
A Learning-based Capacitated Arc Routing Problem Solver Comparable to Metaheuristics While with Far Less Runtimes

Anonymous Author(s)

Affiliation

Address

email

Abstract

1 Recently, neural networks (NN) have made great strides in combinatorial optimization
2 problems (COPs). However, they face challenges in solving the capacitated
3 arc routing problem (CARP) which is to find the minimum-cost tour that covers all
4 required edges on a graph, while within capacity constraints. Actually, NN-based
5 approaches tend to lag behind advanced metaheuristics due to complexities caused
6 by *non-Euclidean graph, traversal direction and capacity constraints*. In this
7 paper, we introduce an NN-based solver tailored for these complexities, which
8 significantly narrows the gap with advanced metaheuristics while with far less
9 runtimes. First, we propose the direction-aware attention model (DaAM) to in-
10 corporate directionality into the embedding process, facilitating more effective
11 one-stage decision-making. Second, we design a supervised reinforcement learning
12 scheme that involves supervised pre-training to establish a robust initial policy for
13 subsequent reinforcement fine-tuning. It proves particularly valuable for solving
14 CARP that has a higher complexity than the node routing problems (NRPs). Finally,
15 a path optimization method is introduced to adjust the depot return positions within
16 the path generated by DaAM. Experiments show that DaAM surpasses heuristics
17 and achieves decision quality comparable to state-of-the-art metaheuristics for the
18 first time while maintaining superior efficiency, even in large-scale CARP instances.
19 The code and datasets are provided in the Appendix.

20 1 Introduction

21 The capacitated arc routing problem (CARP) [7] is a combinatorial optimization problem, frequently
22 arising in domains such as inspection and search-rescue operations. Theoretically, the CARP is
23 established on an undirected connected graph $G = (\mathbf{V}, \mathbf{E}, \mathbf{E}_R)$ that includes a set of nodes \mathbf{V}
24 connected by a set of edges \mathbf{E} , and an edge subset $\mathbf{E}_R \subseteq \mathbf{E}$ that needs to be served, called required
25 edges. Each required edge has a demand value which spends the capacity of the vehicle when it is
26 served. In this context, all vehicles start their routes from the depot node $depot \in \mathbf{V}$ and conclude
27 their journey by returning to the same $depot$. The main goal of a CARP solver is to serve all required
28 edges with the lowest total path cost, while ensuring that the vehicle does not exceed its capacity Q .

29 According to [7], the CARP is recognized as an NP-Hard problem. Among all solver solutions, the
30 Memetic Algorithms (MA) [15, 25], first proposed in 2005, have still maintained unrivaled results in
31 solving the CARP challenge to this day. However, they have struggled with high time costs and the
32 exponential growth of the search space as the problem scale increases. Compared to the traditional
33 methods, NN-based solvers [16, 10, 20] are faster with the assistance of GPU, have therefore gained
34 increasing attention in recent years. However, unlike the decent performance in solving NRP, such as
35 vehicle routing problem (VRP) and travelling salesman problem (TSP), or ARP defined in Euclidean
36 graphs, such as rural postman problem (RPP) and Chinese postman problem (CPP), NN-based

37 methods usually lags far behind the traditional ones in solving CARP. This discrepancy is attributed
38 to the inability of existing methods to effectively reduce the high complexity of solving CARP:

- 39 • **Lack of edge direction in embedding learning:** ARP solvers need to determine the edges
40 to be traversed along with the direction of traversal, easy for humans to achieve in one step
41 but extremely challenging for computers. Existing methods didn't encode edge directionality
42 in embedding, making them have to build edge sequences and determine edges' directions
43 separately and leading to path generation without sufficient consideration.
- 44 • **Ineffective learning for solving CARP:** CARP is more complex than NRPs and Euclidean
45 ARPs owing to the *non-Euclidean input*, *edge direction*, and *capacity constraints*. Thus,
46 learning methods for NRPs and Euclidean ARPs cannot be directly transferred to solve
47 CARP or work well even if adapted, leaving a lack of effective learning strategies for CARP.

48 In this paper, we aim to address both above issues and propose an NN-based solver for CARP that
49 competes with the state-of-the-art MA [25] while with far less runtimes. Firstly, we propose the
50 direction-aware attention model (DaAM). It computes embeddings for directed arcs decomposed
51 from undirected edges to align with the nature of ARP, thus avoiding missing direction information
52 and enabling concise one-stage decision-making. Secondly, we design a supervised reinforcement
53 learning method to learn effective heuristics for solving CARP. DaAM is pre-trained to learn an initial
54 policy by minimizing the difference from the decisions made by experts, and fine-tuned on larger-
55 scale CARP instances by Proximal Policy Optimization with self-critical strategy. Finally, to further
56 boost the path quality, we propose a path optimizer (PO) to re-decide the vehicles' optimal return
57 positions by dynamic programming. In the experiments, our method approaches the state-of-the-art
58 MA with an average gap of 5% and is 4% better than the latest heuristics and gains high efficiency.

59 2 Related Work

60 2.1 Graph Embedding Learning

61 Graph embedding [3] aims to map nodes or edges in a graph to a low-dimensional vector space.
62 This process is commonly achieved via graph neural networks (GNNs) [31]. Kipf *et al.* [13]
63 introduced graph convolutional operations to aggregate information from neighboring nodes for
64 updating node representations. Unlike GCN, GAT [27] allowed dynamic node attention during
65 information propagation by attention mechanisms. Other GNN variants [9, 30] exhibited a similar
66 information aggregation pattern but with different computational approaches. In this paper, since
67 an arc is related to the outgoing arc of its endpoint but irrelevant to the incoming arc of that, we use
68 attention mechanisms to capture the intricate relationships between arcs for arc embedding learning.

69 2.2 Learning for Routing Problems

70 The routing problem is one of the most classic COPs, and it is mainly categorized into two types
71 according to the decision element: node routing problems and arc routing problems.

72 **Node routing problems** (NRPs), such as TSP and VRP, aim to determine the optimal paths that
73 traverse all nodes in the Euclidean space or graphs. As the solutions to these problems are context-
74 dependent sequences of variable size, they cannot be directly modeled by the Seq2Seq model [24].
75 To address this problem, Vinyals *et al.* [28] proposed the Pointer network (PN) to solving Euclidean
76 TSP, which achieves variable-size output dictionaries by neural attention. Due to the scarcity of labels
77 for supervised learning, Bello *et al.* [2] modeled the TSP as a single-step reinforcement learning
78 problem and trained the PN using policy gradient [29] within Advantage Actor-Critic (A3C) [17]
79 framework. Nazari *et al.* [19] replaced the LSTM encoder in PN with an element-wise projection
80 layer and proposed the first NN-based method to solve the Euclidean VRP and its variants. To better
81 extract correlations between inputs, Kool *et al.* [14] utilized multi-head attention for embedding
82 learning and trained the model using REINFORCE [29] with a greedy baseline. To solve COPs
83 defined on graphs, Khalil *et al.* [11] proposed S2V-DQN to learn heuristics, employing structure2vec
84 [5] for graph embedding learning and n-step DQN [18] for training. While the mentioned NN-based
85 approaches have achieved comparable performance to metaheuristics, they cannot be directly applied
86 to solve ARP due to the modeling differences between ARP and NRP.

