

Machine-learning-based arc selection for constrained shortest path problems in column generation

M. Morabit, G. Desaulniers, A. Lodi

G-2021-72

December 2021

La collection *Les Cahiers du GERAD* est constituée des travaux de recherche menés par nos membres. La plupart de ces documents de travail a été soumis à des revues avec comité de révision. Lorsqu'un document est accepté et publié, le pdf original est retiré si c'est nécessaire et un lien vers l'article publié est ajouté.

Citation suggérée : M. Morabit, G. Desaulniers, A. Lodi (Décembre 2021). Machine-learning-based arc selection for constrained shortest path problems in column generation, Rapport technique, Les Cahiers du GERAD G- 2021-72, GERAD, HEC Montréal, Canada.

Avant de citer ce rapport technique, veuillez visiter notre site Web (<https://www.gerad.ca/fr/papers/G-2021-72>) afin de mettre à jour vos données de référence, s'il a été publié dans une revue scientifique.

The series *Les Cahiers du GERAD* consists of working papers carried out by our members. Most of these pre-prints have been submitted to peer-reviewed journals. When accepted and published, if necessary, the original pdf is removed and a link to the published article is added.

Suggested citation: M. Morabit, G. Desaulniers, A. Lodi (December 2021). Machine-learning-based arc selection for constrained shortest path problems in column generation, Technical report, Les Cahiers du GERAD G-2021-72, GERAD, HEC Montréal, Canada.

Before citing this technical report, please visit our website (<https://www.gerad.ca/en/papers/G-2021-72>) to update your reference data, if it has been published in a scientific journal.

La publication de ces rapports de recherche est rendue possible grâce au soutien de HEC Montréal, Polytechnique Montréal, Université McGill, Université du Québec à Montréal, ainsi que du Fonds de recherche du Québec – Nature et technologies.

Dépôt légal – Bibliothèque et Archives nationales du Québec, 2021
– Bibliothèque et Archives Canada, 2021

The publication of these research reports is made possible thanks to the support of HEC Montréal, Polytechnique Montréal, McGill University, Université du Québec à Montréal, as well as the Fonds de recherche du Québec – Nature et technologies.

Legal deposit – Bibliothèque et Archives nationales du Québec, 2021
– Library and Archives Canada, 2021

Machine-learning-based arc selection for constrained shortest path problems in column generation

Mouad Morabit [a, b, c](#)

Guy Desaulniers [a, b](#)

Andrea Lodi [a, b, c, d](#)

^a *Department of Mathematics and Industrial Engineering, Polytechnique Montréal, Montréal (Qc) Canada, H3C 3A7*

^b *GERAD, Montréal (Qc), Canada, H3T 1J4*

^c *Canada Excellence Research Chair in Data Science for Real-Time Decision-Making, Polytechnique Montréal*

^d *Jacobs Technion – Cornell Institute, Cornell Tech and Technion – IIT, New York, NY 10044, United States*

`mouad.morabit@polymtl.ca`

`guy.desaulniers@polymtl.ca`

`andrea.lodi@cornell.edu`

December 2021

Les Cahiers du GERAD

G–2021–72

Copyright © 2021 GERAD, Morabit, Desaulniers, Lodi

Les textes publiés dans la série des rapports de recherche *Les Cahiers du GERAD* n'engagent que la responsabilité de leurs auteurs. Les auteurs conservent leur droit d'auteur et leurs droits moraux sur leurs publications et les utilisateurs s'engagent à reconnaître et respecter les exigences légales associées à ces droits. Ainsi, les utilisateurs:

- Peuvent télécharger et imprimer une copie de toute publication du portail public aux fins d'étude ou de recherche privée;
- Ne peuvent pas distribuer le matériel ou l'utiliser pour une activité à but lucratif ou pour un gain commercial;
- Peuvent distribuer gratuitement l'URL identifiant la publication.

Si vous pensez que ce document enfreint le droit d'auteur, contactez-nous en fournissant des détails. Nous supprimerons immédiatement l'accès au travail et enquêterons sur votre demande.

The authors are exclusively responsible for the content of their research papers published in the series *Les Cahiers du GERAD*. Copyright and moral rights for the publications are retained by the authors and the users must commit themselves to recognize and abide the legal requirements associated with these rights. Thus, users:

- May download and print one copy of any publication from the public portal for the purpose of private study or research;
- May not further distribute the material or use it for any profit-making activity or commercial gain;
- May freely distribute the URL identifying the publication.

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Abstract : Column generation is an iterative method used to solve a variety of optimization problems. It decomposes the problem into two parts: a master problem, and one or more pricing problems (PP). The total computing time taken by the method is divided between these two parts. In routing or scheduling applications, the problems are mostly defined on a network, and the PP is usually an NP-hard shortest path problem with resource constraints. In this work, we propose a new heuristic pricing algorithm based on machine learning. By taking advantage of the data collected during previous executions, the objective is to reduce the size of the network and accelerate the PP, keeping only the arcs that have a high chance to be part of the linear relaxation solution. The method has been applied to two specific problems: the vehicle and crew scheduling problem in public transit and the vehicle routing problem with time windows. Reductions in computational time of up to 40% can be obtained.

Keywords: Column generation, machine learning, arc selection, heuristic pricing, vehicle routing and crew scheduling.

Acknowledgements: We thank Giro Inc. and the Natural Sciences and Engineering Research Council of Canada for their financial support through the grant CRDPJ 520349-17.

1 Introduction

Column generation (CG, [8]) is an iterative method that splits the original problem into two parts. The first part is a restricted master problem (RMP) that corresponds to the original linear problem (LP) but restricted to a subset of variables, which is often solved using the Simplex method. The second part is the subproblem, also called the pricing problem (PP). The role of the PP is to generate new columns of negative reduced cost (in case of a minimization problem) that can improve the current RMP solution. The method stops when no more columns can be generated. Otherwise, the new columns are added to the RMP to start a new iteration. To obtain integer solutions, CG is often embedded within a branch-and-bound process, where CG is used to solve the linear relaxation at each node of the tree. In this case the method is referred to as branch-and-price, see [3, 8].

Column generation has been successfully used to solve a variety of optimization problems, notably crew scheduling and vehicle routing problems arising in freight, urban and airline transportation. For most of these problems, the master problem corresponds to a set partitioning problem with side constraints, and the PP is most likely a shortest path problem with resource constraints (SPPRC) or one of its variants. The success of using CG for this type of problems is also due to the efficient methods for solving the SPPRC, which are mainly dynamic programming algorithms [10, 20]. The goal is to find a path with the least cost between a source and a destination, while respecting the constraints on the resources (i.e., time, load of a vehicle, break duration, ...). It is well known that the SPPRC is NP-hard, but the methods developed to solve it are quite efficient (i.e., pseudopolynomial time complexity), and have allowed to solve large instances in reasonable times. The algorithms based on dynamic programming are also flexible enough to tackle different variants of the problems and take into account complex constraints that influence the feasibility of a route or a schedule. An improvement over the basic SPPRC algorithm is that of [25] who proposed a bi-directional search that propagates labels in both directions at the same time, i.e., from the source node to the destination, and backward from the destination to the source. Forward and backward labels residing in the same intermediate node are possibly merged to form complete feasible paths.

The standard SPPRC often arises in problems that are defined on acyclic time-space networks, among others, in vehicle and crew scheduling problems [7, 17] and rostering problems [15]. Another well-known variant that is more difficult to solve is the elementary SPPRC (ESPPRC) occurring in vehicle routing problems (VRP) [5, 6, 12]. The use of an exact method to solve the ESPPRC is known to yield a tight lower bound, but given the difficulty of the problem, other alternatives based on relaxations and heuristics have been explored; see for example: ng-routes [2], SPPRC-k-cyc [21] and partial elementarity [9]. The goal of these later approaches is to find a trade-off between the quality of the bound and the difficulty of solving the PP. Several SPPRC heuristics have also been proposed in the literature. For instance, some authors relax the dominance rule applied when solving the problem with a labeling algorithm by only dominating on a subset of the resources [9]. Another option is to keep only a limited number of labels at each node [14]. Other proposed strategies are based on reducing the size of the network: one consists of sorting the incoming and outgoing arcs of each node by their reduced cost [9], where only a subset of the arcs with the least costs are retained. In [14], the authors use an extension of Kruskal's algorithm to build disjoint spanning trees from the network, and then, they consider only the arcs used on those trees in addition to the depot outgoing and incoming arcs. All of the above strategies can be used in the same execution and some can even be combined and used simultaneously. In most cases, an exact algorithm is invoked at the last iteration to prove the solution optimality.

Combining machine learning methods with operations research and combinatorial optimization algorithms [4] has been the subject of many studies over the last few years. These studies generally fall into two categories. The first category is based on what is called *imitation learning*, where the learner tries to imitate an expert in order to get a fast prediction to make decisions that are otherwise computationally expensive. In this category one seeks to reproduce decisions that are as close as possible to those of the expert. The presence of data from previous executions of the expert is therefore

necessary. The second category is *reinforcement learning*, also called learning by experience as the name indicates, where the learner (i.e., agent) interacts with the environment and tries to optimize the results (i.e., maximize the rewards) of the decisions (i.e., actions) it makes. In this context, the agent learns by trial and error while optimizing its reward function. After a number of iterations (i.e., epochs), a policy is learned telling the agent what is the best decision to make in each situation (i.e., state). The approach described in this article falls in the first category, that is learning by imitation and is thus based on supervised learning.

Several papers have covered the branch-and-bound method in mixed integer programming context, many of which seek to imitate *strong branching* which is known to be computationally heavy but very effective, such as the works of [1, 16, 22]. Several other methods exist and are covered in the survey of [23]. For a more general overview of ML methods that have been developed and used in the context of combinatorial optimization (CO), readers are invited to check the survey of [4]. This survey covers various machine learning methods and describes how they can be integrated to help solve CO problems. In the context of CG and branch-and-price, there are some very interesting papers worth mentioning. The authors in [26] proposed a regression model for estimating an upper bound on the optimal value of the PP, which is then used to speed up the PP solution process. At each iteration of the CG, features based on the dual values are fed to the model. The learning is done in a so-called “online” way (i.e., the learning is performed at the same time as the optimization), the loss function used during training penalizes the model according to the difference between the predicted bound and the true bound value (a high penalty is applied when the value of the bound is underestimated) and a discount factor is also used to give more importance to the last iterations. Also in the CG context, this time on the master problem side, a recent paper by [24] presents a learned column selector that selects, from a large set of generated columns, the columns to be added to the RMP at each CG iteration. The goal is to add the most promising columns that have a high probability of improving the current solution. The model used is based on graph neural networks, i.e., neural networks applied to graph-structured data. The problem is represented in a bipartite graph to model the variable-constraint relations. Differently from the current paper, the strategy in [24] is designed for problems that take a larger portion of the computing time in solving the RMP.

