

# A Neural Column Generation Approach to the Vehicle Routing Problem with Two-Dimensional Loading and Last-In-First-Out Constraints

Yifan Xia<sup>1</sup>, Xiangyi Zhang<sup>2\*</sup>

<sup>1</sup>State Key Laboratory for Novel Software Technology, Nanjing University

<sup>2</sup>1QB Information Technologies, Inc.

yfxia@smail.nju.edu.cn, xyzhangarccos0@gmail.com

## Abstract

The vehicle routing problem with two-dimensional loading constraints (2L-CVRP) and the last-in-first-out (LIFO) rule presents significant practical and algorithmic challenges. While numerous heuristic approaches have been proposed to address its complexity, stemming from two NP-hard problems: the vehicle routing problem (VRP) and the two-dimensional bin packing problem (2D-BPP), less attention has been paid to developing exact algorithms. Bridging this gap, this article presents an exact algorithm that integrates advanced machine learning techniques, specifically a novel combination of attention and recurrence mechanisms. This integration accelerates the state-of-the-art exact algorithm by a median of 29.79% across various problem instances. Moreover, the proposed algorithm successfully resolves an open instance in the standard test-bed, demonstrating significant improvements brought about by the incorporation of machine learning models. Code is available at <https://github.com/xyffff/NCG-for-2L-CVRP>.

## 1 Introduction

The vehicle routing problem (VRP) is crucial in various logistics applications such as express systems, industrial warehousing, and on-demand delivery [Zong *et al.*, 2021]. The basic form of VRP, or the capacitated vehicle routing problem (CVRP), focuses on optimizing a fleet of vehicles to meet customer demands with the goal of minimizing the total travel costs. An important extension of CVRP is the vehicle routing problem with the two-dimensional loading constraints (2L-CVRP) [Iori *et al.*, 2007]. In 2L-CVRP, the challenge involves transporting rectangular items, each with specific length, width, and weight. The loading constraints in 2L-CVRP revolve around two main aspects: the orientation of items, which concerns whether an item can be rotated, and sequential loading, which concerns whether items belonging to successive customers along a route are allowed to be moved [Fuellerer *et al.*, 2009]. The latter is often referred to as the

last-in-first-out (LIFO) or rear-loading constraints. This problem is especially relevant in scenarios like transporting furniture or large industrial equipment where stacking items is not feasible due to safety or operational concerns. A real-world example of 2L-CVRP’s application is seen at Opein, a company distributing large equipment, highlighting the model’s practical importance. The practical value of 2L-CVRP has motivated numerous studies including both exact algorithms and heuristics, and solving new practical variants [Pollaris *et al.*, 2015]. In this paper, we focus on 2L-CVRP with fixed item orientation and LIFO rule.

The 2L-CVRP encompasses two NP-hard problems: the CVRP and the two-dimensional bin packing problem (2D-BPP) stemming from the loading constraints. Given this complexity, the 2L-CVRP is solved approximately in most of studies [Gendreau *et al.*, 2008; Wei *et al.*, 2018]. However, the development of exact solvers is crucial for a deeper understanding of the problem’s structure and for evaluating the performance gap between approximate solutions and optimal solutions. To achieve this end, recent studies [Côté *et al.*, 2020; Zhang *et al.*, 2022a; Zhang *et al.*, 2022b] have proposed several efficient exact algorithms and closed many open instances. The state-of-the-art (SOTA) exact algorithm for the 2L-CVRP [Zhang *et al.*, 2022b] relies on the column generation (CG), which features repeatedly solving a challenging pricing problem (PP) and thus is regarded as one of the main bottlenecks blocking us from solving more open instances.

Recently, the integration of machine learning (ML) with CG has shown promise in solving combinatorial optimization problems more efficiently, while still aiming for optimal solutions [Morabit *et al.*, 2021; Zhang *et al.*, 2022c; Shen *et al.*, 2022; Chi *et al.*, 2022; Morabit *et al.*, 2023; Yuan *et al.*, 2023]. ML’s ability to learn from data and make probabilistic predictions offers a potential acceleration in the CG process, particularly in problems where traditional methods are either too slow or impractical.

However, applying ML-based CG methods to the 2L-CVRP, especially under the LIFO rule, introduces distinct challenges. Existing ML-based CG algorithms often solve the PP and generate multiple columns by traditional methods, with ML typically applied as a post-processing tool for column selection [Morabit *et al.*, 2021; Chi *et al.*, 2022; Yuan *et al.*, 2023]. Some studies have attempted to bypass solving PP altogether, using ML to directly generate columns

\*Corresponding author.

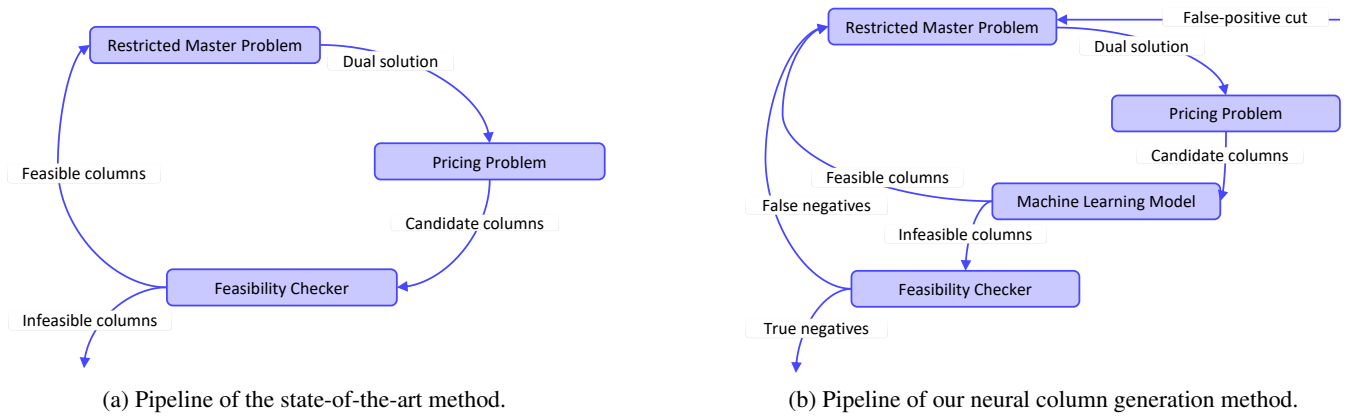


Figure 1: Comparative illustration of the SOTA method and our NCG method for the 2L-CVRP with the LIFO rule.