87 **Arc routing problems** (ARPs) involve determining optimal paths for traversing arcs or edges in
88 graphs, with variants like RPP, CPP, and CARP. Truong *et al.* [26] proposed a DRL framework to

89 address the CPP with load-dependent costs on Euclidean graphs and achieved better solution quality
90 than metaheuristics. However, CARPs are defined on non-Euclidean graphs. Unlike Euclidean graphs
91 with given node coordinates, non-Euclidean graphs require manually extracting and aggregating the
92 node representations, a task that is typically learnable. Although several NN-based algorithms have
93 been proposed, they still lag significantly behind traditional methods. Li *et al.* [16] pioneered the use
94 of the NN-based approach in solving the CARP by transforming it into an NRP. They first determined
95 the sequence of edges and then decided the traversal direction for each edge. Hong *et al.* [10] trained a
96 PN in a supervised manner to select undirected edges in each time step, and also determined the edge
97 traversal direction as post-processing. Ramamoorthy *et al.* [20] proposed to generate an initial tour
98 based on edge embeddings and then split it into routes within capacity constraint. These approaches
99 lack edge directionality encoding, leading to edge selection without sufficient consideration and
100 necessitating a two-stage decision process or an additional splitting procedure.

101 3 Background

102 The attention model (AM) [14] exhibits superior effectiveness in solving classic Euclidean COPs due
103 to its attention mechanisms for extracting correlations between inputs. Therefore, we use the AM
104 as the backbone and give a brief review in terms of the TSP. Given an Euclidean graph $\mathbf{G}=(\mathbf{V}, \mathbf{E})$,
105 the AM defines a stochastic policy, denoted as $\pi(\mathbf{x}|\mathcal{S})$, where $\mathbf{x} = (x_0, \dots, x_{|\mathbf{V}|-1})$ represents a
106 permutation of the node indexes in \mathbf{V} , and \mathcal{S} is the problem instance expressing \mathbf{G} . The AM is
107 parameterized by θ as:

$$\pi_{\theta}(\mathbf{x}|\mathcal{S}) = \prod_{t=1}^{|\mathbf{V}|} \pi_{\theta}(x_t|\mathcal{S}, \mathbf{x}_{0:t-1}) \quad (1)$$

108 where t denotes the time step. Specifically, the AM comprises an encoder and a decoder. The
109 encoder first computes initial d_h -dimensional embeddings for each node in \mathbf{V} as h_i^0 through a learned
110 linear projection. It then captures the embeddings of h_i^0 using multiple attention layers, with each
111 comprising a multi-head attention (MHA) sublayer and a node-wise feed-forward (FF) sublayer.
112 Both types of sublayers include a skip connection and batch normalization (BN). Assuming that
113 $l \in \{1, \dots, N\}$ denotes the attention layer, the l^{th} layer can be formulated as h_i^l :

$$h_i^l = \text{BN}^l(\hat{h}_i + \text{FF}^l(\hat{h}_i)); \quad \hat{h}_i = \text{BN}^l(h_i^{l-1} + \text{MHA}_i^l(h_0^{l-1}, \dots, h_{|\mathbf{V}|-1}^{l-1})) \quad (2)$$

114 The decoder aims to append a node to the sequence \mathbf{x} at each time step. Specifically, a context
115 embedding $h_{(c)}$ is computed to represent the state at the time step t . Then a single attention head is
116 used to calculate the probabilities for each node based on $h_{(c)}$:

$$u_{(c)j} = \begin{cases} C \cdot \tanh\left(d_h^{-\frac{1}{2}} [\mathbf{W}^Q h_{(c)}]^T \mathbf{W}^K h_j^N\right) & \text{if } j \neq x_{t'} (\forall t' < t) \\ -\infty & \text{otherwise,} \end{cases}$$

$$p_i = \pi_{\theta}(x_t = i|\mathcal{S}, \mathbf{x}_{0:t-1}) = u_{(c)i} / \sum_j u_{(c)j} \quad (3)$$

117 where \mathbf{W}^Q and \mathbf{W}^K are the learnable parameters of the last attention layer. $u_{(c)j}$ is an unnormalized
118 log probability with (c) indicating the context node. C is a constant, and p_i is the probability
119 distribution computed by the softmax function based on $u_{(c)j}$.

120 4 Method

121 4.1 Direction-aware Attention Model

122 In this section, we propose the direction-aware attention model (DaAM). Unlike previous methods
123 that separately learn edge embeddings and determine edge directions, our model encodes direction
124 information directly into the embedding, enabling one-stage decision-making. As shown in Fig. 1,
125 the DaAM makes sequential decisions in two phases to select arcs. **The first phase** is a one-time
126 transformation process, in which the arcs of the input graph are represented as nodes in the new
127 directed complete graph. **The second phase** is executed at each time step, in which GAT is used to
128 aggregate the inter-arc weights. Subsequently, AM is used to select the arc of the next action.

129 4.1.1 Arc Feature Formulation via Graph Transformation

130 **Graph Transformation** Motivated by the need to consider direction when traversing edges, we
131 explicitly encode the edge direction by edge-to-arc decomposition. Let $\mathbf{G}=(\mathbf{V}, \mathbf{E}, \mathbf{E}_R)$ denotes

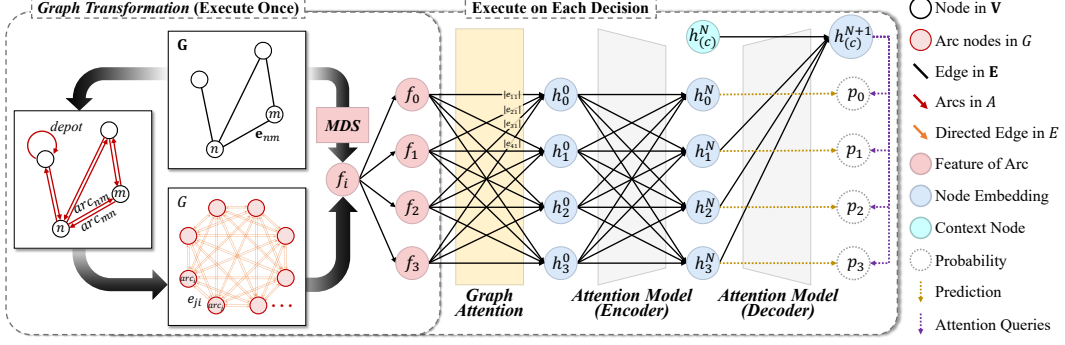


Figure 1: **DaAM Pipeline** consists of two parts. The first part transforms the input graph G by treating the arcs on G as nodes of a new directed graph G , only executing once. The second part leverages the GAT and AM to update arc embeddings and select arcs, executing at each time step.

132 the undirected connected graph as input, where V is the node set of G , E is the edge set of G ,
 133 and $E_R \subseteq E$ is the required edge set. Firstly, given that an edge has two potential traversal
 134 directions, we decompose each edge $e_{nm} = (cost_{nm}, demand_{nm}, allow_serve_{nm}) \in E_R$ into two
 135 arcs $\{arc_{nm}, arc_{mn}\}$ with opposite directions but the same cost, demand and serving state. Here
 136 n, m are the indexes of node in V . To simplify the representation below, we replace nm and mn
 137 with single-word symbols, such as i and j . In this way of edge decomposition, we obtain a set of arcs
 138 denoted as A_R . Secondly, we build a new graph $G = (A_R, E)$. Specifically, each arc in A_R serves
 139 as a node in G , and directed edge set E is created, with $e_{ij} \in E$ representing the edge from node
 140 arc_i to arc_j . The weight $|e_{ij}|$ represents the total cost of the shortest path from the end node of arc_i
 141 to the start node of arc_j . In addition, we treat the depot as a self-loop zero-demand arc that allows
 142 for repeated serving, denoted as arc_0 . Consequently, we transform the input graph G into a directed
 143 complete graph G . By decomposing all edges in E_R into arcs, it is natural to directly select the arcs
 144 from G during the decision-making. Given that the Floyd-Warshall algorithm is used to calculate the
 145 shortest path cost between any pair of nodes in G , the time complexity of our graph transformation is
 146 $\max(\mathcal{O}(|E_R|^2), \mathcal{O}(|V|^3))$.

Table 1: **Feature Detail** of arc_i at time step t for CARP.

