09	21.83	35.65	21.80	17.36
GRBP-5	12.12	26.96	40.91	25.96	24.40

Table 4: Solution count for sampling algorithms on dataset GT

Scenario	Quantiles			Mean	Geomean
	0.1	0.5	0.9		
LSFS-3	4.30	7.50	13.70	8.29	8.69
LSFS-10	3.00	4.50	8.70	5.54	6.13
LSFS-50	2.00	2.00	6.70	3.29	3.83
RSS	19.70	53.50	206.00	104.60	66.44
GRBP-5	2.00	4.00	5.00	3.71	4.53

The ML models are also tuned on the training sets A and C, and we study their performance with the balanced accuracy metric. It is the average of recall obtained on each class, which helps to manage the unbalanced class distribution. Table 6 reveals the effect of features sets on a selection of ML models. We test different combinations of basic (B), node graph (N) and sampling (S) features. The basic features are specific to the given arc, whereas node features contain aggregated information about the graph neighbourhood. The sampling features come from the solutions generated by the sampling routines. The linear binary classifier (Lin) is compared to four LGBM models with different values for `max_depth`, which controls its capacity. The main takeaway from this table is the importance of the sampling features. The average validation accuracy increases from 67% to 92% when sampling features are included. Note that BNS and BS have a similar validation accuracy, but BNS has a higher training accuracy. This effect is more noticeable as the model capacity increases. This is a sign of overfitting. It suggests that the graph features do not help the model to generalize.

Note that we considered other parameters of LGBM such as `min_data_in_leaf`, `learning_rate` and `num_iterations`, but they did not improve the performance. We noticed that LGBM is resilient to overfitting when `min_data_in_leaf` is set to its default value of 20. When we decrease it to 10 or 3, the training accuracy increases, but the validation accuracy decreases. Furthermore, we should mention that LGBM will print a warning “No further splits with positive gain”, when its depth is higher than or equal to 7. This reinforces the idea that many features are not useful or they are redundant. LGBM is able to find high-quality splits in the feature space early in the training process, but it struggles to find meaningful splits as the depth increases. Since both the linear model and LGBM have similar performance, the choice of the model is not critical. We choose to use LGBM because it is more flexible and offers more options to manage the class unbalance.

A deeper look into the raw data shows that misclassifications seem to be caused by a weak point in the sampling solutions. In particular, around 10% of open arcs in the best solution are always closed in sampling solutions. Sampling routines explore the solution space around the optimal solution of continuous relaxations, which is not always near the optimal solution of the original problem. When this is the case, the ML model has to rely on graph features to predict that an arc should be open. It must not be overlooked that the evidence in Table 6 points to the limited predictive power of graph features. This highlights a limitation of our ML-based approach. The model cannot accurately predict that an arc is open unless there is some evidence of that in the sampling solutions. This is a difficult problem to solve because sampling routines cannot reveal the optimal solution with any certainty.

**Bias adjustments.** An important challenge that emerges when using solutions as labels is the unbalanced nature of the dataset. In a solution of any given MCFND instance, it is rarely the case that the number of open arcs is equal

to the number of closed arcs. This is a problem because the ML model can learn to predict the majority class and still achieve a high accuracy. To mitigate this issue, one might consider using a resampling technique. We experimented with a random undersampling technique, but it did not improve performance when the flag `is_unbalance` is set to true in LGBM. The library can handle unbalanced datasets by itself to some extent. The default binary loss, however, is designed to minimize misclassifications with an equal weight for both classes. We suspect that this is inadequate because false negatives are not necessarily as costly as false positives. To address this, we test different weights in the loss (17) to skew the model in favor of one class over the other. This incentivizes the model to classify ambiguous cases as positive or negative depending on the bias. We study its effect with two models: LGBMW1 is biased towards the positive class ( $\omega_1 = 0.75$ ,  $\omega_0 = 0.25$ ), whereas LGBMW-1 is biased towards the negative class ( $\omega_0 = 0.75$ ,  $\omega_1 = 0.25$ ). We quantify the bias via false positive and false negative rates in Tables 7 and 8, respectively. Note that models are trained to predict the optimal reduction (see Section 4.3 for the justification). Hence, the positive label corresponds to removing an arc from the graph. As expected, LGBMW1 has the lowest false negative rate at 3% and LGBMW-1 has the lowest false positive rate at 5%. We demonstrate in Section 4.4 that including the bias adjustments can improve performance.