[Shen *et al.*, 2022]. However, the PP in the context of 2L-CVRP, which combines the elementary shortest path problem with resource constraints (ESPPRC) and the 2D-BPP with the LIFO rule, is non-trivial to be addressed through traditional methods or ML models alone. While [Morabit *et al.*, 2023] proposed an ML approach to prune the ESPPRC pricing graph, applying this strategy to the PP of 2L-CVRP is difficult due to the added complexity of 2D-BPP. A related work to ours is [Zhang *et al.*, 2022c], which tackled a variant of 2L-CVRP without the LIFO constraint, utilizing feed-forward networks (FF) to accelerate the feasibility checking of columns generated by the pricing algorithm. Our study extends this approach by incorporating the LIFO rule, adding further complexity. We propose a novel ML model that leverages attention mechanism [Vaswani *et al.*, 2017] for homogeneous items within the same customer and recurrence mechanism [Cho *et al.*, 2014] for heterogeneous items across different customers. This model is used to predict the feasibility of columns generated from the ESPPRC, ensuring compliance with the 2D-BPP constraints with the LIFO rule. Our method provides a more efficient alternative to the traditional, time-consuming feasibility checker as exemplified in [Zhang *et al.*, 2022b], achieving a median acceleration of 29.79% and successfully solving one challenging open instance for the first time. For a visual representation of our algorithmic pipeline, see Figure 1b.

In summary, the main contributions of this paper are:

- We propose a neural column generation (NCG) algorithm that combines the state-of-the-art column generation for the 2L-CVRP with the newly developed ML model.
- Our novel ML model integrates attention and recurrence mechanisms, along with a symmetry-based data augmentation technique. It effectively tackles the 2D-BPP with the LIFO rule, achieving an overall accuracy of around 95%.
- The NCG algorithm, when tested on standard benchmark instances, demonstrates a significant reduction in runtime, with a median decrease of 29.79% compared to the state-of-the-art column generation algorithm.

- Additionally, the NCG algorithm is incorporated into the state-of-the-art branch-and-price-and-cut (BPC) algorithm, successfully solving an open instance for the first time.

## 2 Related Work

In this section, we review the existing methods for 2L-CVRP and recent advances in enhance CG with ML techniques.

**Non-learning Methods for 2L-CVRP.** The vehicle routing problem with two-dimensional loading constraints has been an active research area due to its practical relevance in logistics. The seminal work by [Iori *et al.*, 2007] introduced a branch-and-cut (B&C) algorithm adapted for 2L-CVRP by introducing infeasible-path constraints. They also propose an exact packing algorithm as a subroutine for the B&C algorithm. The method set a benchmark for the problem, optimally solving instances with up to 35 customers. Subsequent heuristic approaches, such as the tabu search heuristic by [Gendreau *et al.*, 2008] and the guided local search by [Zachariadis *et al.*, 2009], improved solution times and led to several new best solutions. The most effective heuristic for the 2L-CVRP to date involves a simulated annealing heuristic coupled with a local search-based packing algorithm [Wei *et al.*, 2018], which has shown superior solution quality. The 2L-CVRP with unloading sequential constraints has been tackled by [Pinto *et al.*, 2013] through a column-generation-based heuristic. This approach initially omits loading constraints and then applies a heuristic to construct feasible packing solutions. The same authors later developed a more refined branch-and-price (B&P) algorithm [Pinto *et al.*, 2016] that integrates a variable neighborhood search (VNS) algorithm [Pinto *et al.*, 2015] into the pricing problem. Despite these advancements, exact methods still face challenges with larger instances. The B&C algorithm by [Côté *et al.*, 2020] incorporated advanced exact packing algorithms [Côté *et al.*, 2014b], achieving a significant breakthrough in benchmark instances. [Zhang *et al.*, 2022a] further enhanced this approach by introducing a heuristic for separating infeasible set inequalities at fractional nodes, closing several open instances and improving dual bounds. Later on, [Zhang *et al.*, 2022b]

propose an exact CG algorithm addressing the loading constraints with a novel data structure *L-Trie*. The CG algorithm leading to a successful BPC algorithm which closed many open instances for the first time and remains to be the state-of-the-art exact column generation for the 2L-CVRP. For additional studies on 2L-CVRP, please refer to [Wang *et al.*, 2009; Pollaris *et al.*, 2015; Iori *et al.*, 2021].