No.	Features	Field	Description	No.	Features	Field	Description
1	is_depot_i	\mathbb{F}_2	Is arc_i the depot?	5	$allow_serve_i^{(i)}$	\mathbb{F}_2	Is arc_i at time step t allowed to serve?
2	$cost_i$	\mathbb{R}^+	Cost of arc_i .	6	$m_{start(i)}$	\mathbb{R}^d	Euclidean coordinates of arc_i 's start node.
3	$demand_i$	\mathbb{R}^+	Demand of arc_i .	7	$m_{end(i)}$	\mathbb{R}^d	Euclidean coordinates of arc_i 's end node.
4	$ e_{x_{t-1}i} $	\mathbb{R}^+	Edge weight from $arc_{x_{t-1}}$ to arc_i .				

147 **Arc Feature Formulation** To establish a foundation for decision-making regarding arc selection,
 148 the features of the arcs are constructed as input for the subsequent model. Specifically, multi-
 149 dimensional scaling (MDS) is used to project the input graph G into a d -dimensional Euclidean
 150 space. The Euclidean coordinates of arc_i 's start and end nodes, denoted as $m_{start(i)}$ and $m_{end(i)}$,
 151 are then taken as the features of arc_i to indicate its direction. As shown in Table 1, at time step t ,
 152 arc_i can be featured as:

$$F_t^{(i)} = (is_depot_i, cost_i, demand_i, |e_{x_{t-1}i}|, allow_serve_i^{(i)}, m_{start(i)}, m_{end(i)}) \quad (4)$$

153 where x_{t-1} is the index of the selected arc at the last time step, and $t \in [1, +\infty)$. Our feature models
 154 arcs rather than edges and encodes the direction attribute of arcs through MDS. Therefore, it is more
 155 suitable than previous methods [10, 16] for ARPs that need to consider the direction of traversing.

156 4.1.2 Arc Relation Encoding via Graph Attention Network

157 Although AM is efficient in decision-making, according to Eq. (2), it cannot encode the edge weights
 158 between nodes in G , an important context feature, during learning. Therefore, we use graph attention
 159 network (GAT) [27] to encode such weights. At each time step t , for each arc arc_i , we integrate the

160 weights between arc_i and all arcs in A_R along with their features into the initial embedding of arc_i .

$$c_{ij} = softmax(\alpha(\mathbf{W}[F_t^{(i)} || F_t^{(j)} || |e_{ji}|])); \quad h_i^0 = \sigma(\sum_{j=0}^{|A_R|-1} c_{ij} \mathbf{W}F_t^{(j)}) \quad (5)$$

161 where \mathbf{W} is a shared learnable parameter, $[\cdot || \cdot]$ is the horizontal concatenation operator, $\alpha(\cdot)$ is a
 162 mapping from the input to a scalar, and $\sigma(\cdot)$ denotes the activation function. h_i^0 denotes the initial
 163 feature embedding of arc_i , which is taken as the input of subsequent AM. Since G is a complete
 164 graph, we use one graph attention layer to avoid over-smoothing [4].

165 4.1.3 Arc Selection via Attention Model

166 After aggregating the edge weights of G into the initial embeddings, we utilize AM to learn the final
 167 arc embeddings and make arc selection decisions. In the encoding phase described by Eq.2, for each
 168 arc $\{arc_i\}$, we leverage N attention layers to process the initial embeddings $\{h_i^0\}$ and obtain the
 169 output embeddings of the N^{th} layer, i.e., $\{h_i^N\}$. In the decoding phase, we define the context node
 170 applicable to CARP:

$$h_{(c)}^N = \left[|A_R|^{-1} \sum_{i=0}^{|A_R|-1} h_i^N, h_{x_{t-1}}^N, \delta_t, \Delta_t \right], t \in [1, +\infty) \quad (6)$$

171 where x_{t-1} indicates the chosen arc index at time step $t-1$ and x_0 is arc_0 . δ_t is the remaining
 172 capacity at time step t , $\Delta_t = \Delta(\delta_t > \frac{Q}{2})$ is a variable to indicate whether the vehicle's remaining
 173 capacity exceeds half. Finally, according to Eq.(3), the decoder of AM takes the context node $h_{(c)}^N$
 174 and arc embeddings $\{h_i^N\}$ as inputs and calculates the probabilities for all arcs, denoted as p_i . The
 175 serviceable arc selected at time step t , i.e., arc_{x_t} , is determined by sampling or greedy decoding.

176 4.2 Supervised Reinforcement Learning for CARP

177 The decision-making of selecting arcs can be modeled as a Markov decision process with the
 178 following symbols regarding reinforcement learning:

- 179 • **State** s_t is the newest path of arcs selected from G : $(arc_{x_0}, \dots, arc_{x_{t-1}})$, while the terminal state
 180 is s_T with T indicating the final time step.
- 181 • **Action** a_t is the selected arc at time step t , i.e., arc_{x_t} . Selecting the action a_t would add arc_{x_t} to the
 182 end of the current path s_t and tag the corresponding arcs of arc_{x_t} with their features $allow_serve$
 183 changed to 0. Notably, arc_0 can be selected repeatedly but not consecutively.
- 184 • **Reward** r_t is obtained after taking action a_t at state s_t , which equals the negative shortest path
 185 cost from the last arc $arc_{x_{t-1}}$ to the selected arc arc_{x_t} .
- 186 • **Stochastic policy** $\pi(a_t|s_t)$ specifies the probability distribution over all actions at state s_t .

187 We parameterize the stochastic policy of DaAM with θ :

$$\pi(x_t | \mathcal{S}, \mathbf{x}_{0:t-1}) = \pi_\theta(a_t | s_t) \quad (7)$$

188 where \mathcal{S} is a CARP instance. Starting from initial state s_0 , we get a trajectory $\tau =$
 189 $(s_0, a_0, r_0, \dots, r_{T-1}, s_T)$ using π_θ . The goal of learning is to maximize the cumulative reward:
 190 $R(\tau) = \sum_{t=0}^{T-1} r_t$. However, due to the high complexity of CARP, vanilla deep reinforcement learn-
 191 ing methods learn feasible strategies inefficiently. A natural solution is to minimize the difference
 192 between the model's decisions and expert decisions. To achieve this, we employ supervised learning
 193 to learn an initial policy based on labeled data and then fine-tune the model through reinforcement
 194 learning.

195 4.2.1 Supervised Pre-training via Multi-class Classification

196 In the pre-training stage, we consider arc-selection at each time step as a multi-class classification
 197 task, and employ the state-of-the-art CARP method MAENS to obtain high-quality paths as the label.
 198 Assuming that $y_t \in \mathbb{R}^{|A_R|}$ denotes the one-hot label vector at time step t of any path, with $y_t^{(k)}$
 199 indicating each element. We utilize the cross-entropy loss to train the policy represented in Eq. (7):

$$L = - \sum_{t=0}^{T-1} \sum_{k=0}^{|A_R|-1} y_t^{(k)} \log(\pi_\theta(arc_k | s_t)) \quad (8)$$

200 We use the policy optimized by cross-entropy, denoted as π_s , to initialize the policy network π_θ and
 201 as the baseline policy π_b in reinforcement learning.