Table 5: Primal gap (%) for tuning  $\lambda$  on dataset A

Scenario	Quantiles			Mean	Geomean
	0.1	0.5	0.9		
LS*-0.05	0.00	0.00	0.00	0.00	1.00
LS*-0.15	0.02	0.05	0.13	0.07	1.07
LS*-0.25	0.08	0.14	0.23	0.15	1.15
LS*-0.35	0.07	0.19	0.27	0.18	1.17
LS*-0.45	0.14	0.24	0.35	0.26	1.25

Table 6: Train and validation balanced accuracy for different ML models and sets of features

Model	BN		BNS		BS	
	Train	Val	Train	Val	Train	Val
Lin	0.65	0.65	<b>0.92</b>	<b>0.92</b>	<b>0.92</b>	<b>0.92</b>
LGBM-D3	0.66	0.66	<b>0.92</b>	0.91	<b>0.92</b>	<b>0.92</b>
LGBM-D7	0.70	0.68	<b>0.93</b>	<b>0.92</b>	<b>0.93</b>	<b>0.92</b>
LGBM-D11	0.76	0.68	<b>0.96</b>	0.91	0.94	<b>0.92</b>
LGBM-D15	0.81	0.66	<b>0.96</b>	<b>0.92</b>	0.95	<b>0.92</b>
<b>num best out of 5</b>	0	0	5	3	3	5
<b>mean</b>	0.72	0.67	0.94	0.92	0.93	0.92

Table 7: False negative rate for different ML models

Scenario	Quantiles			Mean	Geomean
	0.1	0.5	0.9		
Lin	0.02	0.07	0.42	0.12	1.12
LGBM	0.04	0.10	0.39	0.16	1.15
LGBMW1	0.01	0.02	0.05	0.03	1.03
LGBMW-1	0.11	0.21	0.35	0.23	1.22

Table 8: False positive rate for different ML models

Scenario	Quantiles			Mean	Geomean
	0.1	0.5	0.9		
Lin	0.06	0.11	0.48	0.19	1.18
LGBM	0.06	0.11	0.30	0.15	1.14
LGBMW1	0.17	0.27	0.45	0.29	1.29
LGBMW-1	0.00	0.03	0.13	0.05	1.05

### 4.3. Algorithm Design Analysis

ML-OR hybrid approaches come with a unique challenge, which is related to the integration of a noisy prediction in a deterministic algorithm. To obtain good performance, the algorithm must be designed in a way that is resilient to a relatively high level of noise. Related works (Ding et al., 2020) use the local branching cut (18) to guide the search. We perceive it as a good choice in this context because it is a form of soft variable fixing. For that reason, our first experimentations revolve around the integration of that cut into the original implementation of  $LS^*$  to guide the search around the ML prediction.

In this section, we introduce an algorithm design analysis where we test different ways to integrate the ML predictions in an LS algorithm. To simulate the effect of the ML model, we use our database of solutions (DB) to get the correct label and then we add noise to it. The noise is a parameter of the simulation, and it is the probability that the label is flipped. We test four noise values between 0 and 0.15, which are close to the misclassification rate on the validation set. The three algorithms we test are:

- **LBH**: A local branching cut (18) is added relative to the ML prediction. The distance parameter is set such that 80% of arcs must respect the ML solution. We also set the variable parameters `VarHintVal` in Gurobi to the ML prediction.
- **LSWSH**: Similar to LBH, but the cut (18) and the hints are added to  $LS^*$  instead of Gurobi alone. We use the same distance parameter as LBH.
- **LSR**: This algorithm reduces the problem instead of guiding the search. Arcs are removed from the graph according to the ML prediction. However, if the reduction is applied as-is, the reduced problem is unlikely to be feasible. To guarantee feasibility, we add a repair routine that keeps arcs that are part of the best sampling solution. LSR uses the  $LS^*$  implementation without the CS routine.