**Neural Column Generation.** Recent years researchers have become increasingly interested in ML to accelerate optimization tasks [Bengio *et al.*, 2021], and several learning-based methods have been proposed for specific problems solved by CG. For example, [Morabit *et al.*, 2021] employed a Graph Neural Network (GNN) for column selection in CG as a supervised binary classification task, aiming to mitigate degeneracy in the restricted master problem (RMP). This approach involved representing columns and constraints as a bipartite graph, with the GNN predicting whether to include or exclude each column during CG iterations. Building on this, [Chi *et al.*, 2022; Yuan *et al.*, 2023] modeled CG as a Markov decision process, using reinforcement learning for column selection during CG iterations. These methods, using GNN as Q-function approximator, proved to be more efficient than traditional greedy policies in problems such as the cutting stock problem. All these methods could be viewed as a post-processing step after CG, i.e., applying ML models to select from the generated columns rather than selecting greedily. [Morabit *et al.*, 2023] proposed a supervised ML-based algorithm to prune the pricing network in the form of ESPPRC, alternating between reduced and complete graphs to accelerate the CG process. [Shen *et al.*, 2022] utilized a support vector machine (SVM) for the graph coloring problem, directly generating columns by sampling from the SVM to enhance the CG process. The most related work to ours is [Zhang *et al.*, 2022c], which addressed the vehicle routing problem without the LIFO rule. They utilized FF for heuristic validation of columns generated from the ESPPRC, reducing the dependence on the exact solver [Côté *et al.*, 2014a]. However, our paper addresses a more complex variant of the problem by including additional loading constraints, specifically the LIFO rule, which significantly increases the complexity of the CG process.

### 3 Background

In this section, we first present the mathematical formulation of the 2L-CVRP with the LIFO rule as well as the CG process. Then, we introduce the SOTA exact method for solving the 2L-CVRP with the LIFO rule.

#### 3.1 Problem Formulation

The 2L-CVRP with the LIFO rule is defined on a complete undirected graph  $G = (V, E)$ , where  $V = \{0, 1, 2, \dots, n, n+1\}$  stands for the set of vertices consisting of customers  $V_c = \{1, 2, \dots, n\}$  and the depot 0. Vertex  $n+1$  represents a copy of the depot. The connections between any pair of vertices are depicted by the edge set  $E$ .  $\forall e \in E$ ,  $c_e$  represents the traveling cost associated with edge  $e$ . An alternative representation of an edge  $(v_i, v_j)$  is also used. Set  $K$  represents a fleet of homogeneous vehicles which are available at

the depot. A loading area characterized by  $H$  and  $W$  as the length and width, respectively, is a property attached to each vehicle. The loading areas of all the vehicles are the same. Naturally, we have the total area of the loading surface of any vehicle equal to  $A = H \times W$ . Each vehicle also has a weight capacity denoted as  $Q$ .

As for the customers,  $\forall i \in V_c$  is characterized by a set  $M_i$  of rectangles. Any item  $m \in M_i$  is marked by width  $w_{i,m}$ , length  $h_{i,m}$ , and weight  $q_{i,m}$ . Let  $\nu_i$  and  $c_i$  represent the total area and the total weights of all the items in customer  $i$ . In other words,  $\nu_i = \sum_{m \in M_i} w_{i,m} h_{i,m}$  and  $q_i = \sum_{m \in M_i} q_{i,m}$ . The total number of the items in  $G$  is equal to  $|M|$ , where  $M$  is the union of the item sets of all the customers. The 2L-CVRP calls for planning routes for the fleet such that the demands of the customers are covered while respecting the following constraints:

1. Each customer has to be visited exactly once;
2. The total weight of the items to be delivered by a vehicle cannot exceed  $Q$ ;
3. The carried items have to be packed in the loading area without collision;
4. Items are not allowed to be rotated over the course of packing;
5. Items delivered to subsequent customers cannot be moved when unloading the items for the current customer (the LIFO rule).

The 2L-CVRP can be formulated as a set partitioning (SP) problem:

$$\min \sum_{r \in \Omega} c_r \lambda_r, \quad (1)$$

$$\text{s.t.} \sum_{r \in \Omega} \lambda_r = |K|, \quad (2)$$

$$\sum_{r \in \Omega} a_{i,r} \lambda_r = 1, \quad \forall i \in V_c, \quad (3)$$

$$\lambda_r \in \{0, 1\}, \quad \forall r \in \Omega, \quad (4)$$

where  $\Omega$  represents the collection of feasible routes;  $c_r$  represents the total traveling cost of route  $r$ ;  $\lambda_r$  is a binary decision variable indicating whether route  $r$  is selected as a part of the solution;  $a_{i,r}$  is a binary indicator, where  $a_{i,r} = 1$  if vertex  $i$  is visited in route  $r$ , and  $a_{i,r} = 0$  otherwise.

Equation (1) defines the objective function for the set partitioning problem. Constraint set (2) imposes that there should be exactly  $|K|$  routes to be selected as we assume that there is no idle vehicle in the fleet. This has been a convention when it comes to developing exact algorithms for the 2L-CVRP [Iori *et al.*, 2007]. Constraint set (3) calls for that each customer should be visited exactly once. Constraint set (4) defines the domain of the decision variables.

In practical applications, directly solving the SP formulation is infeasible due to the need to enumerate all routes in  $\Omega$ . Typically, a smaller subset of  $\Omega$  is selected to create a reduced version of the problem, known as the restricted master problem. Solving the RMP yields a solution,  $\lambda^*$ , that minimizes the objective value of the reduced formulation. However,  $\lambda^*$

may not be optimal for the original problem. To potentially improve upon  $\lambda^*$ , a sub-problem called the pricing problem is solved to identify any routes (or columns) in  $\Omega$  that could enhance the solution. This iterative process, known as column generation, continues until no further improvements are found. For a detailed discussion on this approach, readers are referred to Chapter 2 in [Desaulniers *et al.*, 2006].

