

000
001
002
003
004
005
006
007
008
009
010
011
012
013
014
015
016
017
018
019
020
021
022
023
024
025
026
027
028
029
030
031
032
033
034
035
036
037
038
039
040
041
042
043
044
045
046
047
048
049
050
051
052
053

SEAFORMER: A SPATIAL PROXIMITY AND EDGE-AWARE TRANSFORMER FOR REAL-WORLD VEHICLE ROUTING PROBLEMS

Anonymous authors

Paper under double-blind review

ABSTRACT

Real-world Vehicle Routing Problems (RWVRPs) require solving complex, sequence-dependent challenges at scale with constraints such as delivery time window, replenishment or recharging stops, asymmetric travel cost, etc. While recent neural methods achieve strong results on large-scale classical VRP benchmarks, they struggle to address RWVRPs because their strategies overlook sequence dependencies and underutilize edge-level information, which are precisely the characteristics that define the complexity of RWVRPs. We present SEAFormer, a novel transformer that incorporates both node-level and edge-level information in decision-making through two key innovations. First, our Clustered Proximity Attention (CPA) exploits locality-aware clustering to reduce the complexity of attention from $O(n^2)$ to $O(n)$ while preserving global perspective, allowing SEAFormer to efficiently train on large instances. Second, our lightweight edge-aware module captures pairwise features through residual fusion, enabling effective incorporation of edge-based information and faster convergence. Extensive experiments across four RWVRP variants with various scales demonstrate that SEAFormer achieves superior results over state-of-the-art methods. Notably, SEAFormer is the first neural method to solve 1,000+ node RWVRPs effectively, while also achieving superior performance on classic VRPs, making it a versatile solution for both research benchmarks and real-world applications.

1 INTRODUCTION

The Vehicle Routing Problem (VRP) is a fundamental challenge in logistics, where the goal is to deliver goods to many customers from a depot while minimizing cost and respecting constraints such as vehicle capacity. In practice, VRPs appear most prominently in last-mile deliveries, which have surged dramatically in recent years. For example, in 2022, Manhattan saw over 2.4 million delivery requests per day (Blueprint [2022]), averaging over 3,000 deliveries per minute during a 12-hour workday. Although existing solutions perform well on simplified benchmarks, they rarely account for the operational real-world problems, scales, and constraints encountered in practice.

Real-world Vehicle Routing Problems (RWVRPs) extend the VRP and includes variants such as VRPTW (time window [Kallehauge et al., 2005]), EVRPCS (electric vehicles with recharging - which is important as carrying large weight significantly reduces their driving range [Szumska et al., 2021]), VRPRS (replenishment stops where vehicles can restock to continue service [Schneider et al., 2015]) and AVRP (asymmetric travel costs [Vigo, 1996]). RWVRPs incorporate sequence-dependent constraints and have the following properties: **i) Local infeasibility:** A decision's validity cannot be determined in isolation, and it depends on the complete path history. **ii) State accumulation:** Vehicle state (battery level, current time, etc.) evolves dynamically through the route and is only calculable when the visitation sequence is defined. **iii) Tightly coupled constraints:** Violating one constraint propagates through subsequent decisions. These interdependencies transform the capacity-constrained spatial optimization problem of simple VRPs into a sequence-dependent one and form a tightly coupled setup that can only be validated by considering the entire route sequence.

Recent progress in deep learning has delivered impressive results for combinatorial optimization, reaching near-optimal solutions on classical VRP benchmarks. Existing approaches can be grouped

054 into two main paradigms: autoregressive and non-autoregressive models, each offering unique benefits
 055 and facing specific challenges when extended to RWVRPs.

056 Autoregressive methods (Kool et al., 2018; Kwon et al., 2020; Berto et al., 2025; Li et al., 2024; Lin
 057 et al., 2021; Chen et al., 2022; Wang et al., 2024) designed for VRP or RWVRPs, construct solutions
 058 sequentially using transformers and perform well on small-scale benchmarks. However, their perfor-
 059 mance deteriorates as problem size increases. They lack spatial inductive bias, treating all node pairs
 060 uniformly, disregarding the inherent geometric structure of routing problems. In addition, their full
 061 attention incurs an $O(n^2d)$ memory cost, where d is the embedding dimension, limiting training to
 062 hundreds rather than thousands of nodes. Moreover, they cannot represent edge-specific attributes
 063 (e.g., asymmetric costs, energy use, or traffic), limiting applicability to real-world routing where
 064 pairwise relationships drive feasibility and optimality. These architectural limitations make current
 065 autoregressive methods unsuitable for deployment on industrial-scale real-world routing problems.

066 Recent efforts to address the scalability limitations of autoregressive models introduce new chal-
 067 lenges when applied to RWVRPs. Divide-and-conquer methods (Hou et al., 2022; Nasehi et al.,
 068 2025; Zheng et al., 2024) decompose the problem by clustering customers first and then solving
 069 each cluster independently. However, in RWVRPs this creates a fundamental circular dependency:
 070 assessing whether a cluster is feasible requires knowing the vehicle’s state upon arrival (e.g., remain-
 071 ing battery in EVRPCS or current time in VRPTW), but that state depends on the full route, which is
 072 not yet known during clustering. This mismatch often leads to infeasible clusters that require costly
 073 post-hoc repair or prevents the discovery of high-quality routes altogether (refer to Appendix G).
 074 Other scalable architectures (Gao et al., 2024; Luo et al., 2023, 2025) face their own limitations,
 075 including substantial computational overhead or reliance on supervised learning which requires a
 076 set of pre-computed optimal solutions that are themselves computationally expensive to obtain.

077 In contrast, non-autoregressive methods (Kool et al., 2022; Qiu et al., 2022) improve scalability by
 078 predicting all routing decisions simultaneously through learned heatmaps, thereby avoiding the se-
 079 quential bottleneck of autoregressive decoding. While computationally efficient, these methods face
 080 a fundamental challenge similar to divide-and-conquer approaches: they cannot assess constraint
 081 satisfaction without the knowledge of the traversal sequence. Critical constraints, such as battery
 082 levels, are inherently path-dependent and require sequential state accumulation to ensure feasibility.
 083 As a result, non-autoregressive models are inherently hard to adapt to RWVRPs.

084 In this paper, we introduce **SEAFormer** (Spatial proximity and Edge-Aware transFormer), a novel
 085 architecture that combines the representational power required for RWVRPs with the computational
 086 efficiency necessary for real-world use. Unlike existing methods that treat routing as purely node
 087 selection, our model explicitly reasons about both *where to go* (nodes) and *how feasible/costly that*
 088 *transition is* (edges), a distinction critical for handling edge-level information and improving gen-
 089 eralization. SEAFormer introduces two complementary innovations that together enable scalable,
 090 high-quality solutions across diverse RWVRP and VRP variants.