As previously mentioned, CG can be applied to several optimization problems, so both RMP and PP can be different depending on the application at hand. In this article, the focus is more on crew scheduling and vehicle routing problems, where in most cases, the PP is a shortest path problem with resource constraints or one of its variants, and it is most likely the part that consumes the largest portion of the computing time. The purpose of this paper is to propose a heuristic pricing algorithm based on machine learning. The method takes advantage of the data collected from previous executions in order to learn a classifier that will select the arcs that have a greater chance to be used or to be part of an optimal solution. A good reduction of the network size decreases the computing time of the PP without increasing significantly the number of iterations, and thus reducing the overall total time. To ensure the exactness of the overall algorithm, the complete network is used when the reduced network fails to generate new columns.

For our tests, we chose to use two well-known problems in the literature: the vehicle and crew scheduling problem (VCSP) and the vehicle routing problem with time windows (VRPTW). Each of these two problems have a different network structure (i.e., different type of nodes and arcs, etc), but nothing prevents from restricting the selection strategy to a specific type of arcs, making the methods described here general enough to be applied to various problems.

The remainder of the paper is organized as follows. Section 2 is devoted to some preliminaries and details about CG and the dynamic programming method used to solve the SPPRC. In Section 3, we describe the proposed selection strategy. Sections 4 and 5 will cover our two case studies, the VCSP and VRPTW respectively, highlighting the implementation differences between the two problems and reporting the computational results for each of them. Finally, we draw some conclusions in Section 6.

2 Preliminaries

In the CG context, and by focusing on the linear relaxation of the problem, let us consider the following master problem (MP) formulated as follows:

$$z_{MP}^* := \min_x \sum_{p \in \Omega} c_p x_p \quad (1)$$

$$\text{s.t.} \quad \sum_{p \in \Omega} \mathbf{a}_p x_p = \mathbf{b}, \quad (2)$$

$$x_p \geq 0, \quad \forall p \in \Omega, \quad (3)$$

where Ω represents the set of variable indices (e.g., feasible routes or schedules), $c_p \in \mathbb{R}$ the variable cost, $\mathbf{a}_p \in \mathbb{R}^m$ the constraints coefficients and $\mathbf{b} \in \mathbb{R}^m$ the constraints right-hand side vector. When $|\Omega|$ is large and the variables cannot be enumerated explicitly, we consider only a subset $\mathcal{F} \subseteq \Omega$ of the variables, obtaining a restricted version of the problem above called a restricted master problem, RMP. At each CG iteration, the RMP is optimized, and an optimal solution x is obtained along with a dual solution $\pi \in \mathbb{R}^m$ associated with the constraints (2). The dual values π are then used to define the PP $\min_{p \in \Omega} \{c_p - \pi^T \mathbf{a}_p\}$ and find new columns with negative reduced cost by solving it. The method stops when no such columns exist. In our applications, the variables represent either routes or schedules, and they can be generated by solving one or more PP (in our case a SPPRC or an ESPPRC).

2.1 SPPRC formulation

Let $G = (V, A)$ be a directed graph with a set of nodes V that include the source and destination nodes denoted by s and t respectively, and A the set of arcs. Let R be the set of resources. For each node $i \in V$, we define T_i^r as the value of resource $r \in R$ accumulated on a partial path from source node s to node i , and resource windows $[\underline{w}_i^r, \bar{w}_i^r]$, $\underline{w}_i^r, \bar{w}_i^r \in \mathbb{R}$ restricting the values that resource $r \in R$ can take on node i . For each arc $(i, j) \in A$, we define the resource consumptions $t_{ij}^r \in \mathbb{R}$, $r \in R$ and the cost c_{ij} . Let X_{ij} be the decision variables of the model, which represent the flow on the arcs $(i, j) \in A$. The SPPRC is then formulated as follows:

$$\min \sum_{(i,j) \in A} c_{ij} X_{ij} \quad (4)$$

$$\text{s.t.} \quad \sum_{j \in V} X_{ij} - \sum_{j \in V} X_{ji} = \begin{cases} +1 & \text{if } i = s \\ 0 & \text{if } i \in V \setminus \{s, t\} \\ -1 & \text{if } i = t \end{cases} \quad (5)$$

$$X_{ij}(T_i^r + t_{ij}^r - T_j^r) \leq 0, \quad \forall r \in R, \forall (i, j) \in A, \quad (6)$$

$$\underline{w}_i^r \leq T_i^r \leq \bar{w}_i^r, \quad \forall r \in R, \forall i \in V, \quad (7)$$

$$X_{ij} \in \{0, 1\}, \quad \forall (i, j) \in A, \quad (8)$$

where the objective (4) minimizes the total cost of the path, the constraints (5) ensure the flow conservation along the path, while the constraints (6) model the resource consumption for each resource $r \in R$ on arc $(i, j) \in A$. Constraints (7) make sure that the resource values accumulated respect the resource intervals at each node and constraints (8) are the binary requirements on the flow variables X_{ij} . Note that there are more complex SPPRC than those expressed by (4)–(8), i.e., using complex resource extension functions that, given the resource values T_i^r at node i , provide a lower bound on the resource values at node j .

As mentioned above, at each CG iteration, the dual values are used find new columns to add in the RMP. This dual information is included in the SPPRC by using a modified cost $\bar{c}_{ij} = c_{ij} - \sum_{k=1}^m \rho_{ij}^k \pi_k$ on each arc $(i, j) \in A$, where ρ_{ij}^k is a binary parameter, taking value 1 if the arc $(i, j) \in A$ contributes

to the master problem constraint k , 0 otherwise. The cost of a path $p \in \Omega$ can be written in terms of the individual arcs:

$$c_p = \sum_{(i,j) \in A} c_{ij} b_{ij}^p, \quad (9)$$

where b_{ij}^p is a binary parameter equal to 1 if the arc $(i, j) \in A$ is included in path p . The constraints coefficients $a_p^k, k = \{1, 2, \dots, m\}$, where $\mathbf{a}_p = [a_p^1, a_p^2, \dots, a_p^m]^T$ can also be written as

$$a_p^k = \sum_{(i,j) \in A} \rho_{ij}^k b_{ij}^p. \quad (10)$$

The reduced cost c_p of a variable/path can then be defined as the sum of the modified costs of the arcs composing the path, namely

$$\bar{c}_p = c_p - \pi^T \mathbf{a}_p = c_p - \sum_{k=1}^m \pi_k a_p^k = \sum_{(i,j) \in A} b_{ij}^p (c_{ij} - \sum_{k=1}^m \rho_{ij}^k \pi_k) = \sum_{(i,j) \in A} b_{ij}^p \bar{c}_{ij}. \quad (11)$$

Therefore, solving the program (4)–(8) using the modified costs in (4) yields a column with the least reduced cost. However, by solving the program using the dynamic programming formulation, it is possible to obtain several paths at once, which is another strong point for using these methods, as it has proven to speed up the resolution and minimize the number of iterations performed.

2.2 Labeling algorithm

In dynamic programming, shortest path problems are solved by a labeling algorithm. Such an algorithm starts from the trivial path that contains only the source node s , then extends it recursively along all arcs that yields a feasible path until reaching the destination node t . A path $P = (v_0, v_1, \dots, v_p)$ in G represents a sequence of visited nodes such that $(v_{i-1}, v_i) \in A, v_i \in V, \forall i = 1, 2, \dots, p$ where v_p is the last node of the partial path and $v_0 = s$. In the algorithm, a label L encodes information about a partial path P such as the reduced cost $\bar{c}(L)$, the set of nodes $V(L)$ visited by the path, the last node visited by the path $v(L)$ and the accumulated quantity $T^r(L)$ of each resource $r \in R$. For a label L ending at node $v_i (v_i \neq t)$, a new label L' is obtained when an extension is performed along an arc $(i, j) \in A$, such that

$$v(L') = j, \quad (12)$$

$$V(L') = V(L) \cup \{j\}, \quad (13)$$

$$\bar{c}(L') = \bar{c}(L) + \bar{c}_{ij}, \quad (14)$$

$$T^r(L') = \max\{\underline{w}_j^r, T^r(L) + t_{ij}^r\}, \quad \forall r \in R. \quad (15)$$

The label L' also keeps a reference to its predecessor label L , in order to build the complete path when the algorithm ends. The label is considered infeasible if $T^r(L') > \bar{w}_j^r$ and is therefore deleted. In addition, to the infeasible labels that get rejected when extending the labels, a dominance rule can be applied to eliminate non-promising paths. A label L_1 is said to be dominating a label L_2 if the following conditions are verified:

$$v(L_1) = v(L_2), \quad (16)$$

$$V(L_2) \subseteq V(L_1), \quad (17)$$

$$\bar{c}(L_1) \leq \bar{c}(L_2), \quad (18)$$

$$T^r(L_1) \leq T^r(L_2), \quad \forall r \in R \quad (19)$$

The performance of the labeling algorithm is closely linked to the choice of the dominance rule, a good choice of the latter allows to reduce the search space and to eliminate a maximum of labels, allowing to

accelerate the resolution. Note that the validity of a dominance rule depends closely on the variant of the SPPRC being solved, as well as the resource consumptions on the arcs. For simplicity, some details have been omitted, such as resource extension functions (REFs, [19]), which are considered being a more general way to define the resource arc consumptions, and extensions of the basic algorithm for the ESPPRC case to deal with cycle elimination. Readers are referred to [20] for a good overview over the details omitted here.

At the end of the algorithm, several labels are obtained at the destination node t . The labels with a non-negative reduced cost are discarded, and those with a negative reduced cost are kept to build the new columns to be added in the RMP. The different steps discussed in this section are summarized in Figure 1.

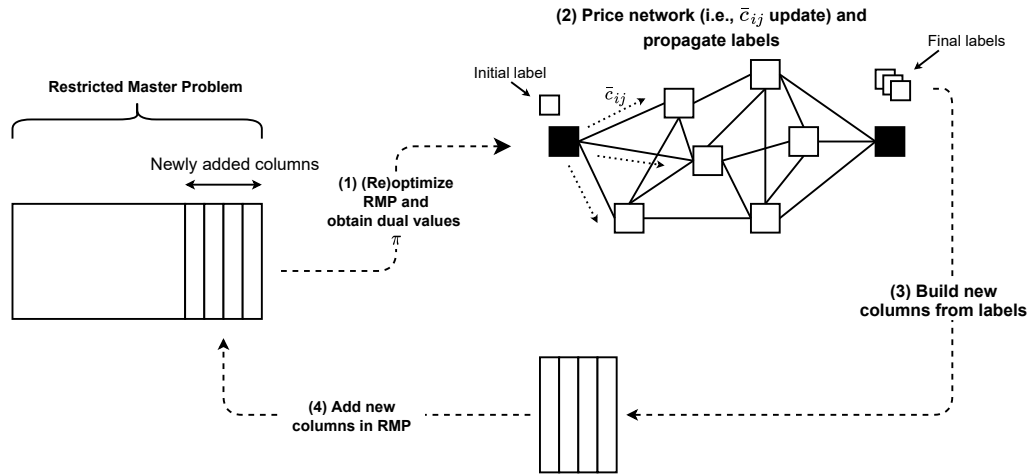


Figure 1: Summary of the different steps of a CG iteration.