202 **4.2.2 Reinforcement Fine-tuning via PPO with self-critical strategy**

203 During the fine-tuning phase, we use Proximal Policy Optimization (PPO) to optimize our model
 204 $\pi_\theta(a_t|s_t)$ due to its outstanding stability in policy updates. Considering the low sample efficiency in
 205 reinforcement learning, we employ a training approach similar to self-critical training [21] to reduce
 206 gradient variance and expedite convergence. Specifically, We use another policy π_b to generate a
 207 trajectory and calculate its cumulative reward, serving as a baseline function. Our optimization
 208 objective is based on PPO-Clip [23]:

$$\mathbb{E}_{(s,a)\sim\pi_b} \left[\min \left(\frac{\pi_\theta(a|s)}{\pi_b(a|s)} (R(\tau_s^\theta) - R(\tau_s^b)), \text{clip} \left(\frac{\pi_\theta(a|s)}{\pi_b(a|s)}, 1-\epsilon, 1+\epsilon \right) (R(\tau_s^\theta) - R(\tau_s^b)) \right) \right] \quad (9)$$

209 where s is used to replace current state s_t for symbol simplification, and a for a_t . $\text{clip}(w, v_{\min}, v_{\max})$
 210 denotes constraining w within the range $[v_{\min}, v_{\max}]$, and ϵ is a hyper-parameter. τ_s^θ denotes a
 211 trajectory sampled by π_θ with s as the initial state, while τ_s^b for the trajectory greedily decoded by π_b .
 212 In greedy decoding, the action with the maximum probability is selected at each step. $R(\tau_s^\theta) - R(\tau_s^b)$
 213 serves as an advantage measure, quantifying the advantage of the current policy π_θ compared to π_b .
 214 We maximize Eq. (9) through gradient descent, which forces the model to select actions that yield
 215 higher advantages. The baseline policy’s parameters are updated if π_θ outperforms π_b .

216 **4.3 Path Optimization via Dynamic Programming**

217 The complexity of the problem is heightened by the increasing capacity constraint, making it challeng-
 218 ing for the neural network to make accurate decisions regarding the depot return positions. In this sec-
 219 tion, we propose a dynamic programming (DP) based strategy to assist our model in optimizing these
 220 positions. Assuming that \mathbf{P} is assigned with the terminal state $s_T = (arc_{x_0}, arc_{x_1}, \dots, arc_{x_{T-1}})$,
 221 representing a generated path. Initially, we remove all the depot arcs in \mathbf{P} to obtain a new
 222 path $\mathbf{P}' = (arc_{x'_0}, arc_{x'_1}, \dots, arc_{x'_{T'-1}})$, where $\{x'_i | i \in [0, T' - 1]\}$ denotes a subsequence of
 223 $\{x_i | i \in [0, T - 1]\}$. Subsequently, we aim to insert several new depot arcs into the path \mathbf{P}'
 224 to achieve a lower cost while adhering to capacity constraints. To be specific, we recursively find the
 225 return point that minimizes the overall increasing cost, which is implemented by the state transition
 226 equation as follows:

$$\begin{aligned} f(\mathbf{P}') &= \min_i (f(\mathbf{P}'_{0:i}) + SC(arc_{x'_i}, arc_0) + SC(arc_0, arc_{x'_{i+1}}) - SC(arc_{x'_i}, arc_{x'_{i+1}})) \\ \text{s.t. } & 0 \leq i < T' - 1, \quad \sum_{j=i+1}^{T'-1} demand_{x'_j} \leq Q \end{aligned} \quad (10)$$

227 where $SC(arc_{x'_i}, arc_0) = |e_{x'_i 0}|$ denotes the shortest path cost from $arc_{x'_i}$ to the depot. Q is the
 228 vehicle capacity. According to Eq. (10), we insert the depot arc arc_0 after an appropriate position
 229 $arc_{x'_i}$, which meets with the capacity constraint of the subpath $\mathbf{P}'_{i+1:T'-1}$. $f(\cdot)$ denotes a state
 230 featuring dynamic programming. By enumerating the position i , we compute the minimum increasing
 231 cost $f(\mathbf{P}')$ utilizing its sub-state $f(\mathbf{P}'_{0:i})$. The final minimum cost for path \mathbf{P} is $f(\mathbf{P}') + g(\mathbf{P}')$, here
 232 $g(\mathbf{P}')$ is the unoptimized cost of \mathbf{P}' . Since \mathbf{P}' includes only the required edges, i.e., $T' = |\mathbf{E}_R|$,
 233 the time complexity of DP is $\mathcal{O}(|\mathbf{E}_R|^2)$. During Path Optimization, we use beam search to generate
 234 two paths with the trained policy, one with capacity-constrained and one without. Both paths are
 235 optimized using DP and the one with the minimum cost is selected as the final result.

236 **5 Experiments**

237 **5.1 Setup**

238 **Problem Instances.** We extracted sub-graphs from the roadmap of Beijing, China, obtained from
 239 OpenStreetMap [8], to create CARP instances for both the training and testing phases. All instances
 240 are divided into seven datasets, each representing different problem scales, as presented in Table 2.
 241 Each dataset consists of 30,000 instances, further divided into two **disjoint** subsets: 20,000 instances
 242 for training and the remaining for testing. For each instance, the vehicle capacity is set to 100.

243 **Implementation Details.** Our neural network is implemented using the PyTorch framework and
 244 trained on a single NVIDIA RTX 3090 GPU. The heuristics and metaheuristics algorithms are

Table 2: **Datasets information.** $|V|$ is the number of nodes, $|E_R|$ is the number of required edges. *demand* represents the demand range for each required edge. Each dataset has 20,000 training instances and 10,000 test instances.

CARP instances	$ V $	$ E_R $	<i>demand</i>	CARP instances	$ V $	$ E_R $	<i>demand</i>
Task 20	25-30	20	5-10	Task 200	205-210	200	1000
Task 40	45-50	40	5-10	Task 300	305-310	300	1000
Task 60	65-70	60	5-10	Task 400	405-410	400	1000
Task 80	85-90	80	5-10	Task 500	505-510	500	1000
Task 100	105-110	100	5-10	Task 600	605-610	600	1000

Table 3: **Solution quality comparison.** All methods are evaluated on 10,000 CARP instances in each scale. We measure the gap (%) between different methods and MAENS. Methods marked with an asterisk were originally proposed for NRP, but we modified them to solve CARP. The best results are indicated in bold, while the second-best results are underlined.

Method	Task20		Task40		Task60		Task80		Task100	
	Cost	Gap (%)	Cost	Gap (%)	Cost	Gap (%)	Cost	Gap (%)	Cost	Gap (%)
MAENS [25]	474	0	950	0	1529	0	2113	0	2757	0
PS [6]	544	14.72	1079	13.56	1879	22.84	2504	18.49	3361	21.90
PS-Ellipse [22]	519	9.49	1006	5.89	1709	11.77	2299	8.80	3095	12.26
PS-Efficiency [1]	514	8.44	1007	<u>6.00</u>	1684	10.14	2282	8.00	3056	10.85
PS-Alt1 [1]	514	8.44	1007	<u>6.00</u>	1685	10.20	2283	8.04	3057	10.88
PS-Alt2 [1]	521	9.92	1009	6.21	1720	12.49	2314	9.51	3102	12.51
S2V-DQN* [11]	590	24.42	1197	26.02	1900	24.23	2820	33.43	3404	23.42
VRP-DL* [19]	528	11.39	1193	25.57	2033	32.96	2898	37.15	3867	40.26
DaAM (SL)	509	7.43	1066	12.24	-	-	-	-	-	-
DaAM (SL+RL)	495	4.48	1009	6.19	1639	<u>7.16</u>	2275	<u>7.67</u>	2980	8.06
DaAM (SL+RL+PO)	482	1.65	992	4.39	1621	5.98	2255	6.70	2958	7.28

245 evaluated on an Intel Core i9-7920X with 24 cores and a CPU frequency of 4.4GHz. We optimize
 246 the model using Adam optimizer [12]. The dimension of MDS coordinates d is set to 8, and the
 247 learning rate is set to $1e^{-4}$. We set ϵ in the PPO training at 0.1. Notably, our PPO training does not
 248 incorporate discounted cumulative rewards, i.e., γ is set to 1.