091 First, we propose **Clustered Proximity Attention (CPA)**, a spatially aware attention mechanism
 092 that improves generalization while reducing the complexity of full attention from $O(n^2)$ to $O(n)$.
 093 Unlike generic sparse attention methods (e.g., Longformer (Beltagy et al., 2020)), CPA leverages
 094 problem-specific spatial patterns to cluster nodes into meaningful partitions and applies attention
 095 within each partition. Our deterministic-yet-diverse clustering strategy offers multiple spatial per-
 096 spectives, to avoid the local optima issues often encountered in sparse attention mechanisms.

097 Second, we introduce a **lightweight edge-aware module** that explicitly models pairwise rela-
 098 tionships between nodes, capabilities largely absent from existing approaches. While node embeddings
 099 capture individual locations and demands, they cannot represent edge-specific attributes. Incorpor-
 100 ating edge-level information not only enables solutions to problems where such details are essential
 101 but also enhances accuracy, generalization, and convergence across models. Our module learns these
 102 relational patterns through a parameter-efficient architecture (increasing number of parameters by
 103 only 7.5%), which is additively combined with the attention decoder to jointly optimize the problem.
 104 **SEAFormer is the first learning-based approach to effectively solve 1000+ node instances across**
 105 **RWVRPs within a single architecture, achieving competitive or superior performance consistently.**

106 We evaluate SEAFormer on four RWVRP variants across a wide range of problem sizes and com-
 107 pare it with state-of-the-art neural and classical methods. For complex variants such as VRPTW,
 EVRPCS, VRPRS, and AVRP, it delivers at least 15% reduction in the objective value over the

108 large-scale instances while fully respecting operational constraints. On standard VRP benchmarks,
 109 SEAFormer consistently surpasses the state-of-the-art approaches. Even on large-scale instances,
 110 SEAFormer preserves solution quality, whereas existing neural solvers often run out of memory or
 111 suffer substantial performance degradation.

112 The contributions of this paper are threefold. (i) We introduce SEAFormer, the first transformer to
 113 jointly optimize spatial proximity and edge-level constraints for real-world VRP variants, achiev-
 114 ing state-of-the-art performance across different benchmarks. (ii) We propose Clustered Proximity
 115 Attention (CPA), a problem-specific sparse attention mechanism that leverages spatial locality in
 116 routing tasks, enhancing generalization and reducing memory complexity through determinis-
 117 tic-yet-diverse clustering. (iii) We propose a parameter-efficient edge-aware module that integrates
 118 pairwise relational information into routing decisions, enabling seamless handling of edge-specific
 119 constraints that are essential for real-world deployment yet missing from existing neural solvers.

120

121 2 RELATED WORK

122 We review existing learning-based approaches to RWVRPs and VRP, focusing on why they struggle
 123 to simultaneously achieve solution quality, computational efficiency, and constraint satisfaction at
 124 real-world scales. A more comprehensive discussion is provided in Appendix J

125

126 2.1 NEURAL SOLVERS WITH LIMITED SCALABILITY

127 Autoregressive neural methods construct solutions sequentially using learned policies, with several
 128 notable approaches in the literature (Kool et al., 2018; Kwon et al., 2020; Berto et al., 2025; Luo
 129 et al., 2023). These models capture sequence dependencies during solution generation. While ef-
 130 fective on small-scale problem instances, their performance deteriorates on large-scale problems.
 131 Moreover, reliance on the full-attention mechanism limits scalability and prevents efficient iterative
 132 refinement on standard hardware.

133

134 2.2 SCALABLE NEURAL APPROACHES

135 Recent approaches for scalable VRP solutions (Hou et al., 2022; Zheng et al., 2024; Nasehi et al.,
 136 2025; Ye et al., 2024) are based on divide-and-conquer method. Such techniques, however, cannot
 137 solve RWVRPs effectively as they cluster nodes before determining visit sequences, yet feasibility of
 138 a solution for RWVRPs depends on those sequences. Whether a vehicle can serve a cluster requires
 139 knowing its state upon arrival, such as battery level in EVRPCS or elapsed time in VRPTW, which
 140 is unavailable during clustering. This circular dependency causes clusters to violate constraints once
 141 sequenced, requiring expensive repairs or prevents finding high-quality solutions.

142

143 Scalable VRP solutions that do not rely on divide-and-conquer strategy face challenges when ap-
 144 plied to RWVRPs. ELG (Gao et al., 2024) restricts attention to nearby nodes using distance-based
 145 penalties, which can limit performance in RWVRPs where vehicles need to reach distant nodes. Fur-
 146 thermore, the local attention approach does not explicitly account for RWVRP-specific constraints,
 147 which may lead to infeasible or suboptimal solutions. Heavy decoder architectures (Luo et al.,
 148 2023; 2025a) rely on supervised learning from near-optimal solutions. For RWVRPs, generating
 149 such training data is computationally expensive. For example, solving 100-node EVRPCS instances
 150 to optimality can be prohibitively time-consuming, and in some cases computationally infeasible,
 151 making these approaches impractical for RWVRPs.

152

153 Non-autoregressive models (Ye et al., 2024); (Qiu et al., 2022); (Kool et al., 2022) generate high-
 154 quality solutions for large-scale VRPs by predicting edge inclusion probabilities via a heatmap.
 155 However, they face specific challenges with sequential constraints in RWVRPs. For instance, in
 156 EVRPCS, edge feasibility depends on the vehicle’s battery state upon arrival, information that is
 157 unavailable during parallel prediction. Applying non-autoregressive methods to sequence-dependent
 158 problems requires costly repair procedures too, which often leads to a drop in solution quality.

159

160 2.2.1 RWVRP-SPECIFIC METHODS

161 A few recent works directly address variations of RWVRPs. Lin et al. (2021) used Transformer-
 162 LSTM to track electric vehicle travel history, while Chen et al. (2022) employed GRUs with two-

stage training to improve charging station routing. Wang et al. (2024) introduced a GAT-based encoder with penalty functions to enforce constraints for EVRPCS. Liu et al. (2024) and Berto et al. (2025) proposed multi-task learning for VRP variants such as VRPTW. Despite these advances, existing methods face at least one of the following limitations: (i) limited scalability, resulting in low-quality solutions for instances larger than 100 nodes; (ii) reliance on manually tuned penalty terms, which slows training and may still produce infeasible solutions requiring costly post-processing; and (iii) sub-optimal strategies for visiting infrastructure nodes, limiting the model’s ability to determine when and where to charge or restock efficiently, leading to lower quality solutions.

2.3 SPARSE ATTENTION MECHANISMS

Recent advances in efficient attention offer potential solutions to scalability challenges. Reformer (Kitaev et al., 2020) leverages locality-sensitive hashing, Longformer (Beltagy et al., 2020) employs sliding windows, and BigBird (Zaheer et al., 2020) combines random, window, and global attention. These mechanisms are designed for text processing and do not naturally align with the geometric structure of routing problems and thus cannot be easily extended to spatial problems. While language tasks mainly involve local sequential dependencies, routing requires radial connectivity between depots and customers. Sparsity patterns based on these methods will disrupt the spatial relationships, limiting performance on routing tasks.