The formulation of the underlying pricing problem is as follows:

$$\min \sum_{e \in E} \bar{d}_e x_e - \pi_f, \quad (5)$$

$$\text{s.t.} \sum_{e \in \delta(i)} x_e = 2, \forall i \in V, \quad (6)$$

$$\sum_{e \in \delta(S)} x_e \geq 2, \forall S \subset V_c, 1 < |S| < n - 1, \forall i \in S, \quad (7)$$

$$\sum_{(i,j) \in E} x_{ij}(q_i + q_j) \leq 2Q, \quad (8)$$

$$\sum_{(i,j) \in E} x_{ij}(\nu_i + \nu_j) \leq 2A, \quad (9)$$

$$\sum_{e \in E(S,\sigma)} x_e \leq |S| - 1, \forall (S, \sigma) \text{ such that } \sigma \notin \Sigma(S), \quad (10)$$

$$x_e \in \{0, 1\}, \forall e \in E, \quad (11)$$

where  $\bar{d}_e$  is the reduced cost defined as  $\bar{d}_{i,j} = c_{i,j} - \frac{1}{2}\pi_i - \frac{1}{2}\pi_j$ ,  $\forall (i, j) \in E$ . The binary decision variable  $x_e$  indicates whether edge  $e$  in  $E$  is used. The set  $\delta(i)$  denotes edges incident to vertex  $i$ , and  $\delta(S)$  denotes edges connecting vertices inside  $S$  with those outside, where  $S$  is a subset of  $V$ .  $\Sigma(S)$  is the set of all feasible permutations of vertices in  $S$ . The route constructed by set  $S$  in order  $\sigma$  is represented as  $(S, \sigma)$ , with  $E(S, \sigma)$  being its edge set. The dual variables  $\pi_i$  and  $\pi_f$  are associated with constraints 3 and 2, respectively.

The objective function (constraint set 5) aims to minimize the reduced cost. Constraint set (6) ensures proper degree constraints for vertices, and constraint set (7) addresses sub-tour elimination [Desrochers and Laporte, 1991]. Constraint sets (8) and (9) ensure that the vehicle's capacity and loading surface area limits are not exceeded. Finally, constraint set (10) imposes restrictions due to the loading constraints.

### 3.2 State-of-the-Art Pipeline

Figure 1a illustrates the pipeline of the SOTA CG-based algorithm for the 2L-CVRP with the LIFO rule, as proposed by [Zhang *et al.*, 2022b]. The process starts with a restricted master problem which essentially enumerates a subset of set  $\Omega$ . The next step involves solving the linear relaxation of this problem to obtain a dual solution, which then facilitates the establishment of the PP. The labeling algorithm, enhanced by trie [Brass, 2010], completion bounds, and ng-route relaxation [Baldacci *et al.*, 2011], efficiently prices out improving columns without checking the loading feasibility. The last step involves filtering out infeasible columns with the exact checker [Côté *et al.*, 2014b] and adding the feasible ones to the restricted master problem, iterating until no feasible col-

umn is found. For a more detailed understanding, please refer to [Zhang *et al.*, 2022b].

## 4 Methodology

In this section, we first introduce the motivation behind our integrated NCG approach as well as the outline of our approach. Then, we elaborate on the architecture of our machine learning model, including two key mechanisms and an important data augmentation technique.

### 4.1 Motivation and Outline

The current SOTA pipeline employs an exact feasibility checker to evaluate each candidate column generated by the labeling algorithm, necessitating solving the 2D-BPP with LIFO rule, a known NP-hard problem. Our proposed approach, however, introduces a ML model to refine this process by reducing the dependency on solving 2D-BPP. As illustrated in Figure 1b, the ML model is positioned before the feasibility checker, categorizing candidate columns into 'feasible' and 'infeasible'. Feasible columns directly enter the restricted master problem, while infeasible ones are further checked to rectify potential false negatives and thus preventing optimal columns being discarded. To manage false positives, the loading feasibility of variables in the optimal basis is checked during master problem resolution, adding false-positive cuts to exclude infeasible variables from the feasible solution space.

This approach leverages the insight that columns not part of the optimal basis can bypass exact checking. However, its success depends on the ML model's accuracy, since frequent erroneous predictions can increase computational iterations.

### 4.2 Machine Learning Model

Note that candidate columns generated by the labelling algorithm might violate the loading constraints. To address this, we frame it as a binary classification problem, developing a ML model to predict the feasibility of each candidate column. Specifically, our model predicts the probability with which each column respects the loading constraints.

Our model's architecture comprises a parallel embedding mechanism and a recursive processing strategy, designed to effectively capture both homogeneous and heterogeneous features of items within each column. For a given input candidate column, represented as a sequence of item sets  $\left[ \{x_{i,m}\}_{1 \leq m \leq |M_i|} \right]_{1 \leq i \leq n}$ , each item  $x_{i,m}$  is represented by normalized dimensions  $\left[ \frac{w_{i,m}}{W}, \frac{h_{i,m}}{H} \right]$ . We employ an attention mechanism [Vaswani *et al.*, 2017] for each customer  $i$  to integrate features of homogeneous items (those belonging to the same customer). To handle heterogeneous features (items across different customers), a GRU [Cho *et al.*, 2014] is utilized to process the sequence, incorporating the order information which is crucial in adhering to the LIFO constraints. Moreover, a symmetry-based data augmentation technique is employed to incorporate permutation invariance into the model. For a visual representation of our model's architecture, please refer to Figure 2.