249 **Metrics and Settings.** For each method and dataset, We compute the mean tour cost across all test
 250 instances, indicated by ‘‘Cost’’. Employing the state-of-the-art MAENS [25] as a baseline, we measure
 251 the ‘‘Cost’’ gap between alternative algorithms and MAENS, indicated by ‘‘Gap’’. We compare our
 252 method against the heuristic Path-Scanning algorithms (PS) [6, 22, 1] and two NN-based algorithms.
 253 In the absence of publicly available code for prior NN-based CARP methods, we modify two NN-
 254 based NRP solvers to suit CARP, i.e, S2V-DQN [11] and VRP-DL [19]. Note that, for S2V-DQN,
 255 we replace `structure2vec` with GAT to achieve more effective graph embedding learning. For our
 256 method, we incrementally add supervised pre-training (SL), reinforcement learning fine-tuning (RL),
 257 and path optimization (PO) to assess the effectiveness of our training scheme and optimization,
 258 respectively. Due to the excessively long computation times of MAENS on larger-scale datasets, SL
 259 is only performed on Task20, Task 30, and Task40. The batch size for SL is set to 128. During the RL
 260 stage, greedy decoding is used to generate solutions, and except for the Task20 dataset, we utilize the
 261 training results obtained from the preceding smaller-scale dataset to initialize the model. The beam
 262 width in the PO stage is set to 2. For each dataset, we compare the mean cost of different methods on
 263 10,000 problem instances.

264 5.2 Evaluation Results

265 **Solution Quality** Table 6 shows the result. Our algorithm outperforms all heuristic and NN-based
 266 methods across all scales, achieving costs comparable to MAENS, trailing by less than 8%. The
 267 advantage over PS demonstrates that neural networks can learn more effective policies than hand-
 268 crafted ones, attributed to our well-designed modeling approach. Moreover, as the problem scale
 269 increases, it becomes time-consuming to obtain CARP annotation by MAENS. Therefore, we leverage
 270 the model pre-trained on small-scale instances as the initial policy for RL fine-tuning on Task50,
 271 Task60, Task80, and Task100, yielding commendable performance. This proves the generalization of
 272 our training scheme across varying problem scales. The performance gap with MAENS highlights
 273 our algorithm’s superiority in CARP-solving approaches.

Table 4: **Generalization to larger problem instances.** All methods are evaluated on 10,000 CARP instances in each scale. For DaAM, we employ the policy trained on Task100. The best results are indicated in bold, while the second-best results are underlined.

Method	Task200 Cost	Task300 Cost	Task400 Cost	Task500 Cost	Task600 Cost
PS-Ellipse [22]	4240	6563	8600	10909	13377
PS-Efficiency [1]	4233	6544	8583	<u>10883</u>	<u>13338</u>
PS-Alt1 [1]	4233	6544	<u>8580</u>	10884	13338
PS-Alt2 [1]	4244	6569	8606	10922	13393
DaAM (SL+RL)	<u>4189</u>	<u>6372</u>	8610	10938	13340
DaAM (SL+RL+PO)	4132	6281	8473	10633	13100




Figure 2: **Run time comparison.** For each dataset, the total run time of each method on 100 CARP instances is shown.

274 **Generalization Ability.** In Table 4, we assess DaAM’s generalization on large-scale CARP in-
 275 stances using the policy trained on Task100. We remove MAENS and PS due to failing to run on
 276 large-scale graphs, and remove S2V-DQN and VRP-DL due to poor performance. Although DaAM
 277 is not trained on large-scale instances, it achieves or even exceeds the performance of PS, which
 278 shows its potential application on larger-scale CARP instances.
 279

280 **Run Time.** We compare the total time required for solving 100 CARP instances across datasets
 281 Task20 to Task100 datasets using our method, MAENS, and PS algorithms, and show the run time in
 282 log space. For datasets Task200 to Task600, we compare the same metric using variants of PS and
 283 out method. For our method, we measured the solving time with and without PO. Fig. 2 demonstrates
 284 that our method exhibits a significant speed advantage over MAENS, even outperforming variants of
 285 PS [1] on most datasets. In comparison, the consumption time of MAENS increases exponentially as
 286 the problem scale increases. Our method efficiently generates paths for large-scale CARP instances
 287 by leveraging GPU data-level parallelism and CPU instruction-level parallelism.

288 **Effectiveness of Combining MDS and GAT.** Table 5: **Costs of DaAM** using different encoder.

Method	Task30	Task40	Task50	Task60
MDS	743	1017	1338	1699
GAT	746	1019	1317	1684
MDS + GAT	741	1011	1322	1683

289 To evaluate the combination of MDS and GAT
 290 for embedding exhibiting, we individually evalu-
 291 ate the performance of models using only MDS
 292 or GAT, as well as their combined performance.
 293 The experiment is conducted on Task30, Task40,
 294 Task50, and Task60 by comparing the average performance of 1,000 instances on each dataset. In the
 295 RL stage, we use the policy pre-trained on Task30 for initialization. Table 5 indicates that using MDS
 296 or GAT individually yields worse quality in most cases, highlighting that combining MDS and GAT
 297 enhances the model’s capacity to capture arc correlations. Fig. 3 depicts the convergence trends in
 298 these scenes, which shows that the synergy of MDS and GAT contributes to the stability of training.

299 **Solution Visualization.** For a more intuitive understanding of the paths generated by different
 300 methods, we visualize and compare the results of our method with PS [6] and MAENS across four
 301 road scenes in Beijing. Fig. 4 visualizes all results alongside scene information. We observe that our
 302 model obtains similar paths with MAENS since we leverage the annotation generated by MAENS for




Figure 3: **Convergence trends** of different methods in reinforcement learning training.




Figure 4: **Qualitative comparison** in four real street scenes. The paths are marked in different colors, with gray indicating roads that do not require service and red points indicating depots.

303 supervised learning. MAENS paths exhibit superior spatial locality, clearly dividing the scene into
 304 regions, whereas PS paths appear more random.

305 6 Conclusion and Limitations

306 In this paper, we propose a learning-based CARP solver that competes with state-of-the-art meta-
 307 heuristics. Firstly, we encode the potential serving direction of edges into embeddings, ensuring
 308 that edge directionality is taken into account in decision-making. Secondly, we present a supervised
 309 reinforcement learning approach that effectively learns policies to solve CARP. With the aid of
 310 these contributions, our method surpasses all heuristics and achieves performance comparable to
 311 metaheuristics for the first time while maintaining excellent efficiency.

312 **Limitations and future work.** Decomposing undirected edges increases the decision elements,
 313 which complicates the problem and may widens the gap between DaAM and traditional state-of-the-
 314 art approaches as the problem instance scale increases. Our future work focuses on designing an
 315 efficient graph transformation method that does not significantly increase problem complexity.