3 CLUSTERED PROXIMITY ATTENTION

Figure 1 shows a VRP solution captured from Hou et al. (2022), highlighting an important structural pattern: nodes within the same route tend to cluster based on their polar coordinates relative to the depot. In particular, they typically follow one of three patterns: (1) similar angles but varying distances from the depot, (2) similar distances but different angles, or (3) close proximity in both angle and distance. This insight motivates our polar clustering approach, as optimal routes naturally reflect these geometric relationships.

We propose **Clustered Proximity Attention (CPA)**, a problem-specific sparse attention mechanism that maintains spatial routing structure while substantially reducing computational complexity. The procedure is detailed below.

Polar-based Spatial Transformation. Given node coordinates $\{x_i \in \mathbb{R}^2\}_{i=0}^n$ with depot n_0 , first we transform each customer location to polar coordinates:

$$r_i = \|x_i - x_0\|_2, \quad \theta_i = \arctan 2(y_i - y_0, x_i - x_0) \in [0, 2\pi), \quad (1)$$

where r_i represents radial distance and θ_i represents angular position relative to the depot.

Partitioning Score. To form clusters capturing the diverse spatial patterns illustrated in Figure 1 and necessary for optimal routing, we introduce a partitioning score that balances radial and angular proximity. After normalizing polar coordinates to $[0, 1]$, we compute a clustering metric as:

$$s_i^{(\alpha)} = \alpha \cdot \bar{\theta}_i + (1 - \alpha) \cdot \bar{r}_i, \quad (2)$$

where $\bar{r}_i = (r_i - r_{\min}) / (r_{\max} - r_{\min})$ and $\bar{\theta}_i = \theta_i / 2\pi$ represent the normalized radial and angular coordinates, respectively, and α is the mixing coefficient that controls the relative importance of distance versus angle in cluster formation.

216 **Deterministic-yet-Diverse Partitioning.** To prevent overfitting to specific cluster configurations,
 217 preserve global context, and capture diverse spatial patterns while maintaining training efficiency,
 218 we use R partitioning rounds with varied mixing coefficients, defining α in Equation 2 as:
 219

$$220 \quad \alpha = t/(R-1), \quad t \in \{0, 1, \dots, R-1\}. \quad (3)$$

221 This creates a spectrum from radial ($\alpha = 0$, grouping nodes at similar distances from depot) to an-
 222 gular clustering ($\alpha = 1$, grouping nodes sharing a similar angle from the depot). Each configuration
 223 captures different spatial patterns shown in Figure 1, entails to a proper global prospective.
 224

225 **Boundary Smoothing for Robust Clustering.** Hard cluster boundaries can artificially separate
 226 nearby nodes, disrupting natural customer groups. We address this through a boundary smoothing
 227 technique. After sorting the calculated partitioning scores $S_\alpha = [s_\alpha^1, \dots, s_\alpha^n]$, given M as the cluster
 228 size (a user-defined parameter analyzed in Appendix F), we apply a circular shift as:
 229

$$230 \quad S'_\alpha = [s_\alpha^{\lfloor M/2 \rfloor + 1}, \dots, s_\alpha^n, s_\alpha^1, \dots, s_\alpha^{\lfloor M/2 \rfloor}] \quad (4)$$

231 This rotation softens the deterministic cluster boundaries that would otherwise push nearby nodes
 232 apart, keeping close nodes together and improving transitions between clusters. After computing
 233 the partitioning scores $S = \{S_{\alpha_1}, S'_{\alpha_1}, S_{\alpha_2}, \dots\}$ across different proximity weights and rotations,
 234 we partition nodes into clusters of size M using the partitioning score. **Note that when n is not**
 235 **perfectly divisible by (i.e., when $\lceil n/M \rceil \times M > n$, we pad the final cluster with depot nodes to**
 236 **ensure uniform cluster size M across all partitions).**

237 **Localized Attention Computation.** Given partition $\mathcal{P} = \{C_{\alpha_1}^1, C_{\alpha_1}^2, \dots, C_{\alpha_1}^k, \dots\}$ where each
 238 C_α^j contains M proximate nodes, we compute attention independently within each cluster as:
 239

$$240 \quad \text{Attention}(Q_j, K_j, V_j) = \text{softmax} \left(\frac{Q_j K_j^\top}{\sqrt{d}} \right) V_j, \quad \text{where } Q_j, K_j, V_j \in \mathbb{R}^{|C_\alpha^j| \times d}. \quad (5)$$

243 **Complexity analysis:** CPA reduces complexity of attention from $O(n^2)$ to $O(n)$ while preserving
 244 the spatial relationships critical for routing decisions. It partitions nodes into clusters of size M .
 245 Each cluster requires $O(M^2)$ operations for attention computation. With $\lceil n/M \rceil$ clusters total:
 246

$$247 \quad \text{CPA complexity} = O \left(\underbrace{R}_{\text{No. rounds}} \times \underbrace{\lceil n/M \rceil}_{\text{clusters}} \times \underbrace{(M^2)}_{\text{per cluster}} \right) = O(nRM) = O(n) \quad (6)$$

250 3.1 VISUAL EXAMPLE OF CPA

252 Figure 2 illustrates the steps of CPA. Figure 2a shows the problem instance, where customer nodes
 253 (blue dots) are distributed in a 2D Euclidean space around the depot (red star). The instance is then
 254 transformed into polar coordinates relative to the depot (Figure 2b). Next, the partitioning score
 255 from Equation 2 is computed (Figure 2c), after which the nodes are sorted by this score and grouped
 256 into clusters based on the selected cluster size (Figure 2d). Finally, Figure 2e compares CPA's
 257 attention scores with standard full attention (Figure 2f), highlighting CPA's locality-aware sparsity
 258 and reduced memory footprint.

260 4 SEAFormer ARCHITECTURE

262 Figure 3 shows the SEAFormer architecture, a dual-stream design that jointly models node-level
 263 spatial structure and edge-level information for RWVRPs. The node stream encodes location fea-
 264 tures using our Clustered Proximity Attention (presented in Section 3), while the edge-aware stream
 265 maintains explicit embeddings for pairwise attributes such as costs and distances. These two streams
 266 converge in the decoder: node embeddings drive the sequential step-by-step node selection, and edge
 267 embedding is used to rank candidate transitions from the current position through a learned heatmap
 268 generation. This separation of roles allows SEAFormer to scale to thousands of nodes, converge
 269 faster, and remain expressive enough to handle complex real-world constraints. The encoder of
 SEAFormer comprises two complementary modules.