We compare each algorithm on the GT dataset in Table 9.  $LS^*-3$  is a version of  $LS^*$  with only 3 columns generated, which significantly reduces the runtime of the CS routine. As expected, the best performing algorithms use DB with no noise. LSWSH, which adds a cut in  $LS^*$ , marginally benefits from the provided solution. This can be explained by an alignment issue. The best known solution is not necessarily feasible in the reduced problem produced by CS. LBH has most of the wins when the noise is null, but its performance decreases rapidly as the noise increases. The reason LBH performs so well with no noise is because `VarHintVal` is set to a high-quality solution that is feasible. LSR is the most resilient algorithm to noise. It is the only algorithm that is able to outperform  $LS^*$  with a noise of 5 or 10%. This is a promising result because it suggests that the ML model can be used to reduce the problem size without significantly affecting the quality of the solutions as long as the misclassification rate is not above 10%.

### 4.4. Main Results

In this section, we assemble the pieces and compare the performance of  $LS^*$  with LSR where the ML prediction is used to reduce the problem size. We begin our analysis with the results on GT instances. Tables 10 and 11 contain the primal gaps and primal integrals, respectively. We select a 2-hour time limit because at least one hour is required to run CS on some GT instances. For reference, we also include Gurobi with `MIPFocus` set to 1 (`grb-f1`), and LB with a cut around the incumbent. We experiment with three ML models for LSR: LGBM, LGBMW1 and LGBMW-1. In terms of primal gap,  $LS^*$  dominates with 16 wins out of 24 and 1.94 geomean. It is followed by LSR-LGBMW-1 with 3 wins and 4.94 geomean. The primal integral, however, reveals a different story.  $LS^*-3$  finds high-quality solutions

Table 9: Primal gap (%) for algorithms using DB on dataset GT for 2-hour runs

Scenario	Quantiles			Mean	Geomean	Wins
	0.1	0.5	0.9			
LS*	0.39	1.36	16.59	4.49	3.25	0
LS*-3	5.65	11.83	30.01	14.46	12.95	0
LSWSH-DB-0.0	0.19	0.70	42.94	17.32	3.21	0
LSR-DB-0.0	0.00	0.00	0.17	0.08	1.07	7
LBH-0.8-DB-0.0	0.00	0.00	0.00	0.02	1.02	10
LSWSH-DB-0.05	0.19	2.64	129.77	28.07	6.34	1
LSR-DB-0.05	0.45	2.48	6.03	3.15	3.53	1
LBH-0.8-DB-0.05	1.61	8.24	14.42	8.33	7.45	0
LSWSH-DB-0.1	0.38	1.48	132.03	32.58	6.00	0
LSR-DB-0.1	0.68	2.13	5.54	2.51	3.11	1
LBH-0.8-DB-0.1	1.26	11.51	32.92	15.77	10.96	0
LSWSH-DB-0.15	0.38	6.13	134.30	30.49	8.38	0
LSR-DB-0.15	1.44	6.27	10.31	6.20	6.22	0
LBH-0.8-DB-0.15	1.63	11.51	32.92	16.46	11.71	0

early with 12 wins and a geomean of 5.5. LSR-LGBMW-1 is the second best with 7 wins and a geomean of 7.13. The best primal integral is achieved by LS\*-3, but at the cost of a higher primal gap. The results reveal that LSR-LGBMW-1 has a good balance between quality and speed. The outperformance of LS\* can be justified by two factors. First, LS\* was extensively tuned on GT instances. Our experiments suggest that the CS routine is not as effective on the other instances that we generated. Second, each instance in GT has a unique structure, where the number of arcs and commodities varies significantly from one instance to the other. This is particularly challenging for ML methods, which are known to struggle when the underlying distribution of the data is not consistent.

We now turn our attention to the generated instances, which are more representative of the target application for this method. ML-CO hybrid approaches seek to perform well on instances sampled from a relatively narrow distribution, which is the case for the generated instances. Tables 12-15 contain the metrics of interest on test sets B and D. In both cases, LSR-LGBMW-1 outperforms LS\* in terms of quality and speed. Its average primal gap is 0.06% and 0.11%, which correspond to 14 and 16 wins, for datasets B and D, respectively. LS\* is the second-best algorithm with respect to this metric. In terms of primal integral, LSR-LGBMW-1 is also the best-performing algorithm with 32 wins (17+15) in total for both datasets. Furthermore, we remark that LB systematically outperforms grb-f1. Hence, it is a baseline that is relatively easy to implement and worth considering. Finally, we ought to mention that the performance of LSR-LGBM is competitive but it is consistently outperformed by LSR-LGBMW-1. This is a sign that the bias adjustments are beneficial.