3 Methodology

The goal of the project is to reduce the size of the underlying SPPRC network by keeping only the most promising arcs, therefore obtaining a reduced network. At each iteration, the same reduced network is used as long as it generates a satisfactory number of new negative reduced cost columns. The arc selection is thus performed only once before starting to solve the problem. In machine learning terms, this is a classification problem (i.e., supervised machine learning) and requires a labeled dataset $\mathcal{D} = \{(\mathbf{x}_1, y_1), (\mathbf{x}_2, y_2), \dots, (\mathbf{x}_n, y_n)\}$, where $n = |\mathcal{D}|$ and each entry is a tuple that represents an input/output pair, the inputs \mathbf{x}_i are called features and y_i are the desired outputs, also called labels (not the same notion of “label” as in a labeling algorithm) given by an expert or known in advance. Given the dataset, the learning algorithm learns to predict the outputs y_i from the feature vectors \mathbf{x}_i . In our case, it is a binary classification problem (i.e., $y_i \in \{0, 1\}$), so each arc is either selected or not. During the training phase, the algorithm tries to minimize the difference between the true known values y_i and the values predicted by the model denoted by \hat{y}_i . However, the goal is to be able to generalize and give good predictions for inputs never encountered during training, so a portion of the dataset is reserved for testing, in order to evaluate more accurately the performance of the model.

3.1 Features

The extracted features represent the characteristics of each individual arc $(i, j) \in A$. They can be different depending on the problem in hand and the structure of the underlying network. For the VCSP and VRPTW considered in this paper, some of the features are similar, namely:

- Cost of the arc c_{ij}

- For each resource $r \in R$, the resource consumption T_{ij}^r along the arc (i, j)
- Number of arcs leaving node i
- Number of arcs entering node j
- For each resource $r \in R$, the minimum, maximum and average resource consumptions along the arcs leaving node i
- For each resource $r \in R$, the minimum, maximum and average resource consumptions along the arcs entering node j
- For each resource $r \in R$, the upper and lower resource bounds for nodes i and j .

The other features specific to each problem will be described in their dedicated sections.

3.2 Labels

In supervised learning, the labels are either known in advance or assigned by an expert. In our case, we have to define what are the promising arcs to keep. The algorithm followed to assign the labels is in most cases computationally expensive, so we try to collect as much data as possible, then train a machine learning model that gives predictions in a reasonably fast computing time. Let \mathcal{C}_i be the set of columns generated at iteration i , and A_c the set of arcs visited by the path represented by column $c \in \mathcal{C}_i$. A promising arc is an arc that has been used at least once to generate columns at any iteration. The idea is that the generated columns must have dominated many other columns during the resolution of the PP and are therefore all good candidates, and so are the arcs composing their paths. According to the tests on the VCSP and VRPTW instances, the percentage of arcs used at least once during the resolution of the linear relaxation by CG does not exceed 18%, and this number decreases when moving to larger instances. Another idea is to consider only the positive basic columns (i.e., columns in the optimal basis of an RMP with a positive value), which results on a larger reduction of the graph, but according to the results, the number of iterations slightly increases compared to considering all the generated columns and the gain is therefore lost. Figure 2 shows a visual comparison of the number of arcs for a VRPTW instance, considering different scenarios: (a) all arcs, (b) arcs in the generated columns, (c) arcs in an RMP solution, and (d) arcs in the linear relaxation solution. The figure also specifies for each scenario the percentage of arcs selected with respect to the complete arc set.




Figure 2: Number of arc comparisons across different scenarios

3.3 Data collection

Now that we have defined the arcs to be selected and the features, we can proceed and put all the pieces together to define our algorithm for data collection described in Algorithm 1.

The steps of the algorithm are fairly straightforward. First, we start with an empty dataset in Step (1) and we initialize all arc labels with the value 0 in Steps (2)–(4). In (5) an initial solution is obtained, either by using a heuristic or a two-phase simplex algorithm. Then, at each CG iteration, the RMP is reoptimized and the PP is solved, generating a new set of negative reduced cost columns

Algorithm 1 Data collection : features and labels extraction

```

1:  $\mathcal{D} = \emptyset$ 
2: for each  $a \in A$  do
3:    $y_a = 0$ 
4: end for
5:  $\pi \leftarrow \text{initialRMPsolution}()$ 
6: for each CG iteration  $i$  do
7:    $\text{priceNetwork}(G, \pi)$ 
8:    $\mathcal{C}_i \leftarrow \text{generateColumns}(G)$ 
9:   if  $\mathcal{C}_i == \emptyset$  then
10:    break
11:   end if
12:   for each  $c \in \mathcal{C}_i$  do
13:     for each  $a \in A_c$  do
14:        $y_a = 1$ 
15:     end for
16:   end for
17:    $\text{addColumns}(\mathcal{C}_i)$ 
18:    $\pi \leftarrow \text{reoptimizeRMP}()$ 
19: end for
20: for each  $a \in A$  do
21:    $\mathbf{x}_a \leftarrow \text{extractData}(a)$ 
22:    $\mathcal{D} \leftarrow \mathcal{D} \cup \{(\mathbf{x}_a, y_a)\}$ 
23: end for

```

} Initialization

} Solution is optimal

} Label assignement
to "promising" arcs

} Data collection

\mathcal{C}_i (Steps (7)–(8)). In case the PP defined over the full network G fails to generate any column, the CG process stops and the current solution is optimal (Steps (9)–(11)). Otherwise, for each newly generated column $c \in \mathcal{C}_i$, the label $y_a = 1$ is assigned to each arc $a \in A_c$ composing the path represented by the column c (Steps (12)–(16)) and the generated columns are added to the RMP which is repotimized (Steps (17)–(18)). Finally, when the CG process ends, for each arc $a \in A$ the pair of features \mathbf{x}_a and labels y_a are collected and added to the dataset \mathcal{D} (Steps (20)–(23)). Note that it is necessary to solve several instances of the problem to optimality to collect enough data. The features can be extracted before solving the problem but for the labels it is necessary to wait until the end of the optimization.

3.4 ML pricing algorithm

Once the data is available, the preprocessing phase is performed (i.e., normalization of the values and the encoding of categorical features) before starting the training. Various classification algorithms can be used in our case: more details about the training phase and its results are provided in the next sections. Once the training phase is completed, the learned model can be used to select the arcs and build the reduced network. Before starting the first CG iteration, the learned model takes the features as an input and gives predictions \hat{y}_a for each arc $a \in A$. The reduced network is defined as $G_r = (V, A_r)$ where $A_r = \{a \in A \mid \hat{y}_a = 1\}$ and it is used to generate columns as long as it yields a satisfactory number of columns, i.e., a number higher than a parameter value η_{min} . When the reduced network fails to generate enough columns, the complete network is used instead. If the latter generates a number of columns higher than a parameter value η_{max} , we switch back to the reduced network until we reach an optimal solution. Algorithm 2 details how the ML model is used in practice.

The algorithm starts by extracting the features of the arcs and getting the model predictions (Steps (1)–(4)). The predictions \hat{y}_a are then used to build the reduced graph G_r that is set as the active graph to be used when solving the PP (Steps (5)–(8)). Note that the active graph G_{active} can either be the original complete graph G or the reduced one G_r . The boolean variable *useReducedG* is also initialized to **True**, indicating that the reduced graph is being used (e.g., equivalent to $G_{active} == G_r$). In Step (9), an initial solution to the RMP is obtained along with dual values π . At each CG iteration, the currently active network is priced and a new set of columns \mathcal{C}_i is generated (Steps (11) and (12)). If the columns were generated by the reduced network and the number of columns $|\mathcal{C}_i| < \eta_{min}$, we switch to using the full network by setting the variables G_{active} to G and *useReducedG* to **False**

Algorithm 2 Machine-learning-based pricing heuristic in CG

```

1: for each  $a \in A$  do
2:    $x_a \leftarrow \text{extractFeatures}(a)$ 
3:    $\hat{y}_a \leftarrow \text{predict}(x_a)$ 
4: end for
5:  $A_r := \{a \in A \mid \hat{y}_a = 1\}$ 
6:  $G_r := (V, A_r)$ 
7:  $G_{active} \leftarrow G_r$ 
8:  $\text{useReducedG} \leftarrow \text{True}$ 
9:  $\pi \leftarrow \text{initialRMPSolution}()$ 
10: for each CG iteration  $i$  do
11:    $\text{priceNetwork}(G_{active}, \pi)$ 
12:    $\mathcal{C}_i \leftarrow \text{generateColumns}(G_{active})$ 
13:   if  $(|\mathcal{C}_i| < \eta_{min} \wedge \text{useReducedG})$  then
14:      $\text{useReducedG} \leftarrow \text{False}$ 
15:      $G_{active} \leftarrow G$ 
16:   else if  $|\mathcal{C}_i| \geq \eta_{max} \wedge (\neg \text{useReducedG})$  then
17:      $\text{useReducedG} \leftarrow \text{True}$ 
18:      $G_{active} \leftarrow G_r$ 
19:   else if  $\mathcal{C}_i == \emptyset \wedge (\neg \text{useReducedG})$  then
20:     Exit
21:   end if
22:   if  $|\mathcal{C}_i| > 0$  then
23:      $\text{addColumns}(\mathcal{C}_i)$ 
24:      $\pi \leftarrow \text{reoptimizeRMP}()$ 
25:   end if
26: end for

```

} Reduced network G_r
 } Network switch cases
 } Solution is optimal

(Steps (13)–(15)). On the contrary, if the full network was used and $|\mathcal{C}_i| \geq \eta_{max}$, we switch back to using the reduced network G_r and set *usedReducedG* to **True** (Steps (16)–(18)). By using the two previous conditions, the algorithm switches between the two graphs until an optimal solution is obtained. Normally, with the right choice of the two parameters η_{min} and η_{max} , at the last iterations of the CG, the reduced network will not generate enough columns (i.e., less than η_{min}) and so the full network G will be used instead. Most likely, the full network will neither generate many columns (i.e., less than η_{max}), so the algorithm will not switch back to the reduced network and will keep using the full network until it generates no more columns, indicating that the obtained solution is optimal (Steps (19)–(21)). If columns were generated, they are added to the RMP which is reoptimized afterwards (Steps (22)–(25)).

4 Application I: Vehicle and crew scheduling problem

In a public transit system, the planning process goes through several stages. The first step in the process is to determine the bus lines, the stops on each line, and the frequency of the trips. Based on these pieces of information, a timetable is created, describing the trips with their corresponding starting and ending times and locations. The next two steps in the process are the construction of vehicle routes and crew schedules, which means solving, respectively, the two scheduling problems: the *Vehicle scheduling problem* and the *Crew scheduling problem*. Traditionally, these two problems are solved in a sequential manner, where the vehicle routes are determined first before the crew schedules. However, this approach is not guaranteed to provide the best solution, because in most cases driver costs dominate the costs of vehicle use. A first formulation that integrates and considers the two problems simultaneously was proposed by [13], giving rise to the VCSP. For solving the problem, the authors described a column generation approach applied to a Lagrangian relaxation of the master problem, since the problem contains a very large number of variables representing the feasible duties. Most of the formulations proposed in the literature are set partitioning ones, where column generation plays an important role to deal with the large number of variables and to generate schedules as needed. Given the difficulty of the problem, several heuristics have been proposed in the literature as well. In

this paper, we consider the formulation described in [17], where the RMP is the linear relaxation of a set partitioning problem with side constraints and the PP is a SPPRC. The RMP formulation contains only the crew schedule variables, but side constraints are added to the problem to ensure that the vehicle schedules can be obtained afterwards in a polynomial time. In addition, the costs of the vehicles are included in the objective function to ensure that an overall optimal solution is obtained.