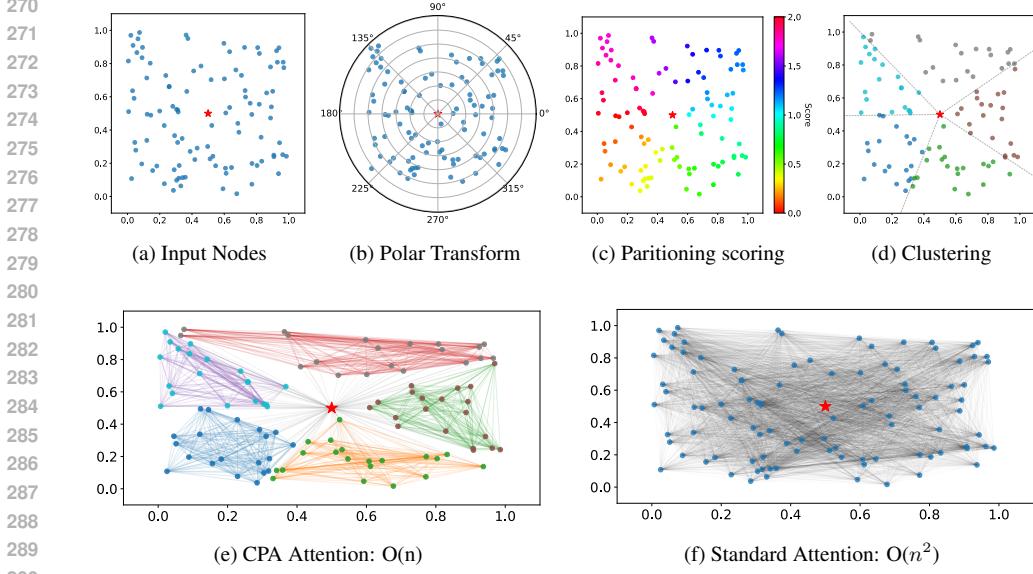


Figure 2: Pipeline of CPA. (a) Input nodes in Cartesian space with depot (red star), (b) Polar transformation relative to depot, (c) Angular scoring with $\alpha = 0$ (pure angular), (d) Nodes sorted by angle and partitioned into fixed-size clusters, (e) Final CPA attention pattern with $O(n)$ complexity, (f) Standard attention with $O(n^2)$ complexity showing all pairwise connections. CPA reduces number of attention calculation for this example from 10,000 to 2100 (approximately 80% reduction).

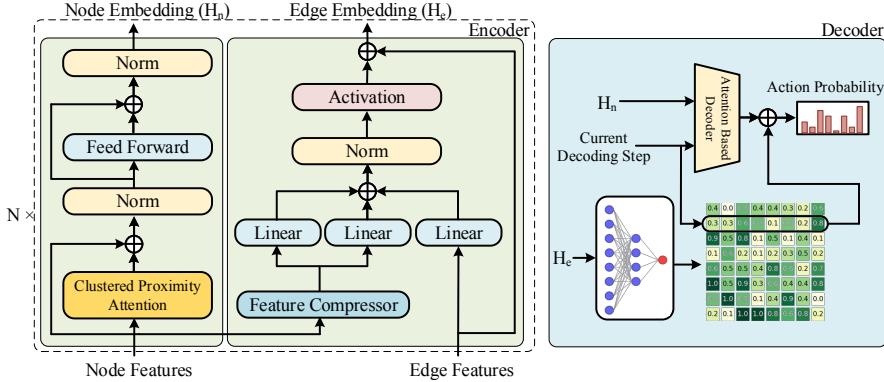


Figure 3: SEAFormer Architecture. The dual-module encoder embeds nodes through CPA and edges through a lightweight residual module. The dual-path decoder combines edge-aware guidance (heatmap) with sequential node selection (attention), unified through logit fusion. This design enables scalable training on large number of nodes while handling diverse RWVRP constraints.

Node embedding through CPA. This module adopts the encoder architecture from [Kwon et al. \(2020\)](#), encoding node features through L layers. Unlike prior works that employ full attention mechanisms, SEAFormer utilizes CPA to produce spatially-aware node embeddings $H_n \in \mathbb{R}^{n \times d}$, while achieving significant memory savings in training and inference.

For problem variants with optional service nodes such as EVRPCS (charging stations) and VRPRS (replenishment stops), we augment the spatial encoder with a specialized optional node processing pathway. This parallel encoding layer generates embeddings for optional nodes $H_o \in \mathbb{R}^{f \times d}$, where

f represents the number of service facilities. These embeddings are then fused with customer embeddings through a learnable mechanism, and the combined representations pass through the same batch normalization and feed-forward layers. This allows each customer embedding to account for nearby optional facilities and their influence on route feasibility, which is crucial in problems where strategic service stops can substantially enhance or even enable high-quality solutions. A comprehensive description of the optional node embedding is provided in Appendix [H](#)

Edge-Aware Embedding Module. We design our edge embedding to preserve the properties that exist for edges while maintaining computational efficiency. The edge-aware embedding module only operates on edges between depots and customers, while optional nodes in EVRPCS and VRPRS do not participate in this process. Given edge features (distance, energy consumption rate, historical traffic) embedding $X_e^{(i,j)} \in \mathbb{R}^{d_e}$, and node embeddings $H_n^{(i)}, H_n^{(j)} \in \mathbb{R}^d$, we compute edge-aware embeddings through residual fusion (Szegedy et al. 2017):

$$H_e^{(i,j)} = X_e^{(i,j)} + \sigma \left(\text{BN} \left(W_1 H_n^{(i)} + W_2 H_n^{(j)} + W_3 X_e^{(i,j)} \right) \right) \quad (7)$$

Where W is a trainable parameter, $W_1 H_n^{(i)}$ captures origin-specific factors, $W_2 H_n^{(j)}$ captures destination-specific factors, $W_3 X_e^{(i,j)}$ captures edge features. Prior to this, we apply a linear transformation to reduce node features to match edge embedding dimensions, enabling the model to extract only relative information from node embeddings. Batch normalization ensures stable training across diverse edge scales and SILU activation (σ) enables modeling of complex non-linear relationships. This architecture anchors edge representations in spatial context while increasing the model's total parameters by only 7.5%. Lastly, a residual connection maintains original edge information while selectively integrating node-level context.

Once edge and node embeddings are generated, SEAFormer's decoder constructs the solution through dual complementary pathways that balance global optimization with sequential decoding.

Edge-Guided Global Heatmap. Before sequential decoding, we produce a static heatmap that encodes global edge preferences. This step occurs once at the onset of the decoding process:

$$\mathcal{H}_{ij} = \tanh(\text{MLP}_\theta(H_e^{(i,j)})) \in \mathbb{R}, \quad \forall i, j \in \{1, \dots, n\}, \quad (8)$$

where MLP_θ is a multi-layer perceptron that maps each edge to a scalar score, followed by the \tanh activation function. This module encodes edge features (e.g., travel times, distances) that are impossible to capture through node-level encoding.