References

- 316
- 317 [1] Rafael Kendy Arakaki and Fabio Luiz Usberti. An efficiency-based path-scanning heuristic for
318 the capacitated arc routing problem. *Computers & Operations Research (COR)*, 103:288–295,
319 2019.
- 320 [2] Irwan Bello*, Hieu Pham*, Quoc V. Le, Mohammad Norouzi, and Samy Bengio. Neural com-
321 binatorial optimization with reinforcement learning. In *International Conference on Learning*
322 *Representations (ICLR)*, 2017.
- 323 [3] Hongyun Cai, Vincent W Zheng, and Kevin Chen-Chuan Chang. A comprehensive survey of
324 graph embedding: Problems, techniques, and applications. *IEEE transactions on knowledge*
325 *and data engineering (TKDE)*, 30(9):1616–1637, 2018.
- 326 [4] Deli Chen, Yankai Lin, Wei Li, Peng Li, Jie Zhou, and Xu Sun. Measuring and relieving
327 the over-smoothing problem for graph neural networks from the topological view. In *AAAI*
328 *conference on artificial intelligence (AAAI)*, 2020.
- 329 [5] Hanjun Dai, Bo Dai, and Le Song. Discriminative embeddings of latent variable models for
330 structured data. In *International conference on machine learning (ICML)*, 2016.
- 331 [6] Bruce L Golden, James S DeArmon, and Edward K Baker. Computational experiments with
332 algorithms for a class of routing problems. *Computers & Operations Research (COR)*, 10(1):47–
333 59, 1983.
- 334 [7] Bruce L Golden and Richard T Wong. Capacitated arc routing problems. *Networks*, 11(3):305–
335 315, 1981.
- 336 [8] Mordechai Haklay and Patrick Weber. Openstreetmap: User-generated street maps. *IEEE*
337 *Pervasive computing*, 7(4):12–18, 2008.
- 338 [9] Will Hamilton, Zhitao Ying, and Jure Leskovec. Inductive representation learning on large
339 graphs. In *Advances in neural information processing systems (NeurIPS)*, 2017.
- 340 [10] Wenjing Hong and Tonglin Liu. Faster capacitated arc routing: A sequence-to-sequence
341 approach. *IEEE Access*, 10:4777–4785, 2022.
- 342 [11] Elias Khalil, Hanjun Dai, Yuyu Zhang, Bistra Dilkina, and Le Song. Learning combinatorial
343 optimization algorithms over graphs. In *Advances in neural information processing systems*
344 *(NeurIPS)*, 2017.
- 345 [12] Diederik Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *International*
346 *Conference on Learning Representations (ICLR)*, 2015.
- 347 [13] Thomas N. Kipf and Max Welling. Semi-supervised classification with graph convolutional
348 networks. In *International Conference on Learning Representations (ICLR)*, 2017.
- 349 [14] Wouter Kool, Herke van Hoof, and Max Welling. Attention, learn to solve routing problems! In
350 *International Conference on Learning Representations (ICLR)*, 2019.
- 351 [15] Natalio Krasnogor and James Smith. A tutorial for competent memetic algorithms: model, taxon-
352 omy, and design issues. *IEEE transactions on Evolutionary Computation (TEVC)*, 9(5):474–488,
353 2005.
- 354 [16] Han Li and Guiying Li. Learning to solve capacitated arc routing problems by policy gradient.
355 In *IEEE Congress on Evolutionary Computation (CEC)*, 2019.
- 356 [17] Volodymyr Mnih, Adria Puigdomenech Badia, Mehdi Mirza, Alex Graves, Timothy Lillicrap,
357 Tim Harley, David Silver, and Koray Kavukcuoglu. Asynchronous methods for deep reinforce-
358 ment learning. In *International conference on machine learning (ICML)*, pages 1928–1937,
359 2016.
- 360 [18] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G
361 Bellemare, Alex Graves, Martin Riedmiller, Andreas K Fidjeland, Georg Ostrovski, et al.
362 Human-level control through deep reinforcement learning. *Nature*, 518(7540):529–533, 2015.

- 363 [19] Mohammadreza Nazari, Afshin Oroojlooy, Lawrence Snyder, and Martin Takáč. Reinforcement
364 learning for solving the vehicle routing problem. In *Advances in neural information processing*
365 *systems (NeurIPS)*, 2018.
- 366 [20] Muhilan Ramamoorthy and Violet R. Syrotiuk. Learning heuristics for arc routing problems.
367 *Intelligent Systems with Applications (ISWA)*, 21:200300, 2024.
- 368 [21] Steven J Rennie, Etienne Marcheret, Youssef Mroueh, Jerret Ross, and Vaibhava Goel. Self-
369 critical sequence training for image captioning. In *IEEE conference on computer vision and*
370 *pattern recognition (CVPR)*, 2017.
- 371 [22] Luís Santos, João Coutinho-Rodrigues, and John R Current. An improved heuristic for the
372 capacitated arc routing problem. *Computers & Operations Research (COR)*, 36(9):2632–2637,
373 2009.
- 374 [23] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal
375 policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- 376 [24] Ilya Sutskever, Oriol Vinyals, and Quoc V Le. Sequence to sequence learning with neural
377 networks. In *Advances in neural information processing systems (NeurIPS)*, 2014.
- 378 [25] Ke Tang, Yi Mei, and Xin Yao. Memetic algorithm with extended neighborhood search for
379 capacitated arc routing problems. *IEEE Transactions on Evolutionary Computation (TEVC)*,
380 13(5):1151–1166, 2009.
- 381 [26] Cong Dao Tran and Truong Son Hy. Graph attention-based deep reinforcement learning for solv-
382 ing the chinese postman problem with load-dependent costs. *arXiv preprint arXiv:2310.15516*,
383 2023.
- 384 [27] Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Liò, and Yoshua
385 Bengio. Graph Attention Networks. In *International Conference on Learning Representations*
386 *(ICLR)*, 2018.
- 387 [28] Oriol Vinyals, Meire Fortunato, and Navdeep Jaitly. Pointer networks. In *Advances in neural*
388 *information processing systems (NeurIPS)*, 2015.
- 389 [29] Ronald J Williams. Simple statistical gradient-following algorithms for connectionist reinforce-
390 ment learning. *Machine learning*, 8:229–256, 1992.
- 391 [30] Felix Wu, Amauri Souza, Tianyi Zhang, Christopher Fifty, Tao Yu, and Kilian Weinberger.
392 Simplifying graph convolutional networks. In *International conference on machine learning*
393 *(ICML)*, 2019.
- 394 [31] Zonghan Wu, Shirui Pan, Fengwen Chen, Guodong Long, Chengqi Zhang, and S Yu Philip. A
395 comprehensive survey on graph neural networks. *IEEE transactions on neural networks and*
396 *learning systems (TNNLS)*, 32(1):4–24, 2020.

397 A Appendix

398 A.1 Source Code and Dataset

399 The source code of DaAM and the datasets used for testing are available at DaAM. Once the paper is
400 accepted, we will promptly release the source code and datasets.

401 A.2 Pseudocode of PPO with self-critical strategy

402 Algorithm 1 presents the pseudocode for the PPO training algorithm we used. In the code implemen-
403 tation, the trajectory τ_s^θ can be replaced by (s, a) 's original trajectory τ_o for efficiency. Once τ_o is
404 sampled, the cumulative rewards from any state $s \in \tau_o$ can be quickly computed.

Algorithm 1 PPO algorithm with self-critical strategy

Input: batch size B , number of episodes K , train instances \mathcal{P} , test instances \mathcal{T}

Initialize policies $\pi_\theta, \pi_b \leftarrow \pi_s$

```

1: for episode  $k = 1$  to  $K$  do
2:   Initialize data batch  $\mathcal{M}, \mathcal{M}' \leftarrow ()$ 
3:   while  $|\mathcal{M}| < B$  do
4:     Sample a CARP instance  $\mathcal{S}$  from  $\mathcal{P}$ 
5:     Sample  $\tau_o = (s_0, a_0, \dots, s_T)$  from  $\mathcal{S}$  using  $\pi_b$ 
6:      $\mathcal{M} \leftarrow \mathcal{M} \cup \{(s_0, a_0), \dots, (s_{T-1}, a_{T-1})\}$ 
7:   end while
8:   for each  $(s, a) \in \mathcal{M}$  do
9:     Generate trajectory  $\tau_s^\theta$  using  $\pi_\theta$  from  $s$  by sampling
10:    Generate trajectory  $\tau_s^b$  using  $\pi_b$  from  $s$  by greedy decoding
11:    Compute advantage  $\mathcal{A}_s = R(\tau_s^\theta) - R(\tau_s^b)$ 
12:     $\mathcal{M}' \leftarrow \mathcal{M}' \cup \{(s, a, \mathcal{A}_s)\}$ 
13:   end for
14:   Update  $\pi_\theta$  using Adam over (9) based on  $\mathcal{M}'$ 
15:   if  $\pi_\theta$  outperforms  $\pi_b$  on  $\mathcal{T}$  then
16:      $\pi_b \leftarrow \pi_\theta$ 
17:   end if
18: end for

```

405 A.3 Experimental Results of Additional Datasets

406 For small-scale problem instances, we generated two additional datasets, Task30 and Task50. In
407 Task30 the range of $|V|$ is 25-30, while in Task50, it spans 55-60. Correspondingly, $|\mathbf{E}_R|$ is set to
408 30 and 50, respectively. The demand for each edge ranges from 5 to 10 in both tasks. Table 6 is the
complete experimental data from the solution quality experiments.