4.1 Network structure

Since the SPPRC is the part of interest for our work, we will take a closer look at the network structure of the problem before diving into the details about the data collection, the instances used and the results obtained.

In this problem, a network is used to generate driver schedules and there is one network for each schedule type (e.g. regular schedule with one meal break or straight without a meal break). Additionally, the arcs can be divided into two categories: 1) bus movement when the driver is driving or attending a bus and 2) walking when the driver is moving on his/her own for repositioning purposes. As mentioned before, each bus line is defined by three different trip nodes: departure node, arrival node and intermediate nodes where an exchange of drivers can be performed. For each bus line, several trips are planned during the day at different times, the exact time at which a bus should arrive at each of the trip nodes is therefore known. To cover the trips, the buses leave the depot in the direction of the trip departure nodes. Each bus can be used for multiple trips, and can be operated by different drivers before returning to the depot at the end of the day. Since the same bus must service the whole trip, a bus moving to the departure node of a trip must be retained until it arrives at the destination node, at which point it can head to the start of the next trip. Therefore, no bus movements leaving the departure or intermediate nodes to a different trip are allowed. In fact, the bus movements represent less than 7% of the total number of arcs, while the majority of the remaining 93% are the walking movements of the drivers. Unlike the bus movements, the walking movements are not as restricted because they can occur at the beginning, the end or the middle of the trips, as long as they are feasible. Figure 3 is a simplified and non-detailed representation of the main components found in a VCSP. In addition to the main components, the network can also include break nodes and their corresponding arcs to model the breaks that drivers can take after a certain working time. Some details are not involved in our selection strategy and therefore not detailed here, but the interested readers can consult [17] for an in-depth overview.

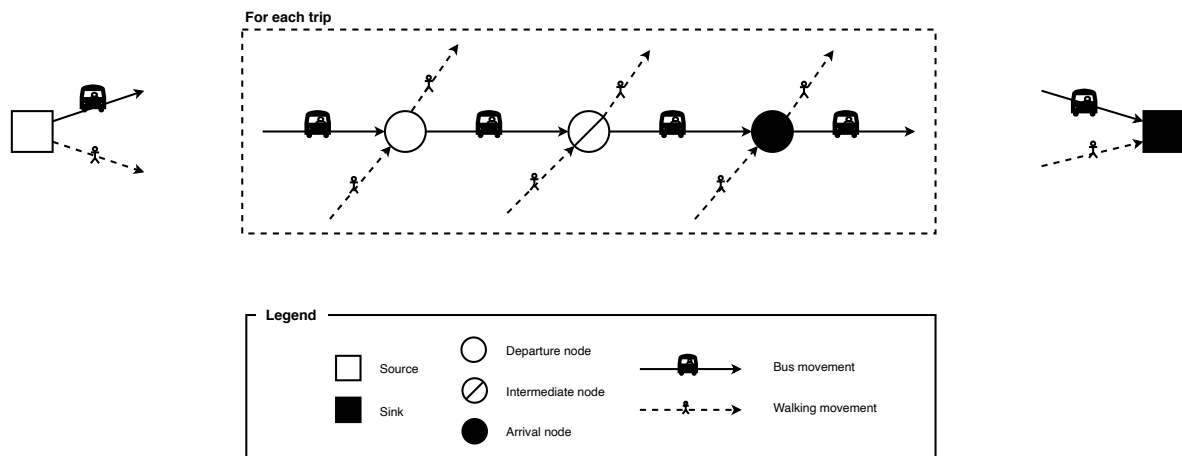


Figure 3: A simplified version of the network components in the VCSP.

For us, the movements we are interested in are the walking movements of the drivers between trip nodes, since they represent more than 90% of the total number of arcs. The selection strategy only covers this subset of arcs, denoted as $A_s \subset A$ and not all arcs in A . For a walking movement arc

linking the trip nodes of two different trips $T1$ and $T2$, the classification model will try to predict whether there is a high chance that a driver leaves his/her current trip $T1$ in the direction of trip $T2$. The arcs entering and leaving the depot, as well as the bus movements entering and leaving the trip nodes are not involved, in addition to any extra arcs for breaks, etc.

4.2 Data collection, features and labels

All the features mentioned in Section 3.1 are valid for the VCSP. For the resource consumptions on the arcs, the two representing the working time and waiting time, are respectively extracted. The first resource is the time taken to walk to the next trip node, while the second resource is the waiting time required before actually starting the trip. Note that the waiting time is bounded and cannot exceed the upper bound indicated by the resource windows. Given that the exact time at which the trip nodes must be visited is known, the times associated with node i and j for each arc $(i, j) \in A_s$ are added to the list of features as well. The difference between j and i times is equal to the time required to walk from i to j , plus the waiting time, but since there are more trips during peak hours, the time of the day information can be useful. Another feature that can be added to the list is the arc type since the arcs we are interested in can link three different types of nodes (i.e., departure, intermediate and arrival nodes) yielding nine possible connection types. Therefore, with the additional features, a total of 22 features is collected for each arc.

As described in Section 3.2, labels are assigned to the arcs that are part of the generated columns. However, only the arcs representing the walking movements between the trip nodes are concerned.

4.3 VCSP instances

To collect enough data, an instance generator is used as in [17]. By taking a total number of trips as an input, the generator splits this number over the different bus lines. The trips are then randomly distributed over the time horizon, with more trips during peak hours. A total of 20 instances with 400 trips were used in the training phase, each instance added about 250,000 data points to the dataset. Other new instances of different sizes (i.e., 300, 350, ..., 600) are generated and used to test the integration of the ML model in the CG algorithm. The instances of size other than 400 also serve to evaluate the generalization of the model to unseen instances of different sizes.

4.4 Computational results

This section is divided into three parts. The first one is devoted to the machine learning phase, where additional details about data collection, the algorithms used and the results obtained are given, whereas the second part presents the results of the integration of the model in the CG algorithm. In the third part, we highlight the results of some additional experiments that are worth mentioning. All the experiments were performed on a Linux machine with an i7-8700 CPU @ 3.20GHz and 64GB of RAM.

4.4.1 Machine learning

With the dataset in hand, some preprocessing is performed such as data normalization and one-hot encoding of categorical features (i.e., arc type). The next step is to choose the appropriate classification algorithm to use. There are several possibilities, from the simplest linear model (i.e., Logistic regression) to the more classic algorithms (SVM, Decision Trees, etc.) and more complex models (Artificial Neural Networks, etc). To compare the models with each other, the main criteria is their accuracy on the test set. However, when there is an imbalance between the output classes (i.e., in our case, only about 15% of the data have label 1 and the remaining have label 0), it is better to consider other metrics such as the true positive rate (TPR, also called the Recall), the true negative rate (TNR),

precision, etc. These metrics provide an idea about the performance obtained for each individual class allowing a better comparison between the different models.

Three different algorithms have been chosen for the training phase, starting with a linear classifier: Logistic regression, then moving to two more advanced algorithms: Random Forest (RF) and Artificial Neural Networks (ANN). Each of these algorithms requires hyperparameters to be tuned, a “cross validation” approach is used to compare the different set of values. The best values obtained are described in Table 1, while Table 2 shows the results of the three algorithms on the test set.

Table 1: Hyperparameters values used in the training phase.

Algorithm	Hyperparameter	Value
Logistic regression	C parameter	1
	Solver	lbfgs
	Class weights	Balanced
Random forest	Bootstrap	True
	Max depth	10
	Max features	7
	Min samples per leaf	50
	Min samples per split	100
	Number of trees	500
	Class weights	Balanced
Neural Network	Learning rate	10^{-3}
	Epochs	1000
	Batch size	32
	Loss function	Binary cross entropy
	Activation function	ReLU
	Architecture	32x32x32x1
	Optimizer	Adam
	Class weights	7:1

Table 2: VCSP metrics of three different algorithms.

Algorithm	Recall	TNR	Balanced accuracy
Logistic regression	64%	81%	72%
Random forest	76%	81%	78%
Neural network	78%	80%	79%

The Recall represents the percentage of arcs with label $y_a = 1$ that are correctly classified, while the TNR represents the accuracy of the opposite class (i.e., classification accuracy of the arcs with label 0). Ideally, a high percentage of both is sought. The last column represents the balanced accuracy which is the average of the two previous columns. To deal with the unbalanced data, weights are used during the optimization of the loss function. The goal is to give a higher penalty to the arcs with label 1 that are misclassified since they are a minority in the dataset, hence the use of the “balanced” setting in the class weights hyperparameter. One can imagine the dataset as a cloud of points, where 85% are red and 15% are green. With logistic regression, one tries to separate these two classes with a linear separator, but since there are more red points than green, there is more chance to have a large number of red points on one side of the separator and thus a high TNR. A high recall can also be obtained if the two classes are easily separable, which does not seem to be the case according to the results obtained (i.e., 64% Recall and 81% TNR). The model slightly underfits but overall, the accuracy obtained is not far from the other two algorithms. On the other hand, the RF and ANN models are able to reach a higher accuracy of 78-79%, with more balance between the Recall and the TNR. The values obtained by the two algorithms are quite close to each other.

4.4.2 CG with arc selection

After the training phase, both the RF and ANN models seem to be good candidates to be integrated in the CG algorithm. Given their almost similar accuracy, one can expect them to give comparable results. When testing the models on several groups of instances with different sizes, the RF model performs better on some instances, but on others it is the ANN model that is better. However, because the RF model is a few percents ahead on average, we retained this candidate for the subsequent tests.

As described in Algorithm 2, the model is used before the beginning of the CG process to build the reduced network G_r . Different scenarios are compared to assess the efficiency of the arc selection, namely:

Baseline-CG: This corresponds to the basic CG with no arc selection, i.e., network G is used in all iterations.

ML-S: Arc selection is used and the CG alternates between the use of the reduced network G_r and the full network G . The reduced network is built from the predictions obtained by the RF model. The parameters used during the CG process are: $\eta_{min} = 30, \eta_{max} = 100$.

Random-S: This scenario corresponds to a completely random selection of the arcs in A_s , the number of selected arcs is the same as with the model. This can be achieved by simply shuffling the values obtained by the model.

Cost-S: For each trip node i , incoming arcs $(n, i) \in A_s$ and outgoing arcs $(i, n) \in A_s$ are sorted in ascending order of their cost, and a subset of arcs with the lowest costs is selected. The total number of selected arcs is approximately the same as with the model.