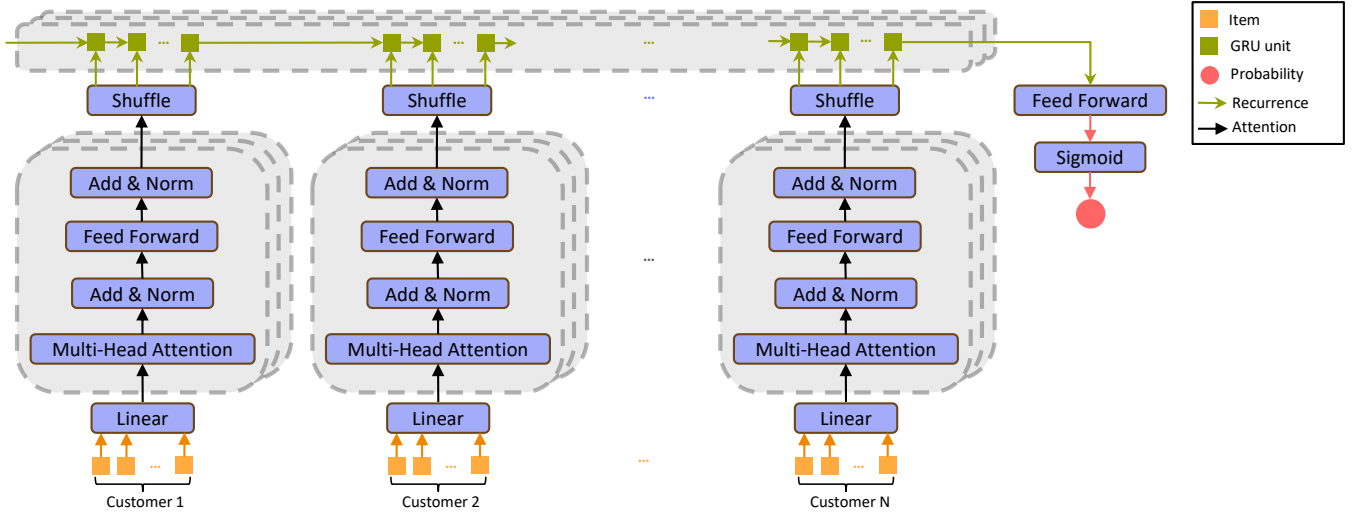


Figure 2: The architecture of the machine learning model employed in our NCG approach, illustrating the integration of attention and GRU mechanisms for column feasibility prediction.

### Item-Level Attention Mechanism

In the 2L-CVRP context, items belonging to the same customer are homogeneous and exempt from LIFO constraints, making the Multi-Head Attention (MHA) mechanism a suitable choice [Kool *et al.*, 2018; Kwon *et al.*, 2020]. Our model uses MHA to capture the shared features of items within the same customer. The initial embedding for the  $i$ -th customer's  $m$ -th item is formulated as  $h_{i,m}^0 = W_0 x_{i,m} + b_0$ , where  $W_0$  is the initial projection matrix and  $b_0$  is the bias vector. The architecture includes skip-connections [He *et al.*, 2015], feed-forward networks, and layer normalization (LN) [Ba *et al.*, 2016] in each sublayer. The embedding for the  $i$ -th customer's  $m$ -th item is iteratively updated in the  $l$ -th layer, as depicted in the following equations:

$$\hat{h}_{i,m}^l = \text{LN}^l \left( h_{i,m}^{l-1} + \text{MHA}_{i,m}^l \left( h_{i,1}^{l-1}, \dots, h_{i,|M_i|}^{l-1} \right) \right), \quad (12)$$

$$h_{i,m}^l = \text{LN}^l \left( \hat{h}_{i,m}^l + \text{FF}^l \left( \hat{h}_{i,m}^l \right) \right). \quad (13)$$

The MHA mechanism at the core of our attention layer is defined as follows:

$$Q_{i,m}^j, K_{i,m}^j, V_{i,m}^j = W_Q^j h_{i,m}, W_K^j h_{i,m}, W_V^j h_{i,m}, \quad (14)$$

$$A_{i,m}^j = \text{softmax} \left( Q_{i,m}^j K_{i,m}^{jT} / \sqrt{d_k} \right) V^j, \quad (15)$$

$$\text{MHA}_{i,m} = \text{Concat} \left( A_{i,m}^1, A_{i,m}^2, \dots, A_{i,m}^H \right) W_O, \quad (16)$$

where  $j = 1, 2, \dots, H$  and  $d_k = d_h/H$ . Here,  $H$  is the number of attention heads,  $d_h$  is the dimension of the item embedding, and  $Q_{i,m}^j, K_{i,m}^j, V_{i,m}^j$  represent the query, key, and value vectors, respectively.  $W_O$  is the projection matrix utilized to project the final MHA output. The final embedding of each item after  $L$  layers is denoted by  $h_{i,m} = h_{i,m}^L$ .

### Customer-Level Recurrence Mechanism

In the context of the 2L-CVRP with the LIFO rule, a sequential order relationship exists among customers, indicating that items from different customers are inherently heterogeneous. This order relationship dictates that for customer  $i$

visited before customer  $i + 1$ , items of customer  $i$  (denoted as  $x_{i,1}, x_{i,2}, \dots, x_{i,|M_i|}$ ) must be loaded into the vehicle after items of customer  $i + 1$  ( $x_{i+1,1}, x_{i+1,2}, \dots, x_{i+1,|M_{i+1}|}$ ).

To model the recursive relationships among items of different customers, we use GRU as follows:

$$\tilde{h}_t = \text{GRU} \left( h_t, \tilde{h}_{t-1} \right), \quad (17)$$

where  $\tilde{h}_t$  denotes the hidden state at time step  $t$  and  $\tilde{h}_0 = \mathbf{0}$ . In our approach, customers are processed sequentially, inputting one item  $h_t$  per time step  $t$  into the GRU, with the total number of time steps  $T$  equaling the total items,  $T = \sum_{i=1}^n |M_i|$ . Upon processing the last item of the last customer, the final state  $\tilde{h}_T$  of the GRU is transformed into a probability via a FF network and a sigmoid function:

$$\text{probability} = \text{sigmoid} \left( \text{FF} \left( \tilde{h}_T \right) \right). \quad (18)$$

### Data Augmentation with Permutation Invariance

As discussed in Section 4.2, items belonging to the same customer are homogeneous and are not constrained by any specific order. This characteristic allows for the application of permutation invariance as a data augmentation strategy, reflecting the symmetry of combinatorial problems [Kwon *et al.*, 2020; Kim *et al.*, 2022]. By permuting the items of each customer, we can generate new, equivalent instances, expanding the training dataset and mitigating early overfitting.