**Key insights:** LS\* dominates on GT instances, as it excels in terms of primal gap but its slow CS routine negatively impacts its primal integral. LSR replaces CS with ML predictions which allows it to find high-quality solutions faster. Its primal integral is an order of magnitude better than LS\* on GT. On generated instances, LSR outperforms for both metrics, particularly when using the ML model with bias adjustments.

## 5. Perspectives and Future Work

The purpose of this study was to explore how to use an ML prediction to improve a SOA heuristic for the MCFND. To do so, we propose a methodology that takes advantage of supervised learning and off-the-shelf ML models. Given that the tools we use are highly accessible, we believe that our method is reasonably easy to reproduce and adapt to other problems similar to MCFND. Also, given that predictions are made on a local level, the method is not computationally expensive even for large instances. It constitutes a stepping stone for hybrid ML-OR approaches that still require further research to be fully operational.

A noteworthy contribution of this research is the development of sampling routines that generate solutions for the MCFND in an efficient manner. Sampling routines are a key component of our method, as they provide the necessary

Table 10: Primal gap (%) with ML models for 2-hour runs on dataset GT

Scenario	Quantiles			Mean	Geomean	Wins
	0.1	0.5	0.9			
LB-0.8	3.34	24.87	41.15	23.83	18.85	0
LS*	0.00	0.00	10.16	8.24	1.94	16
LS*-3	0.00	9.71	33.19	13.57	8.30	5
LSR-LGBM	2.78	8.19	21.95	10.76	9.34	0
LSR-LGBMW-1	0.06	5.91	9.77	6.54	4.94	3
LSR-LGBMW1	4.96	15.34	27.00	16.69	14.76	0
grb-f1	5.21	28.24	50.47	28.36	22.72	0

Table 11: Primal integral (%) with ML models for 2-hour runs on dataset GT

Scenario	Quantiles			Mean	Geomean	Wins
	0.1	0.5	0.9			
LB-0.8	6.69	28.83	39.43	24.95	20.72	0
LS*	0.05	30.94	70.43	34.09	16.57	5
LS*-3	0.26	5.13	16.38	6.93	5.50	12
LSR-LGBM	3.39	10.13	21.05	11.22	10.27	0
LSR-LGBMW-1	1.57	7.32	13.86	8.61	7.13	7
LSR-LGBMW1	5.75	16.80	25.88	16.73	14.99	0
grb-f1	2.97	14.01	24.66	14.07	12.32	0

information for the training process. We demonstrated that sampled solutions form highly informative features, which greatly increase the predictive power of the ML model. Furthermore, our study on feature sets revealed that graph features are not as useful as sampling features. ML models trained on graph features tend to overfit, and exhibit poor generalization properties. This is a key discovery that suggests that GNNs might not be necessary to solve CO problems even when the problem is defined on a graph.

Our experimental analysis contains important insights on the integration of ML predictions in a heuristic. We tested different ways to integrate the ML predictions in LS algorithms, and we showed that the best method is to use them to reduce the problem size by removing arcs from the graph. We found that it is most resilient to the noise in the ML predictions. This method is particularly effective on instances sampled from a uniform distribution, where it outperforms the original heuristic. We also showed that the ML model can be biased to improve the quality of the solutions.

Two limitations of our ML-based approach are worth mentioning. First, the method does not outperform the SOA heuristic on GT instances. A human-designed heuristic remains the best option to find high-quality solutions to the MCFND when the computational time budget is large. A deterministic heuristic is also more reliable than our hybrid approach. ML predictions always come with a level of uncertainty, which translates into a risk of compromising the quality of the final solution. For this reason, we believe that the best use case for our ML approach is to help find high-quality solutions early in the solution process. This is useful within a hyperparameter tuning framework (Hoos, 2012; Su et al., 2021) to find good configurations of a heuristic. Second, in the absence of direct indications from the sampled data, the ML model will fail to accurately predict the target. Often, arcs appear to be closed given the sampled solutions, but they are actually open in the best solution. This is a consequence of the fact that sampling routines cannot definitively identify the optimal solution. This constitutes an opportunity for further research. One possible avenue is to replace the human-designed sampling routines by a ML model that has learned to explore the solution space more efficiently.