The results obtained by the four algorithms are reported in Table 3. The names of the instances are written in the form “VCS_Size_Id”, where *Size* is the instance size, i.e., the number of trips, and *Id* is the instance identifier. For each algorithm, we report the number of CG iterations required to reach the linear relaxation solution, the computing time in seconds of: the PP, the RMP and in total. For the algorithms involving an arc selection, an additional column shows the reduction of the computing time gained in comparison with the “Baseline-CG” algorithm. Additionally, the number of iterations performed using the full network G is indicated in parentheses. After each group of instances of the same size, a line indicating the average values is added.

According to the results, if we start by comparing Baseline-CG and ML-S, we can observe a reduction in computing time ranging from 23 to 27%, mostly on the PP side. One can also notice that the highest reductions are obtained when the number of iterations is lower than the one with Baseline-CG. On average, the number of iterations is very close or slightly higher. Therefore, the ML selection allows a reduction of the size of the PP network without increasing considerably the number of iterations. We can also notice that the performance of the model is maintained across all groups of instances. Considering that the model has been trained only on 400-trip instances, this also shows the ability of the model to generalize to different instance sizes.

For the random selection case (i.e., Random-S), from the negative gains obtained we can conclude that this selection is not very promising. This is reflected by the number of iterations performed that is higher than the one obtained using Baseline-CG, and also by the number of times the full network was used if we compare it to ML-S. On the other hand, the results obtained by Cost-S show that a selection based on the costs is more effective. One can consider that Cost-S is a selection based on a single feature, which is the cost, while ML-S uses several features at the same time to decide whether an arc is to be retained or not. The results lie between those of the two algorithms Random-S and ML-S, i.e., it is better than a completely random selection but worse than the one with a learned model. This is reflected by the total number of iterations, the number of times the full network G was used and the reductions obtained.

Table 3: VCSP results.

Instance	Baseline-CG			ML-S			Random-S			Cost-S												
	#Itr	Time (s)		#Itr	Time (s)		Gain	#Itr	Time (s)		Gain	#Itr	Time (s)		Gain							
		PP	RMP		Total	PP			RMP	Total			PP	RMP		Total						
VCS_300.1	158	133	80	213	146	(12)	75	76	151	29%	180	(59)	114	91	205	4%	168	(34)	94	82	176	17%
VCS_300.2	154	174	85	259	158	(8)	106	91	197	24%	242	(72)	214	122	336	-30%	183	(37)	143	94	237	8%
VCS_300.3	157	182	81	263	154	(14)	115	76	191	27%	204	(70)	198	100	298	-13%	182	(37)	160	86	246	6%
VCS_300.4	176	222	81	303	185	(12)	166	81	247	18%	227	(105)	220	96	316	-4%	220	(66)	241	84	325	-7%
VCS_300.5	199	204	89	293	191	(27)	133	87	220	25%	238	(60)	184	95	279	5%	192	(44)	166	87	253	14%
Average	169	183	83	266	167	(14)	119	82	201	25%	218	(73)	186	101	287	-13%	189	(43)	161	87	248	8%
VCS_350.1	136	183	126	309	143	(15)	112	129	241	22%	199	(67)	207	151	359	-16%	175	(42)	166	144	310	0%
VCS_350.2	148	158	136	294	135	(9)	85	119	204	31%	184	(67)	146	142	288	2%	192	(52)	153	150	303	-3%
VCS_350.3	182	253	179	432	170	(10)	141	167	308	29%	219	(61)	226	207	433	0%	192	(37)	192	178	370	14%
VCS_350.4	171	276	162	438	188	(7)	173	174	347	21%	302	(97)	362	243	605	-38%	227	(51)	291	173	464	-6%
VCS_350.5	201	360	164	524	234	(12)	272	183	455	13%	273	(101)	392	222	614	-17%	249	(48)	347	179	526	0%
Average	168	246	153	399	174	(10)	157	154	311	23%	215	(78)	267	193	460	-5%	207	(46)	230	165	395	1%
VCS_400.1	177	370	253	623	177	(10)	238	258	496	20%	212	(89)	361	293	654	-5%	216	(38)	342	297	639	-3%
VCS_400.2	176	422	246	668	189	(8)	255	258	513	23%	240	(86)	471	298	769	-15%	214	(58)	378	257	635	5%
VCS_400.3	153	266	242	508	163	(9)	177	233	410	19%	176	(64)	263	238	501	1%	184	(42)	245	254	499	2%
VCS_400.4	147	209	206	415	147	(12)	126	190	316	24%	199	(64)	225	224	449	-8%	173	(30)	173	222	395	5%
VCS_400.5	207	518	256	774	180	(12)	283	230	513	34%	266	(100)	558	311	869	-12%	214	(44)	391	252	643	17%
Average	172	357	241	598	171	(10)	216	234	450	24%	218	(80)	376	273	648	-8%	200	(42)	306	256	562	5%
VCS_450.1	340	1114	640	1754	368	(14)	809	636	1445	18%	426	(138)	1034	713	1747	0%	366	(64)	868	637	1505	14%
VCS_450.2	269	1540	544	2084	241	(11)	797	486	1283	38%	358	(126)	1453	663	2116	-2%	281	(45)	1071	565	1636	21%
VCS_450.3	194	570	396	966	194	(10)	319	391	710	27%	243	(81)	497	463	960	1%	251	(64)	508	435	943	2%
VCS_450.4	209	505	387	892	234	(8)	312	440	752	16%	264	(110)	473	446	919	-3%	255	(52)	439	439	878	2%
VCS_450.5	279	1507	551	2058	289	(21)	976	564	1540	25%	376	(124)	1566	681	2247	-9%	350	(61)	1510	568	2078	-1%
Average	258	1047	504	1551	265	(12)	643	503	1146	25%	333	(115)	1005	593	1598	-3%	300	(57)	879	529	1408	8%
VCS_500.1	277	1618	856	2474	304	(8)	1030	902	1932	22%	387	(132)	1738	1021	2759	-12%	308	(60)	1272	916	2188	12%
VCS_500.2	296	1654	620	2274	270	(13)	846	576	1422	37%	366	(130)	1404	738	2142	6%	342	(76)	1318	691	2009	12%
VCS_500.3	203	770	604	1374	221	(10)	453	611	1064	23%	291	(101)	887	736	1623	-18%	266	(54)	688	713	1401	-2%
VCS_500.4	266	1839	748	2587	281	(8)	1235	749	1984	23%	357	(130)	1842	908	2750	-6%	313	(54)	1526	789	2315	11%
VCS_500.5	322	1720	627	2347	340	(23)	1024	656	1680	28%	442	(173)	1635	769	2404	-2%	386	(103)	1486	701	2187	7%
Average	273	1520	691	2211	283	(12)	918	699	1616	27%	368	(133)	1501	834	2336	-8%	323	(69)	1258	762	2020	8%
VCS_550.1	254	1723	842	2565	275	(9)	1125	913	2038	21%	368	(137)	2101	1084	3185	-24%	309	(61)	1421	899	2320	10%
VCS_550.2	280	1347	894	2268	280	(18)	801	971	1772	22%	393	(145)	1583	1116	2699	-19%	324	(87)	1289	976	2265	0%
VCS_550.3	286	1245	822	2067	269	(21)	659	829	1488	28%	326	(103)	960	943	1903	8%	302	(70)	888	907	1795	13%
VCS_550.4	259	1870	863	2733	262	(13)	980	895	1875	31%	406	(132)	2130	1186	3316	-21%	294	(48)	1207	954	2161	21%
VCS_550.5	199	782	774	1556	206	(14)	472	788	1260	19%	242	(67)	632	905	1537	1%	228	(54)	608	858	1466	6%
Average	256	1399	839	2238	258	(15)	807	879	1687	24%	347	(116)	1481	1047	2528	-11%	291	(64)	1083	919	2001	10%
VCS_600.1	290	2840	1229	4069	298	(32)	1694	1235	2929	28%	442	(171)	3416	1595	5011	-23%	320	(69)	1972	1300	3272	20%
VCS_600.2	246	1286	948	2234	249	(20)	815	972	1787	20%	319	(114)	1327	1136	2463	-10%	331	(81)	1324	1185	2509	-12%
VCS_600.3	279	2120	1398	3518	290	(11)	1236	1344	2580	27%	392	(114)	2208	1693	3901	-11%	373	(57)	1937	1537	3474	1%
VCS_600.4	258	2103	1327	3430	272	(15)	1286	1322	2608	24%	379	(104)	2483	1618	4101	-20%	321	(48)	1767	1475	3242	5%
VCS_600.5	247	1722	1186	2908	245	(16)	905	1190	2095	28%	313	(96)	1554	1395	2949	-1%	290	(54)	1510	1315	2825	3%
Average	264	2014	1218	3232	270	(18)	1187	1213	2400	25%	369	(119)	2198	1487	3685	-13%	327	(61)	1702	1362	3064	3%

4.4.3 Additional experiments

This section is dedicated to some additional experiments that have been conducted. Even though they did not give significantly better results, we think that they deserve to be highlighted here.

- The first idea has been previously mentioned, which consists of selecting only the arcs that are part of the columns in the optimal basis of the solved RMPs. In the training phase, this strategy gave slightly worse results, i.e., 75% accuracy, and when integrated in the CG algorithm, the extra gain obtained by reducing further the network size is nullified by doing more iterations.
- Another idea is to assign weights to the arcs according to the number of times they have been used in the generated columns. An accuracy of 78% is obtained in the training phase, distributed as 86% recall and 70% TNR, but again, the results were not any better when integrated in CG.
- We made a small modification to the initial algorithm which consists of skipping the label assignment at the first n iterations, where n is a given parameter. The idea is that some arcs can be interesting at the beginning of the optimization but never used at later iterations. Like the other changes, the results obtained were not really compelling.
- In the same context of combining ML and CG, in our previous paper [24], we were interested in reducing the computing time of the RMPs by selecting the most promising columns at each iteration. Reductions in computation time of up to 30% were achievable. However, the developed approach was more intended for problems that take the majority of the time in solving the RMPs and not the PPs, e.g., the tests were performed on instances that take on average 75% or more of the computing time in solving the RMPs. Nevertheless, it was interesting to check if column selection can yield an additional gain when used with ML-S. The results comparing ML-S and ML-S with column selection, i.e., ML-COL-S, are reported in Table 4.

From these results, we observe a small additional reduction in the average computing time per instance group ranging from 5% to 14%.

5 Application II: Vehicle routing problem with time windows

The vehicle routing problem (VRP) is a well-known and classical problem that has been extensively studied in the literature. Given a fleet of vehicles, the VRP consists of constructing routes to serve geographically dispersed customers, while minimizing the travel costs and respecting the capacity of the vehicles. Each route must start at the depot, visit a set of clients and then return to the depot at the end of the tour. The variant with time windows (VRPTW) adds an additional complexity by introducing a restriction on the times during which the clients can be visited. Each client must be served during the associated time window by exactly one vehicle, and in case the vehicle arrives earlier, it has the possibility to wait until the client is available.