Table 6: **Solution quality comparison.** All methods are evaluated on 10,000 CARP instances in each scale. We measure the gap (%) between different methods and MAENS. Methods marked with an asterisk were originally proposed for NRP, but we modified them to solve CARP. The gray indicates that MAENS is taken as the baseline when calculating ‘‘Gap’’. The best results are indicated in bold, while the second-best results are underlined.

Method	Task20		Task30		Task40		Task50		Task60		Task80		Task100	
	Cost	Gap (%)	Cost	Gap (%)	Cost	Gap (%)	Cost	Gap (%)	Cost	Gap (%)	Cost	Gap (%)	Cost	Gap (%)
MAENS [25]	474	0.00	706	0.00	950	0.00	1222	0.00	1529	0.00	2113	0.00	2757	0.00
PS [6]	544	14.72	859	21.76	1079	13.56	1448	18.45	1879	22.84	2504	18.49	3361	21.90
PS-Ellipse [22]	519	9.49	798	13.03	1006	5.89	1328	8.67	1709	11.77	2299	8.80	3095	12.26
PS-Efficiency [1]	514	8.44	790	11.90	1007	<u>6.00</u>	1311	7.28	1684	10.14	2282	8.00	3056	10.85
PS-Alt1 [1]	514	8.44	791	12.04	1007	<u>6.00</u>	1312	7.36	1685	10.20	2283	8.04	3057	10.88
PS-Alt2 [1]	521	9.92	802	13.60	1009	6.21	1336	9.33	1720	12.49	2314	9.51	3102	12.51
S2V-DQN* [11]	590	24.42	880	24.65	1197	26.02	1520	24.32	1900	24.23	2820	33.43	3404	23.42
VRP-DL* [19]	528	11.39	848	20.11	1193	25.57	1587	29.87	2033	32.96	2898	37.15	3867	40.26
DaAM (SL)	509	7.43	785	11.18	1066	12.24	-	-	-	-	-	-	-	-
DaAM (SL+RL)	495	4.48	741	<u>5.05</u>	1009	6.19	1303	<u>6.58</u>	1639	<u>7.16</u>	2275	<u>7.67</u>	2980	<u>8.06</u>
DaAM (SL+RL+PO)	482	1.65	725	2.73	992	4.39	1283	5.07	1621	5.98	2255	6.70	2958	7.28

410 **A.4 Licences of Assets Used for Experiments**

411 The code we used does not require special consent from the authors. We follow their licenses as
412 specified below:

- 413 • <https://github.com/wouterkool/attention-learn-to-route>: MIT Licence.
- 414 • https://github.com/Hanjun-Dai/graph_comb_opt: MIT Licence.

415 **NeurIPS Paper Checklist**

416 **1. Claims**

417 Question: Do the main claims made in the abstract and introduction accurately reflect the
418 paper's contributions and scope?

419 Answer: [\[Yes\]](#)

420 Justification: See Abstract and Introduction.

421 Guidelines:

- 422 • The answer NA means that the abstract and introduction do not include the claims
423 made in the paper.
- 424 • The abstract and/or introduction should clearly state the claims made, including the
425 contributions made in the paper and important assumptions and limitations. A No or
426 NA answer to this question will not be perceived well by the reviewers.
- 427 • The claims made should match theoretical and experimental results, and reflect how
428 much the results can be expected to generalize to other settings.
- 429 • It is fine to include aspirational goals as motivation as long as it is clear that these goals
430 are not attained by the paper.

431 **2. Limitations**

432 Question: Does the paper discuss the limitations of the work performed by the authors?

433 Answer: [\[Yes\]](#)

434 Justification: See section **Conclusion and Limitations**.

435 Guidelines:

- 436 • The answer NA means that the paper has no limitation while the answer No means that
437 the paper has limitations, but those are not discussed in the paper.
- 438 • The authors are encouraged to create a separate "Limitations" section in their paper.
- 439 • The paper should point out any strong assumptions and how robust the results are to
440 violations of these assumptions (e.g., independence assumptions, noiseless settings,
441 model well-specification, asymptotic approximations only holding locally). The authors
442 should reflect on how these assumptions might be violated in practice and what the
443 implications would be.
- 444 • The authors should reflect on the scope of the claims made, e.g., if the approach was
445 only tested on a few datasets or with a few runs. In general, empirical results often
446 depend on implicit assumptions, which should be articulated.
- 447 • The authors should reflect on the factors that influence the performance of the approach.
448 For example, a facial recognition algorithm may perform poorly when image resolution
449 is low or images are taken in low lighting. Or a speech-to-text system might not be
450 used reliably to provide closed captions for online lectures because it fails to handle
451 technical jargon.
- 452 • The authors should discuss the computational efficiency of the proposed algorithms
453 and how they scale with dataset size.
- 454 • If applicable, the authors should discuss possible limitations of their approach to
455 address problems of privacy and fairness.
- 456 • While the authors might fear that complete honesty about limitations might be used by
457 reviewers as grounds for rejection, a worse outcome might be that reviewers discover
458 limitations that aren't acknowledged in the paper. The authors should use their best
459 judgment and recognize that individual actions in favor of transparency play an impor-
460 tant role in developing norms that preserve the integrity of the community. Reviewers
461 will be specifically instructed to not penalize honesty concerning limitations.

462 **3. Theory Assumptions and Proofs**

463 Question: For each theoretical result, does the paper provide the full set of assumptions and
464 a complete (and correct) proof?

465 Answer: [\[NA\]](#)

466 Justification: This paper does not include theoretical results.

467 Guidelines:

- 468 • The answer NA means that the paper does not include theoretical results.
- 469 • All the theorems, formulas, and proofs in the paper should be numbered and cross-
470 referenced.
- 471 • All assumptions should be clearly stated or referenced in the statement of any theorems.
- 472 • The proofs can either appear in the main paper or the supplemental material, but if
473 they appear in the supplemental material, the authors are encouraged to provide a short
474 proof sketch to provide intuition.
- 475 • Inversely, any informal proof provided in the core of the paper should be complemented
476 by formal proofs provided in appendix or supplemental material.
- 477 • Theorems and Lemmas that the proof relies upon should be properly referenced.

478 4. Experimental Result Reproducibility

479 Question: Does the paper fully disclose all the information needed to reproduce the main ex-
480 perimental results of the paper to the extent that it affects the main claims and/or conclusions
481 of the paper (regardless of whether the code and data are provided or not)?

482 Answer: [Yes]

483 Justification: This paper discusses the detail to reproduce the main experimental results of
484 the paper.

485 Guidelines:

- 486 • The answer NA means that the paper does not include experiments.
- 487 • If the paper includes experiments, a No answer to this question will not be perceived
488 well by the reviewers: Making the paper reproducible is important, regardless of
489 whether the code and data are provided or not.
- 490 • If the contribution is a dataset and/or model, the authors should describe the steps taken
491 to make their results reproducible or verifiable.
- 492 • Depending on the contribution, reproducibility can be accomplished in various ways.
493 For example, if the contribution is a novel architecture, describing the architecture fully
494 might suffice, or if the contribution is a specific model and empirical evaluation, it may
495 be necessary to either make it possible for others to replicate the model with the same
496 dataset, or provide access to the model. In general, releasing code and data is often
497 one good way to accomplish this, but reproducibility can also be provided via detailed
498 instructions for how to replicate the results, access to a hosted model (e.g., in the case
499 of a large language model), releasing of a model checkpoint, or other means that are
500 appropriate to the research performed.
- 501 • While NeurIPS does not require releasing code, the conference does require all submis-
502 sions to provide some reasonable avenue for reproducibility, which may depend on the
503 nature of the contribution. For example
 - 504 (a) If the contribution is primarily a new algorithm, the paper should make it clear how
505 to reproduce that algorithm.
 - 506 (b) If the contribution is primarily a new model architecture, the paper should describe
507 the architecture clearly and fully.
 - 508 (c) If the contribution is a new model (e.g., a large language model), then there should
509 either be a way to access this model for reproducing the results or a way to reproduce
510 the model (e.g., with an open-source dataset or instructions for how to construct
511 the dataset).
 - 512 (d) We recognize that reproducibility may be tricky in some cases, in which case
513 authors are welcome to describe the particular way they provide for reproducibility.
514 In the case of closed-source models, it may be that access to the model is limited in
515 some way (e.g., to registered users), but it should be possible for other researchers
516 to have some path to reproducing or verifying the results.