Specifically, after applying the item-level multi-head attention mechanism to all items of each customer  $i$ , we perform a permutation  $\pi_i$ , shuffling the sequence  $(1, 2, \dots, |M_i|)$  to produce varied item orderings. This process can be represented mathematically as:

$$h_{i,1}, h_{i,2}, \dots, h_{i,|M_i|} = h_{i,\pi_i(1)}, h_{i,\pi_i(2)}, \dots, h_{i,\pi_i(|M_i|)}, \quad (19)$$

This permutation is applied independently to each customer's set of items before they are processed by the GRU.




Figure 3: Box plot depicting the percentage increase in efficiency of our NCG method over the SOTA across different instance categories.

## 5 Experiments

### 5.1 Experimental Settings

**Data Sets.** Our machine learning model is trained on a unique set of packing problem instances, distinct from the benchmark instances used in [Iori *et al.*, 2007], ensuring an unbiased performance evaluation. The training dataset is derived by solving the 2L-CVRP problems as well as 2L-VRPTW (the 2L-CVRP with time window constraints). In particular, we stored all the packing instances in the course of solving the routing problems. There are in total two batches of the training samples: one from 2L-VRPTW instances and another from 2L-CVRP instances. The VRP instances are randomly generated as per the description in [Zhang *et al.*, 2022c], ensuring a diverse distribution distinct from benchmark instances. The benchmark instances are categorized into five families, as defined in [Iori *et al.*, 2007]. For details, please refer to the Supplementary Material. We exclude family 1, as it is equivalent to the one-dimensional loading scenario, and family 5, as highlighted by [Zhang *et al.*, 2022b], due to its focus on very small items and a specialized column generation variant, making it less relevant for testing our algorithm in standard 2L-CVRP contexts. In Section 5, terms ‘family’ and ‘packing class (PC)’ are used interchangeably.

**Evaluation Metrics.** For the 2L-CVRP instances, our experiment compares the NCG approach to the SOTA CG method by solving the linear relaxation of Formulation 1 - 4. Given that both algorithms are exact, their optimal objective values are expected to match. Our primary interest lies in the difference in computational time. Let  $T_{NCG}$  represent the walltime for NCG and  $T_{SOTA}$  for SOTA CG. The percentage gap, calculated as  $\frac{T_{SOTA} - T_{NCG}}{T_{NCG}} \times 100\%$ , serves to quantify the time savings and performance increase, with higher values indicating greater efficiency gains for NCG. Furthermore, we track the number of column generation iterations to evaluate the impact of the ML-induced false-positive cuts on the frequency of column generation.

**Hardware.** During the training phase of the ML model, experiments were conducted on an AMD EPYC 7V13 64-Core CPU @ 2.45GHz with an NVIDIA A100 GPU. Post training,




Figure 4: Box plot demonstrating the percentage increase in column generation counts for our NCG method over the SOTA.

the ML model was serialized and integrated into the NCG algorithm using Torch C++. To ensure a fair comparison, both the NCG algorithm and the SOTA algorithm [Zhang *et al.*, 2022b] were implemented in C++ and evaluated on an Intel i5-10600KF processor @ 4.10 GHz.

### 5.2 Main Results

#### Acceleration of SOTA Algorithm

We evaluated the NCG against the SOTA on benchmark instances with up to 50 customers from families 2 - 4. As depicted in Figure 3, NCG demonstrates a significant advantage over the SOTA algorithm, with a median performance improvement of 29.79% across all tested instances. Notably, the performance gains are more pronounced in PC3 and PC4, with median percentage gaps of 44.97% and 99.22%, respectively. This improvement correlates with the higher accuracy of our ML model in PC3 and PC4, as indicated in Table 2, suggesting fewer iterations required to rectify false predictions. In contrast, the least improvement occurs with PC2, with a median gain of 9.38%, as the ML model reaches the worst accuracy in PC2.

Figure 4 also presents an intriguing observation: an increase in the number of CG iterations due to the incorporation of false positive cuts. Interestingly, for each instance family, the increase in CG iterations inversely correlates with performance improvement. Despite this, the additional time incurred by these extra iterations is more than compensated for by the time savings from the ML model. This highlights the effectiveness of our NCG approach in enhancing overall computational efficiency despite the potential for additional CG iterations. For detailed results on each instance, please refer to the Supplementary Material.

#### Solving Open Instance

Our NCG algorithm was integrated into the BPC framework developed in [Zhang *et al.*, 2022b]. This integrated approach was tested on benchmark instances, leading to the successful resolution of an open instance. As shown in Table 1, for the instance 2304, our NCG-based BPC method obtained an optimal objective value (Optimal Obj) of 1068, while the existing SOTA BPC algorithm was unable to reach an op-

Instance Name	NCG-based BPC		SOTA BPC		
	Optimal Obj	Run Time (s)	Lower Bound	Primal Bound	Run Time (s)
2l-cvrp-2304	1068	49544.6	1066.78	1069	69948.96 <sup>†</sup>

<sup>†</sup> Algorithm terminated early due to memory overload.

Table 1: Performance comparison on the solved open instance 2304 between NCG-based BPC and SOTA BPC.

Algorithm	Overall		PC2		PC3		PC4	
	TPR	TNR	TPR	TNR	TPR	TNR	TPR	TNR
Ours	<b>94.08%</b>	<b>96.80%</b>	<b>84.26%</b>	<b>97.33%</b>	<b>86.79%</b>	95.83%	<b>97.81%</b>	<b>88.81%</b>
w/o augmentation	92.84%	96.78%	81.54%	97.22%	83.07%	<b>96.69%</b>	97.47%	87.31%
w/o attention mechanism	91.44%	95.22%	77.56%	95.88%	80.76%	94.68%	96.43%	82.84%
w/o recurrence mechanism	92.48%	96.07%	81.70%	96.86%	81.65%	93.81%	95.40%	88.06%

Table 2: Comparative results showcasing TPR and TNR metrics across different model configurations for Overall and PC2, PC3, PC4 categories.