The most efficient and recent methods to solve this problem are based on CG, embedded in a branch-and-bound framework, with the addition of cuts to improve the bounds quality, giving rise to what is called “*branch-price-and-cut*” algorithms. CG is used to solve the linear relaxation at each node of the tree, where the master problem is most likely a set partitioning problem and the PP is an ESPPRC solved using dynamic programming algorithms. The elementarity requirement of the PP makes the problem much more difficult to solve than the standard SPPRC, especially for large instances where several clients can be visited by one route. To overcome these difficulties, researchers have investigated several relaxations (see [2, 21] and [9]), seeking a good trade-off between the quality of the bound and the difficulty to solve the problem. Other studies have explored heuristic solutions [9, 14], such as dominating on a subset of resources, limiting the number of labels kept at each node, or reducing the number of arcs in the network based on their reduced cost. Note that many of these techniques can be combined and used simultaneously, allowing a large reduction in computing time. For more details on the exact methods for solving the VRP, the survey of [6] covers a wide range of methods

Table 4: VCSP results with column selection.

Instance	ML-S				ML-COL-S				Gain
	#Itr	Time (s)			#Itr	Time (s)			
		PP	RMP	Total		PP	RMP	Total	
VCS_300.1	146 (12)	75	76	151	123 (10)	64	70	134	11%
VCS_300.2	158 (8)	106	91	197	140 (7)	103	78	181	8%
VCS_300.3	154 (14)	115	76	191	135 (10)	101	66	167	13%
VCS_300.4	185 (12)	166	81	247	161 (6)	128	67	195	21%
VCS_300.5	191 (27)	133	87	220	142 (9)	102	76	178	19%
Average	167 (14)	119	82	201	140 (8)	99	71	171	14%
VCS_350.1	143 (15)	112	129	241	126 (7)	106	125	231	4%
VCS_350.2	135 (9)	85	119	204	124 (9)	81	105	186	9%
VCS_350.3	170 (10)	141	167	308	163 (7)	139	156	295	4%
VCS_350.4	188 (7)	173	174	347	179 (9)	198	148	346	0%
VCS_350.5	234 (12)	272	183	455	187 (9)	234	155	389	15%
Average	174 (10)	157	154	311	155 (8)	151	137	289	6%
VCS_400.1	177 (10)	238	258	496	161 (7)	229	232	461	7%
VCS_400.2	189 (8)	255	258	513	175 (11)	255	228	483	6%
VCS_400.3	163 (9)	177	233	410	148 (7)	176	203	379	8%
VCS_400.4	147 (12)	126	190	316	122 (6)	108	166	274	13%
VCS_400.5	180 (12)	283	230	513	200 (7)	316	236	552	-8%
Average	171 (10)	216	234	450	161 (7)	216	213	429	5%
VCS_450.1	368 (14)	809	636	1445	319 (10)	723	513	1236	14%
VCS_450.2	241 (11)	797	486	1283	232 (6)	833	430	1263	2%
VCS_450.3	194 (10)	319	391	710	173 (9)	314	355	669	6%
VCS_450.4	234 (8)	312	440	752	190 (9)	280	349	629	16%
VCS_450.5	289 (21)	976	564	1540	263 (8)	918	489	1407	9%
Average	265 (12)	643	503	1146	235 (8)	613	427	1040	9%
VCS_500.1	304 (8)	1030	902	1932	292 (7)	986	888	1874	3%
VCS_500.2	270 (13)	846	576	1422	259 (8)	852	570	1422	0%
VCS_500.3	221 (10)	453	611	1064	177 (7)	416	550	966	9%
VCS_500.4	281 (8)	1235	749	1984	247 (6)	1134	670	1804	9%
VCS_500.5	340 (23)	1024	656	1680	299 (11)	916	604	1520	10%
Average	283 (12)	918	699	1616	254 (7)	860	656	1517	6%

and explores different variants of the problem. Similar to the recent formulations proposed in the literature, the model used in this work is a set partitioning problem and the PP is a SPPRC with 2-cycle elimination [11] solved using a labeling algorithm.

In this section, we will first describe the network structure of the problem, before giving more details about the data collection and the instances used. Finally, we will discuss the results obtained.

5.1 Network structure

We consider the basic network structure of the VRPTW with a single depot, where only one type of arcs exists, representing the possible vehicle movements between the depot and the clients and between the clients. Therefore, all the arcs will be targeted by the selection strategy, with the exception of the ones entering and leaving the depot. For other problem variants and depending on the network structure, it would be possible to limit the selection to only a subset of arcs as we did for the VCSP.

5.2 Data collection, features and labels

In this section, a few additional details specific to the VRPTW are given regarding the features described in Section 3.1. Starting with the resource consumptions on the arcs, two resources are

considered as features: time and capacity (i.e., the customer demand). For each arc $(i, j) \in A, i \neq s, j \neq t$, the lower and upper bounds of the time windows of both client nodes i and j are collected. For the capacity resource, since the windows are the same for all clients (i.e., $[0, Q]$ where Q is the vehicle capacity), their bounds are not included in the features.

Finally, the labels are assigned as described in Section 3.2, i.e., the label 1 is assigned to the arcs that are part of the generated routes, and 0 for the others.

5.3 VRPTW instances

The VRPTW instances used are based on the Gehring & Homberger instances [18]. Three classes of instances are available: the R instances, standing for “Random”, where the clients positions are randomly scattered, the C instances, standing for “Cluster”, where clients are grouped together in clusters and thus have more chance to be served by the same vehicle, and finally the RC instances that are a combination of both. Moreover, for each of the three classes, there are two categories numbered 1 and 2 ($R1, C1, RC1$ and $R2, C2, RC2$). The difference between the two is in the width of the time windows: the category 1 instances tend to have larger time windows than those of category 2, and are therefore more difficult to solve.

Since we are only interested in solving the linear relaxation of the problem, and that most of the $R1, C1$ and $RC1$ instances we considered, i.e., 200 and 400 clients, are actually very easy to solve and take only a few seconds, we decided to only use the $R2, C2$ and $RC2$ instances. Some of these instances have been slightly modified, more precisely, the width of the time windows has been slightly reduced so that the computing time becomes reasonable and not very high. This is because some instances take hours to be solved, especially when no heuristics are used during the data collection phase.

Regarding the size of the instances, we used those of 200 and 400 clients. Half of the 200-client instances were used for training and testing in the ML phase (e.g., from the three classes), while the other half and the 400-client instances were used for testing the model obtained when integrated in the CG algorithm.

5.4 Computational results

As for the VCSP, this section will be divided into three parts. The first part is devoted to the ML phase where additional details about the features and data collection are covered as well as the results obtained. The second part presents the results of the model integration in the CG process, as well as a comparison with other algorithms. In the third part, we include some additional experiments that are worth mentioning. All the experiments were performed on a Linux machine with an i7-8700 CPU @ 3.20GHz and 64GB of RAM.

5.4.1 Machine learning

Because of the noticeable differences in the feature value range from one instance to another, especially between instances of different sizes, a preprocessing step is performed to scale and normalize the data of each instance individually. A random forest model is fit on the training data, and the hyperparameters are tuned through a cross validation approach. A neural network model was also tested, but without obtaining a significantly better accuracy, so we decided to retain the RF model as we did for the VCSP. The best hyperparameter values obtained are described in Table 5, whereas the results on the test set are reported in Table 6.

The hyperparameter values are slightly different from those used for the VCSP, more precisely the `max_depth` and `max_features` parameters. The class weights are used to balance the two classes as before, since there is only 14% and 9% of arcs with label 1 in the 200 and 400 instances, respectively. According to the results in Table 6, we notice a higher accuracy than the one obtained for the VCSP.

Table 5: Hyperparameters values used in the training phase.

Hyperparameter	Value
Bootstrap	True
Max depth	5
Max features	5
Min samples per leaf	50
Min samples per split	100
Number of trees	500
Class weights	Balanced

Table 6: Metrics of the model obtained for the VRPTW.

Metric	Value
Recall	93%
TNR	87%
Balanced Accuracy	90%

The arcs to be selected are correctly predicted with 93% accuracy (i.e., Recall), whereas the TNR value is 87%, which means that the number of arcs selected by the model will not differ much from the number selected by the expert.

Note that the VCSP and the VRPTW are two completely different problems and they are not supposed to give comparable accuracies, especially considering that the arcs to be selected play different roles in each problem.

5.4.2 CG with arc selection

Now let us evaluate the ML model performance when integrated in the CG process. To do so, we have chosen to compare different algorithms described as follows:

Baseline-CG: This corresponds to the standard CG algorithm without arc selection where the complete network is used in all iterations.

ML-S: This is the CG algorithm with the arc selection activated as described in Algorithm 2, but with a small adjustment. In fact, for the VRPTW we noticed that towards the end of the optimization, the reduced network fails to generate columns for several consecutive iterations. Thus, the algorithm ends up solving the PP with both networks at multiple iterations (i.e., with G_r followed by G), which slows down the optimization. Instead, we decided to disable the use of the reduced network as soon as it fails to generate a column until the end of the optimization.

RedCost-S: This corresponds to the CG algorithm using a known heuristic pricing strategy that reduces the number of arcs in the network. It has been used by several researchers, and has shown to be very effective. For each node, excluding the depot nodes s and t , the heuristic consists in setting a minimum number N^{min} of incoming and outgoing arcs to keep. At each iteration, the arcs are sorted by their reduced cost, the N^{min} incoming and outgoing arcs with the least cost are kept, while the others are removed from the network. It is also possible to define multiple values $N_1^{min}, N_2^{min}, \dots, \infty$ where $N_1^{min} < N_2^{min} < \dots < \infty$. The PP tries first to generate columns with the parameter value N_1^{min} , if it fails it moves to the next parameter value, and so on. The values used in this experiment are $N_1^{min} = 10$, $N_2^{min} = 20$, $N_3^{min} = \infty$. Note that this algorithm performs the selection at each iteration while ML-S performs the selection only once before the beginning of the optimization.

The results obtained to compare the three algorithms are described in Table 7. For each algorithm, we report the same information as in Table 3. An additional column presents the time gain in percentage obtained by ML-S and RedCost-S in comparison to Baseline-CG. For ML-S, the number of iterations with the full network is indicated between parentheses.

Table 7: VRPTW results.