Table 12: Primal gap (%) with ML models for 2-hour runs on dataset B

Scenario	Quantiles			Mean	Geomean	Wins
	0.1	0.5	0.9			
LB-0.8	0.07	0.33	0.56	0.32	1.31	1
LS*	0.00	0.04	0.28	0.10	1.10	11
LS*-3	0.71	1.45	1.98	1.35	2.30	0
LSR-LGBM	0.04	0.14	0.41	0.20	1.19	2
LSR-LGBMW-1	0.00	0.01	0.19	0.06	1.06	14
LSR-LGBMW1	0.06	0.22	0.53	0.27	1.25	2
grb-f1	0.08	0.40	0.67	0.40	1.38	1

Table 13: Primal integral (%) with ML models for 2-hour runs on dataset B

Scenario	Quantiles			Mean	Geomean	Wins
	0.1	0.5	0.9			
LB-0.8	0.25	0.39	0.60	0.41	1.41	0
LS*	0.11	0.22	0.37	0.23	1.22	5
LS*-3	0.75	1.22	1.45	1.19	2.17	0
LSR-LGBM	0.11	0.19	0.42	0.24	1.23	6
LSR-LGBMW-1	0.05	0.14	0.28	0.15	1.14	17
LSR-LGBMW1	0.12	0.23	0.53	0.29	1.28	4
grb-f1	0.41	0.67	0.80	0.64	1.63	0

Table 14: Primal gap (%) with ML models for 2-hour runs on dataset D

Scenario	Quantiles			Mean	Geomean	Wins
	0.1	0.5	0.9			
LB-0.8	0.28	0.46	0.69	0.49	1.48	0
LS*	0.00	0.12	0.34	0.14	1.13	6
LS*-3	1.73	2.69	3.81	2.75	3.66	0
LSR-LGBM	0.00	0.16	0.36	0.18	1.17	4
LSR-LGBMW-1	0.00	0.00	0.41	0.11	1.10	16
LSR-LGBMW1	0.00	0.20	0.58	0.24	1.23	4
grb-f1	0.48	0.99	1.26	0.90	1.88	0

Table 15: Primal integral (%) with ML models for 2-hour runs on dataset D

Scenario	Quantiles			Mean	Geomean	Wins
	0.1	0.5	0.9			
LB-0.8	0.28	0.44	0.66	0.48	1.47	0
LS*	0.18	0.34	0.46	0.32	1.31	3
LS*-3	1.63	2.35	2.96	2.37	3.32	0
LSR-LGBM	0.11	0.22	0.38	0.25	1.24	5
LSR-LGBMW-1	0.07	0.16	0.40	0.20	1.19	15
LSR-LGBMW1	0.14	0.24	0.57	0.30	1.29	7
grb-f1	0.50	0.95	1.08	0.87	1.85	0