Figure 5: Comparison of validation losses over epochs for the baseline and augmented models.

timal solution, yielding only a lower bound of 1066.78 and a primal bound of 1069, despite running for approximately 1.4 times longer than our approach. Furthermore, the SOTA BPC algorithm suffered from memory overload, whereas our NCG-based BPC approach did not, indicating its more efficient memory usage. This efficiency is attributed to the reduced number of routes that needed to be checked by the exact packing algorithm, resulting in lower memory demands for the BPC algorithm.

### 5.3 Ablation Study

To validate the contributions of the attention and recurrence mechanisms and the data augmentation technique described in Section 4.2, we performed an ablation study. We compared the following configurations:

1. The full NCG model that integrates both attention and recurrence mechanisms with data augmentation.
2. A variant without data augmentation.
3. A model where the attention mechanism is replaced by a multilayer perceptron (MLP).
4. A model where the recurrence mechanism is replaced by a transformer encoder with sinusoidal positional encoding [Vaswani *et al.*, 2017].

As shown in Table 2, the complete NCG model outperforms the other variants in terms of true positive rate (TPR) and true negative rate (TNR). These results suggest that the item-level attention mechanism contributes the most, followed by the customer-level recurrence mechanism, and then the data augmentation technique. Despite the transformer encoder’s capability to process sequential information, it underperforms compared to the GRU-based model, highlighting the importance of explicit modeling of sequential relationships in the 2L-CVRP with the LIFO rule.

Figure 5 shows that, compared to the baseline model which overfits sooner, the model with permutation invariance maintains a lower validation loss for a longer period, confirming the effectiveness of the data augmentation strategy in enhancing generalization capacity.

## 6 Conclusion

In this work, we introduced a ML-based exact algorithm to solve the 2L-CVRP with the LIFO rule. Our approach integrates neural column generation into the SOTA pipeline, formulating the feasibility checking of candidate columns as a binary classification problem. To preserve solution optimality, we implemented post-processing steps for handling false negatives and positives. The ML model in our framework utilizes an attention mechanism for capturing homogeneous item features and a recurrence mechanism for heterogeneous features. Additionally, we employed data augmentation exploiting the problem’s symmetry.

Experimental evaluations on benchmark instances demonstrate that our NCG method notably accelerates the SOTA algorithm by a median of 29.79% across various problem instances. Significantly, our approach also successfully solves an open instance, marking a substantial contribution to the field. This achievement highlights the potential of integrating machine learning techniques into traditional optimization problems for enhanced performance and efficiency.

## References

- [Ba *et al.*, 2016] Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E Hinton. Layer normalization. *arXiv preprint arXiv:1607.06450*, 2016.
- [Baldacci *et al.*, 2011] Roberto Baldacci, Aristide Mingozzi, and Roberto Roberti. New route relaxation and pricing strategies for the vehicle routing problem. *Operations Research*, page 1269–1283, Oct 2011.
- [Bengio *et al.*, 2021] Yoshua Bengio, Andrea Lodi, and Antoine Prouvost. Machine learning for combinatorial optimization: A methodological tour d’horizon. *European Journal of Operational Research*, 290(2):405–421, 2021.
- [Brass, 2010] Peter Brass. *Advanced Data Structures*, page 68–72. Aug 2010.
- [Chi *et al.*, 2022] Cheng Chi, Amine Mohamed Aboussalah, Elias Boutros Khalil, Juyoung Wang, and Zoha Sherkat-Masoumi. A deep reinforcement learning framework for column generation. In *Advances in Neural Information Processing Systems*, 2022.
- [Cho *et al.*, 2014] Kyunghyun Cho, Bart Van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. Learning phrase representations using rnn encoder-decoder for statistical machine translation. *arXiv preprint arXiv:1406.1078*, 2014.
- [Côté *et al.*, 2020] Jean-François Côté, Michel Gendreau, and Jean-Yves Potvin. The vehicle routing problem with stochastic two-dimensional items. *Transportation Science*, 54(2):453–469, 2020.
- [Côté *et al.*, 2014a] Jean-François Côté, Mauro Dell’Amico, and Manuel Iori. Combinatorial benders’ cuts for the strip packing problem. *Operations Research*, page 643–661, Jun 2014.
- [Côté *et al.*, 2014b] Jean-François Côté, Michel Gendreau, and Jean-Yves Potvin. An exact algorithm for the two-dimensional orthogonal packing problem with unloading constraints. *Operations Research*, page 1126–1141, Oct 2014.
- [Desaulniers *et al.*, 2006] Guy Desaulniers, Jacques Desrosiers, and Marius M Solomon. *Column generation*, volume 5. Springer Science & Business Media, 2006.
- [Desrochers and Laporte, 1991] Martin Desrochers and Gilbert Laporte. Improvements and extensions to the miller-tucker-zemlin subtour elimination constraints. *Operations Research Letters*, 10(1):27–36, 1991.
- [Fuellerer *et al.*, 2009] Guenther Fuellerer, Karl F. Doerner, Richard F. Hartl, and Manuel Iori. Ant colony optimization for the two-dimensional loading vehicle routing problem. *Computers & Operations Research*, 36(3):655–673, Mar 2009.
- [Gendreau *et al.*, 2008] Michel Gendreau, Manuel Iori, Gilbert Laporte, and Silvaro Martello. A tabu search heuristic for the vehicle routing problem with two-dimensional loading constraints. *Networks: An International Journal*, 51(1):4–18, 2008.
- [He *et al.*, 2015] Kaiming He, X. Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 770–778, 2015.
- [Iori *et al.*, 2007] Manuel Iori, Juan-José Salazar-González, and Daniele Vigo. An exact approach for the vehicle routing problem with two-dimensional loading constraints. *Transportation Science*, page 253–264, May 2007.
- [Iori *et al.*, 2021] Manuel Iori, Vinícius L. de Lima, Silvano Martello, Flávio K. Miyazawa, and Michele Monaci. Exact solution techniques for two-dimensional cutting and packing. *European Journal of Operational Research*, 289(2):399–415, Mar 2021.
- [Kim *et al.*, 2022] Minsu Kim, Junyoung Park, and Jinkyoo Park. Sym-NCO: Leveraging symmetry for neural combinatorial optimization. In *Advances in Neural Information Processing Systems*, 2022.
- [Kool *et al.*, 2018] Wouter Kool, Herke van Hoof, and Max Welling. Attention, learn to solve routing problems! In *International Conference on Learning Representations*, 2018.
- [Kwon *et al.*, 2020] Yeong-Dae Kwon, Jinho Choo, Byoungjip Kim, Iljoo Yoon, Youngjune Gwon, and Seungjai Min. Pomo: Policy optimization with multiple optima for reinforcement learning. In *Advances in Neural Information Processing Systems*, volume 33, pages 21188–21198. Curran Associates, Inc., 2020.
- [Morabit *et al.*, 2021] Mouad Morabit, Guy Desaulniers, and Andrea Lodi. Machine-learning-based column selection for column generation. *Transportation Science*, 55(4):815–831, 2021.
- [Morabit *et al.*, 2023] Mouad Morabit, Guy Desaulniers, and Andrea Lodi. Machine-learning-based arc selection for constrained shortest path problems in column generation. *INFORMS Journal on Optimization*, 5(2):191–210, 2023.
- [Pinto *et al.*, 2013] Telmo Pinto, Cláudio Alves, and J Valério de Carvalho. Column generation based heuristic for a vehicle routing problem with 2-dimensional loading constraints: a prototype. In *XI Congresso Galego de Estatística e Investigación de Operacións, Spain*, 2013.
- [Pinto *et al.*, 2015] Telmo Pinto, Cláudio Alves, and José Valério de Carvalho. Variable neighborhood search for the elementary shortest path problem with loading constraints. In *Computational Science and Its Applications—ICCSA 2015: 15th International Conference, Banff, AB, Canada, June 22–25, 2015, Proceedings, Part II 15*, pages 474–489. Springer, 2015.
- [Pinto *et al.*, 2016] Telmo Pinto, Cláudio Alves, and José Valério de Carvalho. A branch-and-price algorithm for the vehicle routing problem with 2-dimensional loading constraints. In *International Conference on Computational Logistics*, pages 321–336. Springer, 2016.