Size	Instance	Baseline-CG				ML-S				Gain	RedCost-S				Gain
		#Itr	Time (s)			#Itr	Time (s)				#Itr	Time (s)			
			PP	RMP	Total		PP	RMP	Total			PP	RMP	Total	
200 clients	R2.2.1	197	173	8	181	259 (165)	52	8	54	68%	439	23	12	34	81%
	R2.2.2	159	312	5	317	211 (150)	99	6	105	67%	179	25	5	29	91%
	R2.2.3	198	128	7	135	301 (233)	36	8	44	68%	441	20	8	28	79%
	R2.2.4	114	210	2	212	152 (101)	69	3	72	66%	152	34	2	37	83%
	R2.2.5	239	238	10	248	275 (166)	44	9	53	79%	367	22	11	33	87%
	Average	181	216	6	233	240 (163)	59	7	65	70%	316	25	7	32	84%
	RC2.2.1	204	203	5	208	243 (172)	123	5	128	38%	374	32	6	38	82%
	RC2.2.2	119	138	3	141	156 (107)	49	3	52	63%	218	17	3	20	86%
	RC2.2.3	243	211	7	218	314 (214)	82	9	91	59%	852	40	11	51	77%
	RC2.2.4	198	189	7	196	239 (166)	150	7	157	20%	609	34	13	47	76%
	RC2.2.5	230	295	6	301	278 (192)	112	6	118	61%	313	32	6	39	87%
	Average	199	207	6	213	246 (170)	103	6	109	48%	473	31	8	39	81%
	C2.2.1	288	177	8	185	398 (268)	143	10	153	18%	760	72	11	83	55%
	C2.2.2	149	353	5	358	195 (126)	255	5	260	27%	316	121	6	126	65%
	C2.2.3	211	190	6	196	359 (176)	146	11	157	20%	272	54	7	61	69%
C2.2.4	164	267	6	273	351 (163)	230	9	239	12%	181	65	5	70	74%	
C2.2.5	158	469	6	475	420 (161)	440	11	451	5%	237	140	6	146	69%	
Average	194	291	6	297	345 (178)	243	9	252	16%	353	90	7	97	66%	
400 clients	R2.4.1	397	643	143	786	447 (239)	154	104	258	67%	373	66	119	185	76%
	R2.4.2	248	318	85	403	330 (165)	90	76	166	59%	289	45	84	129	68%
	R2.4.3	202	1058	42	1100	252 (188)	333	47	380	65%	235	97	37	134	88%
	R2.4.4	149	734	21	755	194 (112)	203	23	226	70%	239	142	20	162	79%
	R2.4.5	355	316	131	447	439 (244)	124	107	231	48%	361	51	127	178	60%
	R2.4.6	240	293	76	369	331 (162)	75	67	142	62%	277	44	78	122	67%
	R2.4.7	245	629	47	676	303 (179)	227	50	277	59%	218	81	44	125	82%
	R2.4.8	187	632	25	657	203 (116)	243	27	270	59%	242	144	23	167	75%
	R2.4.9	334	870	138	1008	472 (253)	277	111	388	62%	322	101	116	217	78%
	R2.4.10	308	988	113	1101	461 (257)	521	102	623	43%	380	154	108	262	76%
	Average	267	648	82	730	343 (191)	255	71	296	59%	294	93	76	168	75%
	RC2.4.1	392	281	123	404	555 (288)	144	116	260	36%	390	45	102	147	64%
	RC2.4.2	301	282	65	347	434 (216)	108	70	178	49%	295	37	67	104	70%
	RC2.4.3	202	1224	29	1253	293 (151)	288	33	321	74%	301	108	27	135	89%
	RC2.4.4	155	361	17	378	231 (125)	185	22	207	45%	183	81	16	97	74%
	RC2.4.5	797	337	198	535	1301 (848)	271	199	470	12%	1164	84	186	270	50%
	RC2.4.6	370	622	98	720	543 (295)	229	101	330	54%	434	98	118	216	70%
	RC2.4.7	340	615	87	702	510 (235)	264	92	356	49%	415	100	98	198	72%
	RC2.4.8	373	629	74	703	478 (254)	245	73	318	55%	435	90	72	162	77%
	RC2.4.9	305	460	83	543	417 (263)	188	91	279	49%	371	83	91	174	68%
RC2.4.10	411	219	101	320	458 (347)	128	89	217	32%	438	40	91	131	59%	
Average	365	503	88	591	522 (302)	205	89	294	46%	443	77	87	163	69%	
C2.4.1	599	373	136	509	1125 (660)	424	178	602	-18%	966	144	183	327	36%	
C2.4.2	270	214	74	288	484 (237)	197	78	275	5%	445	71	88	159	45%	
C2.4.3	254	543	50	593	342 (180)	496	58	554	7%	484	176	55	231	61%	
C2.4.4	153	574	26	600	408 (121)	588	52	640	-7%	332	267	30	297	51%	
C2.4.5	261	632	62	694	714 (222)	637	93	730	-5%	417	194	63	257	63%	
Average	307	467	69	536	614 (284)	468	91	560	-4%	528	170	83	254	51%	

Starting with Baseline-CG, we can notice that most of the computing time is spent on the PP side. If we take a closer look at the time per iteration of one of the instances (R2.2.1 as an example), as shown in Figure 4a, the PP computing time decreases very rapidly, starting at 22 seconds at the first iteration and reaching less than 0.1 second after a hundred iterations. In fact, according to the cumulative time shown in Figure 4b, we can see that 90% of the computing time is spent in the first 50 iterations, which is just a quarter of the total number of iterations. The progressive decrease of computing time can be explained by looking at the number of labels created during the resolution

of the PP as shown in Figure 4c. This shows that more labels are dominated as we progress in the optimization, and with less labels there is less processing and therefore less computing time.

Moving on to the ML-S results, we can notice a higher number of iterations compared to Baseline-CG. The number of iterations with the full network (i.e., in parentheses) is also quite high, representing more than 50% of the total number. Nevertheless, a significant reduction of computing time on the PP side is obtained. Since most of the iterations using the full network are performed in the second half of the optimization (i.e., when the use of the reduced network is disabled), they are not time consuming as previously shown in Figure 4. We can see a significant gain on the *R2* instances. However, this gain decreases on the instances where the clients are more clustered (i.e., *RC2* and *C2*) and also when moving from the 200 instances to the 400 ones. If we compare the increase in the number of iterations between ML-S and Baseline-CG for the 200-client instances, we can notice an increase of 32% for *R2*, 24% for *RC2* and 73% for *C2*. It is true that the *C2* instances seem to perform way more iterations than the other groups, and this may be one reason explaining the drop in performance. However, we see that the *RC2* instances have a lower increase in the number of iterations than *R2* and yet their results are less good.




Figure 4: PP computing time and number of labels with Baseline-CG.

Let us consider that the optimization goes through two stages when using ML-S. The first stage starts from the beginning until the reduced network fails to generate any columns (i.e., the reduced network is thus disabled), and the second stage comes afterwards until the end of the optimization. By taking an *R2* instance and comparing the computing time per iteration during the two stages, Figure 5a shows that there are multiple peaks of computing time during the first stage, which occur when we switch to the full network (i.e., when the reduced network does not generate enough columns). We can also see that the peaks heights decreases rapidly as we progress in the optimization. In the second stage (e.g., starting at the dashed line), the computing time is not too high, and it is also more stable since we are only using the full network. Figures 5b and 5c give some additional information, namely the cumulative time and the number of labels created, respectively. According to the cumulative time, we can see a step line in the first stage where the computing time increases significantly each time we switch to the full network, whereas in the second stage the line is almost linear. Notice that the total time spent on the first stage is larger than the second one. On the other hand, the number of labels follows exactly the same trend depicted in Figure 5a, which makes sense.

If we try to visualize the same data for a *C2* instance in Figure 6 and compare it to the previously seen data for an *R2* instance (i.e., Figure 5), some differences can be noticed. We can observe that the computing time per iteration during the first stage decreases more slowly and is still quite high when we reach the second stage. The same thing can be noticed for the number of labels created. Looking at the cumulative time, we can see that the second stage takes more time than the first and the line is also steeper.

Finally, for the RedCost-S algorithm, we can notice that the results are significantly better than ML-S. A drop in performance is observed when dealing with the more clustered instances, but it is not as aggressive as ML-S. The number of iterations for the 200-client instances is larger or very close to the one of ML-S, but for the 400-client instances it is less on average.




Figure 5: The PP computing time and number of labels with ML-S for instance R2.2.1.




Figure 6: The PP computing time and number of labels with ML-S for instance C2.2.1.

Now the question is whether we can achieve better results than those obtained with RedCost-S. To answer that, we tried two additional strategies:

ML-RedCost-S: This corresponds to the combination of the two strategies ML-S and RedCost-S. In other words, the arc reduction according to the reduced cost is performed at each iteration on the two graphs G and G_r depending on which one is used.

ML-RedCost-S-2: This corresponds to an ML model learned on the data generated using RedCost-S. This means that RedCost-S is considered the new expert. By going back to the first phase, i.e., data collection, the label 1 is assigned to the arcs that are part of the columns generated using RedCost-S, and 0 for the others. The training is performed on the new data in order to obtain a new model, which is then used to build the reduced network.

For both strategies, the same values of $N_1^{min} = 10$, $N_2^{min} = 20$, $N_3^{min} = \infty$ are used. The results comparing the two strategies to RedCost-S are reported in Table 8, providing the same information as in Table 7 except that the time gains are measured with respect to the time obtained by RedCost-S.

From the results, we can see that the gains obtained by the two new strategies are quite comparable. As before, a decrease in performance is noticed when dealing with the $RC2$ and $C2$ instances. On average, ML-RedCost-S-2 seems to be slightly better as it gives a few larger gains for some groups of instances. Overall, the results show that it is possible to get improvements on top of RedCost-S, up to 41% for the 200-client $R2$ instances and 28% for the 400-client ones, especially when the PP is much more time consuming than the RMP. However, it seems to be less effective on $RC2$ and $C2$ instances.

5.4.3 Additional experiments

Another attempt to deal with the inefficiency of the model on the $C2$ instances was conducted. On the ML side, one can think that the $C2$ instances are quite different from the other groups (i.e., RC and R), and that despite the overall accuracy of 90% obtained, the accuracy is perhaps less for the $C2$ instances. It turns out that this is not the case, the metrics reported for the $C2$ instances alone are the following: 94% recall, 79% TNR and 87% accuracy. It is true that the accuracy is 3% below the average, but these are still solid results. Moreover, even with a new model that was only trained on the $C2$ instances, the results obtained when integrated to the CG algorithm were not much better.

Table 8: Additional VRPTW results.