## References

- Albertoni, R., Colantonio, S., Skrzypczyński, P., Stefanowski, J., 2023. Reproducibility of Machine Learning: Terminology, Recommendations and Open Issues. doi:p10.48550/arXiv.2302.12691. arXiv preprint arXiv:2302.12691.
- An, K., Lo, H.K., 2016. Two-phase stochastic program for transit network design under demand uncertainty. *Transportation Research Part B: Methodological* 84, 157–181. doi:p10.1016/j.trb.2015.12.009.
- Applegate, D., Bixby, R., Chvatal, V., Cook, W., 2006. Concorde TSP solver.
- Basallo-Triana, M.J., Bravo-Bastidas, J.J., Contreras, I., Cordeau, J.-F., Vidal-Holguín, C.J., 2023. Intermodal hub network design with generalized capacity constraints and non-synchronized train–truck operations. *Transportation Research Part B: Methodological* 174, 102770. doi:p10.1016/j.trb.2023.05.007.
- Bengio, Y., Lodi, A., Prouvost, A., 2021. Machine learning for combinatorial optimization: A methodological tour d’horizon. *European Journal of Operational Research* 290, 405–421. doi:p10.1016/j.ejor.2020.07.063.
- Blum, C., Pinacho, P., López-Ibáñez, M., Lozano, J.A., 2016. Construct, Merge, Solve & Adapt A new general algorithm for combinatorial optimization. *Computers & Operations Research* 68, 75–88. doi:p10.1016/j.cor.2015.10.014.
- Borisov, V., Leemann, T., Seßler, K., Haug, J., Pawelczyk, M., Kasneci, G., 2022. Deep Neural Networks and Tabular Data: A Survey. *IEEE Transactions on Neural Networks and Learning Systems* 1, 1–21. doi:p10.1109/TNNLS.2022.3229161.
- Caruana, R., Niculescu-Mizil, A., 2006. An empirical comparison of supervised learning algorithms, in: *Proceedings of the 23rd international conference on Machine learning*, Association for Computing Machinery, New York, NY, USA. pp. 161–168. doi:p10.1145/1143844.1143865.
- Chen, T., Guestrin, C., 2016. XGBoost: A Scalable Tree Boosting System, in: *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, Association for Computing Machinery, New York, NY, USA. pp. 785–794. doi:p10.1145/2939672.2939785.
- Chouman, M., Crainic, T.G., Gendron, B., 2017. Commodity Representations and Cut-Set-Based Inequalities for Multicommodity Capacitated Fixed-Charge Network Design. *Transportation Science* 51, 650–667. doi:p10.1287/trsc.2015.0665.
- Crainic, T.G., Frangioni, A., Gendron, B., 2001. Bundle-based relaxation methods for multicommodity capacitated fixed charge network design. *Discrete Applied Mathematics* 112, 73–99. doi:p10.1016/S0166-218X(00)00310-3.
- Crainic, T.G., Gendron, B., 2021. Exact Methods for Fixed-Charge Network Design, in: Crainic, T.G., Gendreau, M., Gendron, B. (Eds.), *Network Design with Applications to Transportation and Logistics*. Springer International Publishing, Cham, pp. 29–89. doi:p10.1007/978-3-030-64018-7-3.
- Crainic, T.G., Gendron, B., Hernu, G., 2004. A Slope Scaling/Lagrangian Perturbation Heuristic with Long-Term Memory for Multicommodity Capacitated Fixed-Charge Network Design. *Journal of Heuristics* 10, 525–545. doi:p10.1023/B:HEUR.0000045323.83583.bd.
- Dai, H., Khalil, E.B., Zhang, Y., Dilkina, B., Song, L., 2017. Learning combinatorial optimization algorithms over graphs, in: *Proceedings of the 31st International Conference on Neural Information Processing Systems*, Curran Associates Inc., Red Hook, NY, USA. pp. 6351–6361.
- Dhar, P., 2020. The carbon impact of artificial intelligence. *Nature Machine Intelligence* 2, 423–425. doi:p10.1038/s42256-020-0219-9.