Node-Guided Sequential Attention. The sequential decoder constructs solutions autoregressively, selecting one node at each step t based on the current state of the partially generated solution. To enable the model to differentiate between vehicle states across problem variants, we adapt the query vector to each RWVRP variant's state representation $\mathbf{q}_t = [\mathbf{h}_t; c_t; \xi_t]$, where \mathbf{h}_t is the last visited node embedding and ξ_t represents variant-specific constraints: none for VRP, battery level b_t for EVRPCS, remaining travel length τ_t for VRPRS, and current time w_t for VRPTW. To ensure stable convergence and prevent bias from large values, we normalize ξ_t by its maximum value to maintain the range $[0, 1]$. Finally, the attention mechanism computes compatibility scores with all unvisited nodes through:

$$\alpha_{ti} = \begin{cases} \frac{\exp(\mathbf{q}_t^\top \mathbf{k}_i / \sqrt{d})}{\sum_{j \in \mathcal{U}_t} \exp(\mathbf{q}_t^\top \mathbf{k}_j / \sqrt{d})} & \text{if } i \in \mathcal{U}_t \\ 0 & \text{otherwise} \end{cases} \quad (9)$$

where \mathcal{U}_t denotes the set of unvisited feasible nodes. We implement a proactive masking function that enhances solution quality, reduces search space, and prevents the model from generating infeasible solutions during inference (see Appendix [B](#) for detailed description). Finally, at each step t , we combine logits from both paths:

$$\ell_t = \ell_t^{\text{seq}} + \mathcal{H}_{i_t,:}, \quad (10)$$

where ℓ_t^{seq} is sequential attention logits, $\mathcal{H}_{i_t,:}$ is the heatmap row for current node i_t . For optional nodes not included in the heatmap, we define heatmap values as $\mathcal{H}_{i_t,j} = -2 \cdot \frac{r_t}{R_{\max}}$, $j \in \mathcal{O}$ to encourage model to visit such nodes more frequently as resources deplete. Here, r_t denotes the remaining resource level (battery charge in EVRPCS or available driving range in VRPRS), R_{\max} represents the maximum resource capacity (full battery in EVRPCS or max driving range in VRPRS), and \mathcal{O} is set of optional nodes.

378

5 EXPERIMENTS

380 To verify the applicability of SEAFormer on different variations of RWVRP and VRP, we evaluate
 381 SEAFormer on 5 combinatorial optimization problems, including VRPTW, EVRPCS, VRPRS,
 382 AVRPs, and VRP. Detailed formulations of these problems are provided in Appendix A.

383 **Benchmarks.** We evaluate SEAFormer across four datasets: (i) Random RWVRP benchmarks,
 384 comprising 100 instances per problem size with up to 7K nodes (results up to 1k are shown in
 385 table 1 and larger values in appendix C) following Kwon et al. (2020); Zhou et al. (2023); (ii)
 386 Random VRP benchmarks with up to 7K nodes following Zhou et al. (2023); (iii) real-world VRP
 387 instances from CVRPLib; (iv) cross-distribution datasets generated by Zhou et al. (2023).

388 **Implementation.** For all problems, we train the model for 2000 epochs on 100-customer instances,
 389 then 200 epochs on 500-customer instances, and 100 epochs on 1000-customer instances. As in
 390 prior work, problem instances are uniformly sampled from the $[0, 1]^2$ space, and demands drawn
 391 from a discrete uniform distribution on $[1, 10]$. Asymmetric VRPs are created using our dataset gen-
 392 eration procedure (refer to appendix A.4.1). The Adam optimizer (Kingma 2014) is employed for
 393 training, with detailed hyperparameters for each problem instance as well as training times provided
 394 in Appendix K and Appendix L.

395 **Metrics.** We report the mean objective (Obj.), gap (G), and inference time (T) for each approach.
 396 Objective represents solution length, where lower values signify superior performance. Gap quanti-
 397 fies the deviation from solutions generated by one of OR-tools, LKH, or HGS. Time denotes the
 398 total runtime across the entire dataset, measured in seconds (s), minutes (m), or hours (h). Runtimes,
 399 measured on identical hardware (NVIDIA A100 GPU for neural methods, 32-core CPU for classical
 400 solvers), exclude model loading and represent the total solution time over each dataset.

401 **Inference.** In the inference phase, we evaluate SEAFormer using two strategies. First, greedy decod-
 402 ing with 8-fold augmentation, and second, Simulation-Guided Beam Search (SGBS) (Choo et al.,
 403 2022), which requires additional computation time but yields superior results as it explores multiple
 404 solution trajectories simultaneously. The greedy variant of SEAFormer operates significantly faster
 405 than the SGBS version at the expense of marginally reduced solution quality.

406 **Baselines.** We compare SEAFormer with 1) **Classical Solvers:** HGS-PyVRP (Wouda et al., 2024),
 407 LKH (Helsgaun, 2017), and OR-Tools; 2) **Construction-based NCO Methods:** POMO (Kwon
 408 et al., 2020), MTPOMO (Liu et al., 2024), EVPRPL (Lin et al., 2021), EVGAT (Wang et al., 2024),
 409 LEHD (Luo et al., 2023), RELD (Huang et al., 2025), ELG (Gao et al., 2024), L2C-insert (Luo
 410 et al., 2025b), UniteFormer (Meng et al., 2025), Eformer (Meng et al., 2025) DAR (Wang et al., 2025),
 411 BLEHD (Luo et al., 2025a), and RouteFinder (Berto et al., 2025); 3) **Divide-and-conquer based**
 412 **methods:** GLOP (Ye et al., 2024), DeepMDV (Nasehi et al., 2025), and UDC (Zheng et al., 2024).
 413 Not all existing methods apply to every RWVRP variant, so comparisons are limited to applicable
 414 approaches. For details on baselines and their integration, see Appendix M.

416

5.1 EVALUATION ON RWVRPs

417 Table 1 presents a comprehensive evaluation of SEAFormer against state-of-the-art methods across
 418 four challenging RWVRPs. SEAFormer consistently achieves superior or competitive performance
 419 across all problem variants and scales. Notably, using SGBS as a search method, SEAFormer es-
 420 tablishes new state-of-the-art results in learning-based methods across all 12 test configurations be-
 421 tween, at the cost of higher processing time, particularly with impressive gains on larger instances.

422 While SEAFormer achieves superior performance on VRPTW, EVRPCS, and VRPRS across all set-
 423 tings where existing large-scale solutions fail, the AVRPs results especially showcase SEAFormer’s
 424 architectural strengths. Whereas learning-based approaches including POMO, LEHD, BLEHD, and
 425 UDC struggle with asymmetric distance matrices and show gaps up to 3% with OR-Tools on 100-
 426 customer instances, SEAFormer achieves the best performance with a 1.7% improvement over OR-
 427 Tools and increasingly larger gains on bigger instances. This highlights our model’s superior capac-
 428 ity to capture complex spatial dependencies inherent in asymmetric routing problems.