Size	Instance	RedCost-S				ML-RedCost-S					ML-RedCost-S-2				
		#Itr	Time (s)			#Itr	Time (s)			Gain	#Itr	Time (s)			Gain
			PP	RMP	Total		PP	RMP	Total			PP	RMP	Total	
200 clients	R2.2.1	439	23	12	34	578 (403)	13	12	24	31%	420 (235)	8	10	18	48%
	R2.2.2	179	25	5	29	309 (127)	13	7	21	30%	343 (158)	16	8	24	17%
	R2.2.3	441	20	8	28	406 (255)	8	9	17	38%	378 (219)	8	8	16	44%
	R2.2.4	152	34	2	37	217 (98)	8	3	11	69%	245 (113)	9	4	12	66%
	R2.2.5	367	22	11	33	430 (244)	11	12	23	32%	444 (245)	11	12	23	32%
	Average	316	25	7	32	388 (225)	11	8	19	40%	366 (194)	10	8	18	41%
	RC2.2.1	372	32	6	38	586 (378)	18	10	28	28%	568 (376)	17	9	26	32%
	RC2.2.2	218	17	3	20	332 (179)	12	5	17	15%	282 (142)	11	5	15	25%
	RC2.2.3	852	40	11	51	912 (695)	28	11	39	23%	670 (446)	22	10	32	38%
	RC2.2.4	609	34	13	47	675 (476)	18	14	33	30%	531 (352)	17	13	29	37%
	RC2.2.5	313	32	6	39	451 (226)	15	7	23	41%	497 (270)	20	9	29	25%
	Average	473	31	8	39	591 (391)	18	10	28	28%	510 (317)	17	9	26	31%
	C2.2.1	760	72	11	83	924 (552)	54	15	69	17%	1020 (676)	62	15	77	6%
	C2.2.2	316	121	6	126	416 (242)	76	8	84	34%	385 (210)	71	8	79	38%
	C2.2.3	272	54	7	61	505 (208)	49	11	60	2%	489 (190)	51	11	62	-1%
	C2.2.4	181	65	5	70	442 (113)	61	11	72	-3%	464 (122)	65	12	77	-10%
	C2.2.5	237	140	6	146	501 (159)	127	12	139	5%	479 (174)	121	11	132	10%
	Average	353	90	7	97	558 (255)	73	11	84	11%	567 (274)	74	11	85	10%
400 clients	R2.4.1	373	66	119	185	647 (180)	39	119	158	15%	678 (230)	41	122	163	12%
	R2.4.2	289	45	84	129	436 (127)	24	81	105	19%	429 (117)	24	81	105	19%
	R2.4.3	235	97	37	134	392 (111)	40	54	94	30%	440 (157)	44	55	99	0%
	R2.4.4	239	142	20	162	297 (106)	40	29	69	57%	314 (134)	48	29	77	52%
	R2.4.5	361	51	127	178	619 (165)	32	114	146	18%	573 (159)	30	114	144	19%
	R2.4.6	277	44	78	122	523 (196)	32	94	126	-3%	460 (139)	25	81	106	13%
	R2.4.7	218	81	44	125	390 (114)	38	60	98	22%	384 (116)	31	53	84	33%
	R2.4.8	242	144	23	167	328 (131)	60	33	93	44%	309 (119)	45	31	76	54%
	R2.4.9	322	101	116	217	615 (168)	56	131	187	14%	632 (188)	50	121	171	21%
	R2.4.10	380	154	108	262	691 (230)	76	119	195	26%	719 (201)	64	116	180	31%
	Average	294	93	76	168	494 (153)	44	83	127	24%	494 (156)	40	80	121	28%
	RC2.4.1	390	45	102	147	762 (232)	42	124	166	-13%	801 (310)	50	132	182	-24%
	RC2.4.2	295	37	67	104	630 (210)	37	76	113	-9%	651 (309)	43	89	132	-27%
	RC2.4.3	301	108	27	135	446 (188)	52	35	87	36%	486 (206)	53	36	89	34%
	RC2.4.4	183	81	16	97	263 (76)	29	23	52	46%	278 (86)	31	23	54	44%
	RC2.4.5	1164	84	186	270	1600 (882)	83	196	279	-3%	1810 (1084)	101	251	352	-30%
	RC2.4.6	434	98	118	216	773 (244)	62	130	192	11%	867 (314)	62	122	184	15%
	RC2.4.7	415	100	98	198	792 (348)	79	118	197	1%	674 (274)	60	119	179	10%
	RC2.4.8	435	90	72	162	834 (357)	78	84	162	0%	669 (240)	56	74	130	20%
	RC2.4.9	371	83	91	174	723 (304)	67	116	183	-5%	952 (561)	111	108	219	-26%
	RC2.4.10	438	40	91	131	679 (274)	39	118	157	-20%	689 (266)	35	91	126	4%
	Average	443	77	87	163	750 (312)	57	102	159	4%	788 (365)	60	105	165	2%
	C2.4.1	966	144	183	327	1692 (669)	162	190	352	-8%	1785 (768)	170	186	356	-9%
	C2.4.2	445	71	88	159	867 (438)	99	117	216	-36%	726 (239)	68	110	178	-12%
	C2.4.3	484	176	55	231	741 (333)	215	78	293	-27%	913 (463)	210	87	297	-29%
	C2.4.4	332	267	30	297	522 (249)	327	44	371	-25%	536 (287)	290	47	337	-13%
	C2.4.5	417	194	63	257	887 (266)	248	101	349	-36%	895 (278)	222	101	323	-26%
	Average	529	170	84	254	942 (391)	210	106	316	-26%	971 (407)	192	106	298	-17%

6 Conclusion

In this paper, a new pricing heuristic based on ML was presented. The goal is to speed up the CG method for problems where the PP is a SPPRC or one of its variants defined on a network, and solved using a labeling algorithm. The method consists in reducing the network size, by keeping only the most promising arcs that have a high chance to be part of good columns. The ML model is trained on the data collected from previous executions through a supervised learning approach. The arcs selected by the trained model are used to build a new reduced network that is on average 15% to 25% the size of the original network. The new network is used at each iteration as long as it generates a satisfactory number of columns. If not, the full network is used instead, especially in the last iterations.

The approach is fairly general and can be used for different problems. We chose to demonstrate it on two well-known problems, the VCSP and VRPTW. For the VCSP, the selection was limited to the arcs representing walking movements, since they represent more than 95% of the arcs in the network. Whereas for the VRPTW, all the arcs were targeted by the selection. The results showed reductions

in computing time of up to 27% for the VCSP, and 40% for the VRPTW. The resulting ML models have also shown the ability to generalize to new instances not encountered during the training phase. However, the ML model had some limitations when dealing with the C instances of the VRPTW.

References

- [1] Alejandro Alvarez, Quentin Louveaux, and Louis Wehenkel. A machine learning-based approximation of strong branching. *INFORMS Journal on Computing*, 29:185–195, 01 2017.
- [2] Roberto Baldacci, Aristide Mingozzi, and Roberto Roberti. New route relaxation and pricing strategies for the vehicle routing problem. *Operations Research*, 59(5):1269–1283, 2011.
- [3] Cynthia Barnhart, Ellis L. Johnson, George L. Nemhauser, Martin W. P. Savelsbergh, and Pamela H. Vance. Branch-and-price: Column generation for solving huge integer programs. *Operations Research*, 46:316–329, 1996.
- [4] Yoshua Bengio, Andrea Lodi, and Antoine Prouvost. Machine learning for combinatorial optimization: a methodological tour d’horizon. *European Journal of Operational Research*, 2020.
- [5] Claudio Contardo, Guy Desaulniers, and François Lessard. Reaching the elementary lower bound in the vehicle routing problem with time windows. *Networks*, 65(1):88–99, 2015.
- [6] Luciano Costa, Claudio Contardo, and Guy Desaulniers. Exact branch-price-and-cut algorithms for vehicle routing. *Transportation Science*, 53(4):946–985, 2019.
- [7] Guy Desaulniers, Jacques Desrosiers, Irina Ioachim, Marius M. Solomon, François Soumis, and Daniel Vileneuve. A Unified Framework for Deterministic Time Constrained Vehicle Routing and Crew Scheduling Problems, pages 57–93. Springer US, Boston, MA, 1998.
- [8] Guy Desaulniers, Jacques Desrosiers, and Marius M. Solomon. Column generation. Springer, New York, January 2005.
- [9] Guy Desaulniers, François Lessard, and Ahmed Hadjar. Tabu search, partial elementarity, and generalized k-path inequalities for the vehicle routing problem with time windows. *Transportation Science*, 42(3):387–404, 2008.
- [10] M. Desrochers and F. Soumis. A generalized permanent labeling algorithm for the shortest path problem with time windows. *Information Systems and Operations Research*, 26(3):191–212, 1988.
- [11] Martin Desrochers, Jacques Desrosiers, and Marius Solomon. A new optimization algorithm for the vehicle routing problem with time windows. *Operations Research*, 40(2):342–354, 1992.
- [12] Dominique Feillet, Pierre Dejax, Michel Gendreau, and Cyrille Gueguen. An exact algorithm for the elementary shortest path problem with resource constraints: Application to some vehicle routing problems. *Networks*, 44(3):216–229, 2004.
- [13] Richard Freling, Albert P. M. Wagelmans, and José M. Pinto Paixão. An overview of models and techniques for integrating vehicle and crew scheduling. In Nigel H. M. Wilson, editor, *Computer-Aided Transit Scheduling*, pages 441–460, Berlin, Heidelberg, 1999. Springer Berlin Heidelberg.
- [14] Ricardo Fukasawa, Jens Lysgaard, Marcus Poggi de Aragão, Marcelo Reis, Eduardo Uchoa, and Renato F. Werneck. Robust branch-and-cut-and-price for the capacitated vehicle routing problem. *Mathematical Programming*, 106:491–511, 2006.
- [15] Michel Gamache, François Soumis, Gérald Marquis, and Jacques Desrosiers. A column generation approach for large-scale aircrew rostering problems. *Operations Research*, 47(2):247–263, 1999.
- [16] Maxime Gasse, Didier Chételat, Nicola Ferroni, Laurent Charlin, and Andrea Lodi. Exact combinatorial optimization with graph convolutional neural networks. In *Advances in Neural Information Processing Systems* 32, 2019.
- [17] Knut Haase, Guy Desaulniers, and Jacques Desrosiers. Simultaneous vehicle and crew scheduling in urban mass transit systems. *Transportation Science*, 35(3):286–303, 2001.
- [18] Jörg Homberger and Hermann Gehring. A two-phase hybrid metaheuristic for the vehicle routing problem with time windows. *European Journal of Operational Research*, 162(1):220–238, 2005.
- [19] Stefan Irnich. Resource extension functions: properties, inversion, and generalization to segments. *OR Spectrum*, 30(1):113–148, Jan 2008.
- [20] Stefan Irnich and Guy Desaulniers. Shortest path problems with resource constraints. In Guy Desaulniers, Jacques Desrosiers, and Marius M. Solomon, editors, *Column Generation*, pages 33–65. Springer US, Boston, MA, 2005.

- [21] Stefan Irnich and Daniel Villeneuve. The shortest-path problem with resource constraints and k-cycle elimination for $k \geq 3$. *INFORMS Journal on Computing*, 18(3):391–406, 2006.
- [22] Elias B. Khalil, Pierre Le Bodic, Le Song, George Nemhauser, and Bistra Dilkina. Learning to branch in mixed integer programming. In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence*, AAAI’16, pages 724–731. AAAI Press, 2016.
- [23] Andrea Lodi and Giulia Zarpellon. On learning and branching: a survey. *TOP*, 25(2):207–236, Jul 2017.
- [24] Mouad Morabit, Guy Desaulniers, and Andrea Lodi. Machine-learning-based column selection for column generation. *Transportation Science*, 55(4):815–831, 2021.
- [25] Giovanni Righini and Matteo Salani. Symmetry helps: Bounded bi-directional dynamic programming for the elementary shortest path problem with resource constraints. *Discrete Optimization*, 3(3):255–273, 2006. *Graphs and Combinatorial Optimization*.
- [26] Roman Václavík, Antonín Novák, Přemysl Šůcha, and Zdeněk Hanzálek. Accelerating the branch-and-price algorithm using machine learning. *European Journal of Operational Research*, 271(3):1055–1069, 2018.