- Ding, J.Y., Zhang, C., Shen, L., Li, S., Wang, B., Xu, Y., Song, L., 2020. Accelerating Primal Solution Findings for Mixed Integer Programs Based on Solution Prediction. *Proceedings of the AAAI Conference on Artificial Intelligence* 34, 1452–1459. doi:p10.1609/aaai.v34i02.5503.
- Fischetti, M., Lodi, A., 2003. Local branching. *Mathematical Programming* 98, 23–47. doi:p10.1007/s10107-003-0395-5.
- Friedman, J.H., 2001. Greedy Function Approximation: A Gradient Boosting Machine. *The Annals of Statistics* 29, 1189–1232.
- Gendron, B., Hanafi, S., Todosijevic, R., 2018. Matheuristics based on iterative linear programming and slope scaling for multicommodity capacitated fixed charge network design. *European Journal of Operational Research* 268, 70–81. doi:p10.1016/j.ejor.2018.01.022.
- Ghamlouche, I., Crainic, T.G., Gendreau, M., 2003. Cycle-Based Neighbourhoods for Fixed-Charge Capacitated Multicommodity Network Design. *Operations Research* 51, 655–667. doi:p10.1287/opre.51.4.655.16098.
- Ghamlouche, I., Crainic, T.G., Gendreau, M., 2004. Path Relinking, Cycle-Based Neighbourhoods and Capacitated Multicommodity Network Design. *Annals of Operations Research* 131, 109–133. doi:p10.1023/B:ANOR.0000039515.90453.1d.
- Ghanei, S., Contreras, I., Cordeau, J.-F., 2023. A two-stage stochastic collaborative intertwined supply network design problem under multiple disruptions. *Transportation Research Part E: Logistics and Transportation Review* 170, 102944. doi:p10.1016/j.tre.2022.102944.
- Gurobi Optimization, LLC, 2024. Gurobi Optimizer Reference Manual. URL: <https://www.gurobi.com>.
- Hamilton, W.L., 2020. *Graph Representation Learning*. Morgan & Claypool Publishers.
- Hanafi, S., Todosijevic, R., 2017. Mathematical programming based heuristics for the 0–1 MIP: a survey. *Journal of Heuristics* 23, 165–206. doi:p10.1007/s10732-017-9336-y.
- Hastie, T., Friedman, J., Tibshirani, R., 2001a. Kernel Methods, in: Hastie, T., Friedman, J., Tibshirani, R. (Eds.), *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*. Springer, New York, NY. Springer Series in Statistics, pp. 165–192. doi:p10.1007/978-0-387-21606-5-6.
- Hastie, T., Friedman, J., Tibshirani, R., 2001b. Linear Methods for Classification, in: Hastie, T., Friedman, J., Tibshirani, R. (Eds.), *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*. Springer, New York, NY. Springer Series in Statistics, pp. 79–113. doi:p10.1007/978-0-387-21606-5-4.
- Hewitt, M., Lehuédé, F., 2023. New formulations for the Scheduled Service Network Design Problem. *Transportation Research Part B: Methodological* 172, 117–133. doi:p10.1016/j.trb.2023.04.002.
- Hewitt, M., Nemhauser, G.L., Savelsbergh, M.W.P., 2010. Combining Exact and Heuristic Approaches for the Capacitated Fixed-Charge Network Flow Problem. *INFORMS Journal on Computing* 22, 314–325. doi:p10.1287/ijoc.1090.0348.
- Hoos, H.H., 2012. Automated Algorithm Configuration and Parameter Tuning, in: Hamadi, Y., Monfroy, E., Saubion, F. (Eds.), *Autonomous Search*. Springer, Berlin, Heidelberg, pp. 37–71. doi:p10.1007/978-3-642-21434-9-3.
- Kaplan, J., McCandlish, S., Henighan, T., Brown, T.B., Chess, B., Child, R., Gray, S., Radford, A., Wu, J., Amodei, D., 2020. Scaling Laws for Neural Language Models. doi:p10.48550/arXiv.2001.08361. arXiv preprint arXiv:2001.08361.
- Karimi-Mamaghan, M., Mohammadi, M., Pirayesh, A., Karimi-Mamaghan, A.M., Irani, H., 2020. Hub-and-spoke network design un-