429 The results highlight SEAFormer’s exceptional scalability: while most baselines suffer severe degra-
 430 dation on 1,000-customer RWVRPs, SEAFormer preserves solution quality and remains the fastest
 431 method. The SGBS variant, though requiring additional computation (95-100 minutes for 1k cus-

Table 1: Objective function (Obj.), Gap to the OR-tools (Gap), and solving time (Time) on 100, 500, and 1,000-node RWVRPs. All test sets contain 100 instances following settings in Zhou et al. (2023). The overall best performance is in bold and the best learning method is marked by shade. OR-Tools results are not optimal, as execution was stopped early due to time limits. Methods not tailored to EVRPCS or VRPRS are excluded due to sub-optimal performance. (Refer to Section G).

RWVRP	METHODS	100 CUSTOMERS			500 CUSTOMERS			1K CUSTOMERS		
		Obj. \downarrow	G(%)	T	Obj. \downarrow	G(%)	T	Obj. \downarrow	G(%)	T
VRPTW	OR-TOOLS	26.34	0.00	1H	87.3	0.00	5H	151.4	0.00	10H
	HGS-PyVRP	26.04	-1.1	1H	83.8	-4.00	5H	142.6	-5.6	10H
	POMO	26.81	1.78	5S	93.2	6.75	30S	193.2	27.6	1M
	DAR	27.1	2.88	5S	97.2	11.3	30S	212.2	40.1	1M
	MTPOMO	27.02	2.58	5S	96.8	10.8	30S	207.1	36.7	1M
	ROUTEFINDER	26.8	1.74	5S	91.2	4.46	30S	171.4	11.3	1M
EVRPCS	SEAFORMER	26.75	1.55	5S	87.4	0.11	30S	149.7	-1.1	1M
	SEAFORMER-SGBS	26.5	0.6	10S	85.0	-2.6	13M	145.2	-4.1	1.6H
	OR-TOOLS	16.35	0.00	1H	31.1	0.00	5H	47.8	0.00	10H
	EVPRPL	16.54	1.16	10S	34.3	10.3	90S	62.1	29.9	3M
	EVGAT	16.9	3.36	30S	35.2	13.2	3M	56.3	17.8	6M
	SEAFORMER	16.36	0.06	5S	30.8	-1.0	30S	45.8	-4.2	1M
VRPRS	SEAFORMER-SGBS	16.14	-1.3	15S	30.2	-2.9	13M	44.9	-6.1	1.7H
	OR-TOOLS	11.2	0.00	1H	23.4	0.00	5H	36.11	0.00	10H
	EVPRPL	11.43	2.05	10S	25.16	7.52	90S	47.5	31.5	3M
	EVGAT	11.76	5.0	30S	25.65	9.6	3M	42.1	16.6	6M
	SEAFORMER	11.24	0.35	5S	22.93	-2.0	30S	34.87	-3.4	1M
	SEAFORMER-SGBS	10.97	-2.0	15S	22.33	-4.6	13M	33.95	-6.0	1.7H
AVRP	OR-TOOLS	19.37	0.00	1H	40.27	0.00	5H	47.6	0.00	10H
	POMO	19.5	0.7	10S	42.12	4.59	20S	53.44	12.3	1M
	MTPOMO	19.63	1.3	10S	44.4	10.2	3M	56.2	18.1	1M
	ROUTEFINDER	19.6	1.2	5S	41.52	3.1	20S	48.3	1.47	1M
	LEHD	19.96	3.0	5S	40.71	1.1	20S	45.89	-3.6	1.2M
	BLEHD-PRC50	-	-	-	42.13	4.61	2.5M	45.23	-5.0	6M
	UDC ₂₅₀ ($\alpha = 50$)	-	-	-	40.06	-0.5	30M	45.17	-5.1	1.2H
	SEAFORMER	19.48	0.6	5S	40.23	-0.1	30S	45.14	-5.2	1M
	SEAFORMER-SGBS	19.04	-1.7	10S	38.97	-3.2	13M	44.07	-7.4	1.7H

tomers), consistently produces the best solutions across all scales, suggesting an effective trade-off between solution quality and computational resources.

5.2 EVALUATION ON VRP

Table 2 reports SEAFormer’s performance on VRP instances with 100–1000 customers, showing both strong scalability and high solution quality. On small instances, SEAFormer achieves 0.96% gap in 5 seconds, and with SGBS it matches POMO, and slightly outperform UniteFormer and Eformer. The advantage becomes clearer at larger scales: on 500-customer instances, SEAFormer achieves a 4.64% gap in 15 seconds, which outperforms POMO and rivaling methods like GLOPLKH3. With SGBS, SEAFormer further improves to a 2.98% gap, surpassing all learning-based approaches. On the 1,000-customer setting, where most neural methods degrade, SEAFormer proves robust, achieving a 0.62% gap in 30 seconds, surpassing LEHD (1.05% in 2 minutes) and very recent methods such as RELD (2.36%) and L2C-Insert (0.82%). With SGBS, SEAFormer outperforms HGS by a 0.9% gap and matches UDC and BLEHD.

Moreover, we evaluate SEAFormer on benchmarks across very large-scale (Appendix C), real-world VRP, EVRPCS, and VRPTW datasets (Appendix D), and cross-distribution (Appendix E) settings. In all cases, SEAFormer performs well, establishing itself as a robust solution for research benchmarks and real-world applications.

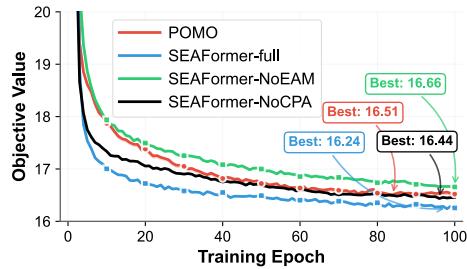
5.3 ABLATION STUDY

Figure 4 shows that both EAM and CPA are essential for SEAFormer’s performance. SEAFormer converges faster than POMO, whereas removing EAM results in noticeably slower convergence. This slowdown occurs because clustering modules such as CPA partition the attention space, which naturally reduces convergence speed. SEAFormer-NoCPA uses POMO’s general encoder and therefore converges faster than SEAFormer-NoEAM, reaching performance close to POMO, indicating

486 Table 2: Performance comparison of various methods on 100 VRP instances with 100, 500, and
 487 1,000 customers. Best values are bolded while the best learning-based solutions are highlighted.
 488