517 5. Open access to data and code

518 Question: Does the paper provide open access to the data and code, with sufficient instruc-
519 tions to faithfully reproduce the main experimental results, as described in supplemental
520 material?

521
522
523
524
525
526
527
528
529
530
531
532
533
534
535
536
537
538
539
540
541
542
543
544
545
546
547
548
549
550
551
552
553
554
555
556
557
558
559
560
561
562
563
564
565
566
567
568
569
570
571
572

Answer: [Yes]

Justification: The source code and datasets are provided in the Appendix.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental Setting/Details

Question: Does the paper specify all the training and test details (e.g., data splits, hyper-parameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: See section **Experiments**.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. Experiment Statistical Significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [No]

Justification: Since the experimental results are deterministic, we did not repeat the experiments multiple times. However, to reduce errors, we calculated the average over 10,000 problem instances for each dataset of any scale.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).

- 573 • It should be clear whether the error bar is the standard deviation or the standard error
574 of the mean.
- 575 • It is OK to report 1-sigma error bars, but one should state it. The authors should
576 preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis
577 of Normality of errors is not verified.
- 578 • For asymmetric distributions, the authors should be careful not to show in tables or
579 figures symmetric error bars that would yield results that are out of range (e.g. negative
580 error rates).
- 581 • If error bars are reported in tables or plots, The authors should explain in the text how
582 they were calculated and reference the corresponding figures or tables in the text.

583 8. Experiments Compute Resources

584 Question: For each experiment, does the paper provide sufficient information on the com-
585 puter resources (type of compute workers, memory, time of execution) needed to reproduce
586 the experiments?

587 Answer: [Yes]

588 Justification: See section **Experiments**.

589 Guidelines:

- 590 • The answer NA means that the paper does not include experiments.
- 591 • The paper should indicate the type of compute workers CPU or GPU, internal cluster,
592 or cloud provider, including relevant memory and storage.
- 593 • The paper should provide the amount of compute required for each of the individual
594 experimental runs as well as estimate the total compute.
- 595 • The paper should disclose whether the full research project required more compute
596 than the experiments reported in the paper (e.g., preliminary or failed experiments that
597 didn't make it into the paper).

598 9. Code Of Ethics

599 Question: Does the research conducted in the paper conform, in every respect, with the
600 NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines?>

601 Answer: [Yes]

602 Justification:

603 Guidelines:

- 604 • The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- 605 • If the authors answer No, they should explain the special circumstances that require a
606 deviation from the Code of Ethics.
- 607 • The authors should make sure to preserve anonymity (e.g., if there is a special consid-
608 eration due to laws or regulations in their jurisdiction).

609 10. Broader Impacts

610 Question: Does the paper discuss both potential positive societal impacts and negative
611 societal impacts of the work performed?

612 Answer: [No]

613 Justification: The aim of our work is to provide a better solution for a class of combinatorial
614 optimization problems, though it is difficult to predict its impact on society.

615 Guidelines:

- 616 • The answer NA means that there is no societal impact of the work performed.
- 617 • If the authors answer NA or No, they should explain why their work has no societal
618 impact or why the paper does not address societal impact.
- 619 • Examples of negative societal impacts include potential malicious or unintended uses
620 (e.g., disinformation, generating fake profiles, surveillance), fairness considerations
621 (e.g., deployment of technologies that could make decisions that unfairly impact specific
622 groups), privacy considerations, and security considerations.

- 623
- 624
- 625
- 626
- 627
- 628
- 629
- 630
- 631
- 632
- 633
- 634
- 635
- 636
- 637
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
 - The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
 - If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

638 11. Safeguards

639 Question: Does the paper describe safeguards that have been put in place for responsible
640 release of data or models that have a high risk for misuse (e.g., pretrained language models,
641 image generators, or scraped datasets)?

642 Answer: [NA]

643 Justification:

644 Guidelines:

- 645
- 646
- 647
- 648
- 649
- 650
- 651
- 652
- 653
- 654
- The answer NA means that the paper poses no such risks.
 - Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
 - Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
 - We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

655 12. Licenses for existing assets

656 Question: Are the creators or original owners of assets (e.g., code, data, models), used in
657 the paper, properly credited and are the license and terms of use explicitly mentioned and
658 properly respected?

659 Answer: [Yes]

660 Justification: See the **Appendix**.

661 Guidelines:

- 662
- 663
- 664
- 665
- 666
- 667
- 668
- 669
- 670
- 671
- 672
- 673
- 674
- The answer NA means that the paper does not use existing assets.
 - The authors should cite the original paper that produced the code package or dataset.
 - The authors should state which version of the asset is used and, if possible, include a URL.
 - The name of the license (e.g., CC-BY 4.0) should be included for each asset.
 - For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
 - If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
 - For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.

675 • If this information is not available online, the authors are encouraged to reach out to
676 the asset’s creators.

677 **13. New Assets**

678 Question: Are new assets introduced in the paper well documented and is the documentation
679 provided alongside the assets?

680 Answer: [Yes]

681 Justification: See the **Appendix**.

682 Guidelines:

- 683 • The answer NA means that the paper does not release new assets.
- 684 • Researchers should communicate the details of the dataset/code/model as part of their
685 submissions via structured templates. This includes details about training, license,
686 limitations, etc.
- 687 • The paper should discuss whether and how consent was obtained from people whose
688 asset is used.
- 689 • At submission time, remember to anonymize your assets (if applicable). You can either
690 create an anonymized URL or include an anonymized zip file.

691 **14. Crowdsourcing and Research with Human Subjects**

692 Question: For crowdsourcing experiments and research with human subjects, does the paper
693 include the full text of instructions given to participants and screenshots, if applicable, as
694 well as details about compensation (if any)?

695 Answer: [NA]

696 Justification:

697 Guidelines:

- 698 • The answer NA means that the paper does not involve crowdsourcing nor research with
699 human subjects.
- 700 • Including this information in the supplemental material is fine, but if the main contribu-
701 tion of the paper involves human subjects, then as much detail as possible should be
702 included in the main paper.
- 703 • According to the NeurIPS Code of Ethics, workers involved in data collection, curation,
704 or other labor should be paid at least the minimum wage in the country of the data
705 collector.

706 **15. Institutional Review Board (IRB) Approvals or Equivalent for Research with Human
707 Subjects**

708 Question: Does the paper describe potential risks incurred by study participants, whether
709 such risks were disclosed to the subjects, and whether Institutional Review Board (IRB)
710 approvals (or an equivalent approval/review based on the requirements of your country or
711 institution) were obtained?

712 Answer: [NA]

713 Justification:

714 Guidelines:

- 715 • The answer NA means that the paper does not involve crowdsourcing nor research with
716 human subjects.
- 717 • Depending on the country in which research is conducted, IRB approval (or equivalent)
718 may be required for any human subjects research. If you obtained IRB approval, you
719 should clearly state this in the paper.
- 720 • We recognize that the procedures for this may vary significantly between institutions
721 and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the
722 guidelines for their institution.
- 723 • For initial submissions, do not include any information that would break anonymity (if
724 applicable), such as the institution conducting the review.