- der congestion: A learning based metaheuristic. *Transportation Research Part E: Logistics and Transportation Review* 142, 102069. doi:p10.1016/j.tre.2020.102069.
- Katayama, N., 2020. MIP Neighborhood Search Heuristics for a Capacitated Fixed-Charge Network Design Problem. *Asia-Pacific Journal of Operational Research* 37, 2050009. doi:p10.1142/S0217595920500098.
- Katayama, N., Chen, M., Kubo, M., 2009. A capacity scaling heuristic for the multicommodity capacitated network design problem. *Journal of Computational and Applied Mathematics* 232, 90–101. doi:p10.1016/j.cam.2008.10.055.
- Ke, G., Meng, Q., Finley, T., Wang, T., Chen, W., Ma, W., Ye, Q., Liu, T.Y., 2017. LightGBM: A Highly Efficient Gradient Boosting Decision Tree. in: *Advances in Neural Information Processing Systems*, Curran Associates, Inc.
- Khalil, E., Bodic, P.L., Song, L., Nemhauser, G., Dilkina, B., 2016. Learning to Branch in Mixed Integer Programming. *Proceedings of the AAAI Conference on Artificial Intelligence* 30. doi:p10.1609/aaai.v30i1.10080.
- Khalil, E.B., Morris, C., Lodi, A., 2022. MIP-GNN: A Data-Driven Framework for Guiding Combinatorial Solvers. *Proceedings of the AAAI Conference on Artificial Intelligence* 36, 10219–10227. doi:p10.1609/aaai.v36i9.21262.
- Kidd, M.P., Darvish, M., Coelho, L.C., Gendron, B., 2024. A relax-and-restrict matheuristic for supply chain network design with facility location and customer due date flexibility. *Transportation Research Part E: Logistics and Transportation Review* 182, 103370. doi:p10.1016/j.tre.2023.103370.
- Kotary, J., Fioretto, F., Van Hentenryck, P., 2021. Learning hard optimization problems: A data generation perspective. *Advances in Neural Information Processing Systems* 34, 24981–24992.
- La Rocca, C.R., Cordeau, J.-F., Frejinger, E., 2024. One-shot Learning for MIPs with SOS1 Constraints. doi:p10.48550/arXiv.2403.09815. arXiv preprint arXiv:2403.09815.
- Larsen, E., BISAILLON, S., Cordeau, J.-F., Frejinger, E., 2023. Pseudo-random Instance Generators in C++ for Deterministic and Stochastic Multi-commodity Network Design Problems. doi:p10.48550/arXiv.2312.10717. arXiv preprint arXiv:2312.10717.
- Liu, D., Perreault, V., Hertz, A., Lodi, A., 2022. A machine learning framework for neighbor generation in metaheuristic search. doi:p10.48550/arXiv.2212.11451. arXiv preprint arXiv:2212.11451.
- Lodi, A., Zarpellon, G., 2017. On learning and branching: a survey. *TOP* 25, 207–236. doi:p10.1007/s11750-017-0451-6.
- Lubin, M., Dowson, O., Garcia, J.D., Huchette, J., Legat, B., Vielma, J.P., 2023. JuMP 1.0: recent improvements to a modeling language for mathematical optimization. *Mathematical Programming Computation* 15, 581–589. doi:p10.1007/s12532-023-00239-3.
- Magnanti, T.L., Wong, R.T., 1984. *Network Design and Transportation Planning: Models and Algorithms*. *Transportation Science* 18, 1–55. doi:p10.1287/trsc.18.1.1.
- Minoux, M., 1989. Networks synthesis and optimum network design problems: Models, solution methods and applications. *Networks* 19, 313–360. doi:p10.1002/net.3230190305.
- Paraskevopoulos, D.C., Gürel, S., Bektaş, T., 2016. The congested multicommodity network design problem. *Transportation Research Part E: Logistics and Transportation Review* 85, 166–187. doi:p10.1016/j.tre.2015.10.007.
- Quilliot, A., 2014. *Network Design Problems: Fundamental Methods*, in: *Applications of Combinatorial Optimization*. John Wiley & Sons, Ltd, pp. 253–289. doi:p10.1002/9781119005384.ch9.
- Real, L.B., Contreras, I., Cordeau, J.-F., de Camargo, R.S., de Miranda, G., 2021. Multimodal hub network design with flexible routes. *Transportation Research Part E: Logistics and Transportation Review* 146, 102188. doi:p10.1016/j.tre.2020.102188.
- Shaw, P., 1998. Using Constraint Programming and Local Search Methods to Solve Vehicle Routing Problems, in: Maher, M., Puget, J.F. (Eds.), *Principles and Practice of Constraint Programming — CP98*, Springer, Berlin, Heidelberg. pp. 417–431. doi:p10.1007/3-540-49481-2-30.
- Su, M., Liang, B., Ma, S., Xiang, C., Zhang, C., Wang, J., 2021. Automatic Machine Learning Method for Hyper-parameter Search. *Journal of Physics: Conference Series* 1802, 032082. doi:p10.1088/1742-6596/1802/3/032082.
- Sun, Y., Ernst, A., Li, X., Weiner, J., 2021. Generalization of machine learning for problem reduction: A case study on travelling salesman problems. *OR Spectrum* 43, 607–633. doi:p10.1007/s00291-020-00604-x.
- Turner, M., Koch, T., Serrano, F., Winkler, M., 2023. Adaptive Cut Selection in Mixed-Integer Linear Programming. doi:p10.48550/arXiv.2202.10962. arXiv preprint arXiv:2202.10962.
- Wallingford, M., Ramanujan, V., Fang, A., Kusupati, A., Mottaghi, R., Kembhavi, A., Schmidt, L., Farhadi, A., 2023. Neural Priming for Sample-Efficient Adaptation. doi:p10.48550/arXiv.2306.10191. arXiv preprint arXiv:2306.10191.
- Yaghini, M., Foroughi, A., 2014. ACO-Based Neighborhoods for Fixed-charge Capacitated Multi-commodity Network Design Problem. *International Journal of Transportation Engineering* 1, 311–334. doi:p10.22119/ijte.2014.6301.
- Yaghini, M., Rahbar, M., Karimi, M., 2013. A hybrid simulated annealing and column generation approach for capacitated multicommodity network design. *Journal of the Operational Research Society* 64, 1010–1020. doi:p10.1057/jors.2012.114.