METHODS	100 CUSTOMERS			500 CUSTOMERS			1K CUSTOMERS		
	OBJ. \downarrow	G(%)	T	OBJ. \downarrow	G(%)	T	OBJ. \downarrow	G(%)	T
HGS-PyVRP	15.5	0.00	40M	36.84	0.00	4H	43.5	0.00	8H
POMO	15.72	1.41	5s	44.8	21.6	20s	101	132	3M
GLOP-LKH3	21.3	36.5	30s	42.45	15.22	3M	45.9	5.51	2M
DEEPMDV-LKH3	16.2	4.51	90s	40.2	9.12	4M	45.0	3.44	8M
ELG	15.8	1.93	30s	38.34	4.07	2.6M	43.58	0.18	15M
UNITEFORMER	15.74	1.54	5s	41.2	11.8	3M	61.8	42.0	22M
EFORMER	15.77	1.74	6s	46.8	27.03	3.3M	87.4	100	27M
DAR	-	-	-	38.21	3.7	30s	43.82	0.73	3M
LEHD	16.2	4.51	5s	38.41	4.26	20s	43.96	1.05	2M
BLEHD-PRC50	-	-	-	41.50	12.64	2.5M	43.13	-0.8	6M
RELD	15.75	1.61	5s	38.33	4.04	30s	44.53	2.36	50s
L2C-INSERT	15.72	1.41	2M	38.72	5.1	7M	43.86	0.82	13M
UDC ₅₀ ($\alpha = 50$)	-	-	-	38.34	4.07	7M	43.48	0.00	14M
UDC ₂₅₀ ($\alpha = 50$)	-	-	-	37.99	3.12	30M	43.00	-1.1	1.2H
SEAFormer	15.82	2.06	5s	38.55	4.64	15s	43.77	0.62	30s
SEAFormer-SGBS	15.72	1.41	30s	37.94	2.98	12M	43.10	-0.9	1.5H

502
 503
 504 that EAM combined with a general encoder can perform slightly better than models without EAM.
 505 Together, CPA and EAM create a complementary effect: CPA provides scalability through its $O(n)$
 506 clustering mechanism and geometrically informed attention, while EAM delivers edge guidance that
 507 accelerates learning. Neither component achieves these advantages alone.
 508



520 Figure 4: Training curves on 100-customer VRP
 521 (NoCPA/NoEAM = Without CPA/EAM).

523 To validate CPA’s effectiveness, we replace it with alternative clustering strategies in SEAFormer
 524 (Table 3). Reformer with LSH4 yields the smallest gap (0.3–1.27%), showing hash-based clustering
 525 can be competitive, though CPA remains superior. Multi-round K-Means improves over single-
 526 round, reducing the gap from 28.7% to 3.9% on VRP1000, underscoring the importance of iterative
 527 refinement. Grid-based clustering performs poorly, with up to 24.05% gap on VRP1000. These
 528 results confirm CPA’s advantage for routing. For a more extensive ablation study, see Appendix F.

6 CONCLUSION AND FUTURE WORK

533 In this paper, we introduce SEAFormer, a scalable, edge-aware transformer for VRP and RWVRP
 534 variants that combines Clustered Proximity Attention with a lightweight edge module. SEAFormer
 535 outperforms state-of-the-art RWVRP methods by at least 15% on large-scale instances, converges
 536 rapidly, and achieves strong performance on standard VRPs. Although the current approach has
 537 limitations, such as its single-depot focus and the training cost associated with handling different
 538 VRP variants, future work may explore multi-task learning across RWVRPs to develop a unified
 539 model trained once, extensions to the MDVRP setting, and heatmap-guided search strategies beyond
 greedy decoding to further leverage our edge-aware representations.

540 REFERENCES

541 Irwan Bello, Hieu Pham, Quoc V Le, Mohammad Norouzi, and Samy Bengio. Neural combinatorial
 542 optimization with reinforcement learning. *arXiv preprint arXiv:1611.09940*, 2016.

543

544 Iz Beltagy, Matthew E Peters, and Arman Cohan. Longformer: The long-document transformer.
 545 *arXiv preprint arXiv:2004.05150*, 2020.

546

547 Federico Berto, Chuanbo Hua, Nayeli Gast Zepeda, André Hottung, Niels Wouda, Leon Lan, Juny-
 548 oung Park, Kevin Tierney, and Jinkyoo Park. Routefinder: Towards foundation models for vehicle
 549 routing problems. *arXiv preprint arXiv:2406.15007*, 2025.

550 E-Commerce Blueprint. In NY daily news: Manhattan BP calls for delivery reforms as NYC resi-
 551 dents, businesses receive more than 2.4 million packages per day. <https://shorturl.at/1hX8B>, 2022.

552

553 Jinbiao Chen, Huanhuan Huang, Zizhen Zhang, and Jiahai Wang. Deep reinforcement learning with
 554 two-stage training strategy for practical electric vehicle routing problem with time windows. In
 555 *International Conference on Parallel Problem Solving from Nature*, pp. 356–370. Springer, 2022.

556

557 Jinho Choo, Yeong-Dae Kwon, Jihoon Kim, Jeongwoo Jae, André Hottung, Kevin Tierney, and
 558 Youngjune Gwon. Simulation-guided beam search for neural combinatorial optimization. *Ad-*
 559 *vances in Neural Information Processing Systems (NeurIPS)*, 35:8760–8772, 2022.

560

561 Tri Dao, Dan Fu, Stefano Ermon, Atri Rudra, and Christopher Ré. Flashattention: Fast and memory-
 562 efficient exact attention with io-awareness. *Advances in Neural Information Processing Systems*
 563 (*NeurIPS*), 35:16344–16359, 2022.

564

565 Han Fang, Zhihao Song, Paul Weng, and Yutong Ban. Invit: A generalizable routing problem solver
 566 with invariant nested view transformer. In *International Conference on Machine Learning*, pp.
 567 12973–12992. PMLR, 2024.

568

569 Ángel Felipe, M Teresa Ortuño, Giovanni Righini, and Gregorio Tirado. A heuristic approach for the
 570 green vehicle routing problem with multiple technologies and partial recharges. *Transportation
 571 Research Part E: Logistics and Transportation Review*, 71:111–128, 2014.

572

573 Zhang-Hua Fu, Kai-Bin Qiu, and Hongyuan Zha. Generalize a small pre-trained model to arbitrarily
 574 large tsp instances. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*,
 575 volume 35, pp. 7474–7482, 2021.

576

577 Chengrui Gao, Haopu Shang, Ke Xue, Dong Li, and Chao Qian. Towards generalizable neural
 578 solvers for vehicle routing problems via ensemble with transferrable local policy. In *Proceedings*
 579 *of the Thirty-Third International Joint Conference on Artificial Intelligence (IJCAI)*, 2024.

580

581 Hermann Gehring and Jörg Homberger. Extended VRPTW benchmark. <https://www.sintef.no/projectweb/top/vrptw/homberger-benchmark/> Accessed: 2025-
 582 11-30.