- [Pollaris *et al.*, 2015] Hanne Pollaris, Kris Braekers, An Caris, Gerrit K. Janssens, and Sabine Limbourg. Vehicle routing problems with loading constraints: state-of-the-art and future directions. *OR Spectrum*, 37(2):297–330, Mar 2015.
- [Shen *et al.*, 2022] Yunzhuang Shen, Yuan Sun, Xiaodong Li, Andrew Eberhard, and Andreas Ernst. Enhancing column generation by a machine-learning-based pricing heuristic for graph coloring. In *Proceedings of the 36th AAAI Conference on Artificial Intelligence, AAAI 2022*, volume 36, pages 9926–9934, United States of America, June 2022.
- [Vaswani *et al.*, 2017] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.
- [Wang *et al.*, 2009] Fan Wang, Yi Tao, and Ning Shi. A survey on vehicle routing problem with loading constraints. *2009 International Joint Conference on Computational Sciences and Optimization*, 2:602–606, 2009.
- [Wei *et al.*, 2018] Lijun Wei, Zhenzhen Zhang, Defu Zhang, and Stephen C.H. Leung. A simulated annealing algorithm for the capacitated vehicle routing problem with two-dimensional loading constraints. *European Journal of Operational Research*, page 843–859, Mar 2018.
- [Yuan *et al.*, 2023] Haofeng Yuan, Lichang Fang, and Shiji Song. A reinforcement-learning-based multiple-column selection strategy for column generation. *arXiv preprint arXiv:2312.14213*, 2023.
- [Zachariadis *et al.*, 2009] Emmanouil E. Zachariadis, Christos D. Tarantilis, and Christos T. Kiranoudis. A guided tabu search for the vehicle routing problem with two-dimensional loading constraints. *European Journal of Operational Research*, 195(3):729–743, Jun 2009.
- [Zhang *et al.*, 2022a] Xiangyi Zhang, Lu Chen, Michel Gendreau, and André Langevin. A branch-and-cut algorithm for the vehicle routing problem with two-dimensional loading constraints. *European Journal of Operational Research*, 302(1):259–269, 2022.
- [Zhang *et al.*, 2022b] Xiangyi Zhang, Lu Chen, Michel Gendreau, and André Langevin. A branch-and-price-and-cut algorithm for the vehicle routing problem with two-dimensional loading constraints. *Transportation Science*, 56(6):1618–1635, 2022.
- [Zhang *et al.*, 2022c] Xiangyi Zhang, Lu Chen, Michel Gendreau, and André Langevin. Learning-based branch-and-price algorithms for the vehicle routing problem with time windows and two-dimensional loading constraints. *INFORMS Journal on Computing*, 34(3):1419–1436, 2022.
- [Zong *et al.*, 2021] Zefang Zong, Tao Feng, Tong Xia, Depeng Jin, and Yong Li. Deep reinforcement learning for demand driven services in logistics and transportation systems: A survey. *arXiv preprint arXiv:2108.04462*, 2021.