583

584 Keld Helsgaun. An extension of the lin-kernighan-helsgaun tsp solver for constrained traveling
 585 salesman and vehicle routing problems. *Roskilde: Roskilde University*, 12:966–980, 2017.

586

587 Qingchun Hou, Jingwei Yang, Yiqiang Su, Xiaoqing Wang, and Yuming Deng. Generalize learned
 588 heuristics to solve large-scale vehicle routing problems in real-time. In *The Eleventh International
 589 Conference on Learning Representations (ICLR)*, 2022.

590

591 Ziwei Huang, Jianan Zhou, Zhiguang Cao, and Yixin XU. Rethinking light decoder-based solvers
 592 for vehicle routing problems. In *The Thirteenth International Conference on Learning Represen-*
 593 *tations (ICLR)*, 2025.

594

595 Eric Jang, Shixiang Gu, and Ben Poole. Categorical reparameterization with gumbel-softmax. *arXiv
 596 preprint arXiv:1611.01144*, 2016.

597

598 Chaitanya K Joshi, Thomas Laurent, and Xavier Bresson. On learning paradigms for the travelling
 599 salesman problem. *arXiv preprint arXiv:1910.07210*, 2019.

594 Brian Kallehauge, Jesper Larsen, Oli BG Madsen, and Marius M Solomon. Vehicle routing problem
 595 with time windows. In *Column Generation*, pp. 67–98. Springer, 2005.
 596

597 Surendra Reddy Kanchala and Gitakrishnan Ramadurai. An adaptive large neighborhood search
 598 approach for electric vehicle routing with load-dependent energy consumption. *Transportation in
 599 Developing Economies*, 4:1–9, 2018.

600 Elias Khalil, Hanjun Dai, Yuyu Zhang, Bistra Dilkina, and Le Song. Learning combinatorial opti-
 601 mization algorithms over graphs. *Advances in Neural Information Processing Systems (NeurIPS)*,
 602 30, 2017.

603 Diederik P Kingma. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*,
 604 2014.

605

606 Nikita Kitaev, Lukasz Kaiser, and Anselm Levskaya. Reformer: The efficient transformer. In
 607 *International Conference on Learning Representations (ICLR)*, 2020.

608

609 Çağrı Koç, Ola Jabali, Jorge E Mendoza, and Gilbert Laporte. The electric vehicle routing problem
 610 with shared charging stations. *International Transactions in Operational Research*, 26(4):1211–
 611 1243, 2019.

612 Wouter Kool, Herke van Hoof, and Max Welling. Attention, learn to solve routing problems! In
 613 *International Conference on Learning Representations (ICLR)*, 2018.

614

615 Wouter Kool, Herke van Hoof, Joaquim Gromicho, and Max Welling. Deep policy dynamic pro-
 616 gramming for vehicle routing problems. In *International Conference on Integration of Constraint
 617 Programming, Artificial Intelligence, and Operations Research*, pp. 190–213, 2022.

618

619 Yeong-Dae Kwon, Jinho Choo, Byoungjip Kim, Iljoo Yoon, Youngjune Gwon, and Seungjai Min.
 620 POMO: Policy optimization with multiple optima for reinforcement learning. *Advances in Neural
 621 Information Processing Systems (NeurIPS)*, 33:21188–21198, 2020.

622

623 Jinqi Li, Bing Tian Dai, Yunyun Niu, Jianhua Xiao, and Yaoxin Wu. Multi-type attention for solving
 624 multi-depot vehicle routing problems. *IEEE Transactions on Intelligent Transportation Systems
 (TITS)*, 2024.

625

626 Sirui Li, Zhongxia Yan, and Cathy Wu. Learning to delegate for large-scale vehicle routing. *Ad-
 627 vances in Neural Information Processing Systems (NeurIPS)*, 34:26198–26211, 2021.

628

629 Wei Li, Bing Tian Dai, Xueming Yan, Junying Zou, Zhijie Liang, and Jingwen Li. Learning to solv-
 630 ing vehicle routing problems via local–global feature fusion transformer. *Complex & Intelligent
 Systems*, 11(9):392, 2025.

631

632 Bo Lin, Bissan Ghaddar, and Jatin Nathwani. Deep reinforcement learning for the electric vehicle
 633 routing problem with time windows. *IEEE Transactions on Intelligent Transportation Systems
 (TITS)*, 23(8):11528–11538, 2021.

634

635 Fei Liu, Xi Lin, Zhenkun Wang, Qingfu Zhang, Tong Xialiang, and Mingxuan Yuan. Multi-task
 636 learning for routing problem with cross-problem zero-shot generalization. In *Proceedings of the
 637 30th ACM Conference on Knowledge Discovery and Data Mining (SIGKDD)*, pp. 1898–1908,
 2024.

638

639 Fu Luo, Xi Lin, Fei Liu, Qingfu Zhang, and Zhenkun Wang. Neural combinatorial optimization with
 640 heavy decoder: Toward large scale generalization. *Advances in Neural Information Processing
 641 Systems (NeurIPS)*, 36:8845–8864, 2023.

642

643 Fu Luo, Xi Lin, Yaoxin Wu, Zhenkun Wang, Tong Xialiang, Mingxuan Yuan, and Qingfu Zhang.
 644 Boosting neural combinatorial optimization for large-scale vehicle routing problems. In *The Thir-
 645 teenth International Conference on Learning Representations (ICLR)*, 2025a.

646

647 Fu Luo, Xi Lin, Mengyuan Zhong, Fei Liu, Zhenkun Wang, Jianyong Sun, and Qingfu Zhang.
 Learning to insert for constructive neural vehicle routing solver. *Advances in Neural Information
 Processing Systems (NeurIPS)*, 2025b.

648 Michalis Mavrovouniotis, Georgios Ellinas, and Marios Polycarpou. Ant colony optimization for the
 649 electric vehicle routing problem. In *2018 IEEE Symposium series on computational intelligence*
 650 (*SSCI*), pp. 1234–1241. IEEE, 2018.

651 Michalis Mavrovouniotis, Charalambos Menelaou, Stelios Timotheou, Georgios Ellinas, Christos
 652 Panayiotou, and Marios Polycarpou. A benchmark test suite for the electric capacitated vehicle
 653 routing problem. In *2020 IEEE Congress on evolutionary computation (CEC)*, pp. 1–8. IEEE,
 654 2020.

655 Dian Meng, Zhiguang Cao, Jie Gao, Yaxin Wu, and Yaqing Hou. Uniteformer: Unifying node
 656 and edge modalities in transformers for vehicle routing problems. In *The Thirty-ninth Annual*
 657 *Conference on Neural Information Processing Systems*.

658 Dian Meng, Zhiguang Cao, Yaxin Wu, Yaqing Hou, Hongwei Ge, and Qiang Zhang. Eformer: An
 659 effective edge-based transformer for vehicle routing problems. *IJCAI*, 2025.

660 Saeed Nasrabi, Farhana Choudhury, Egemen Tanin, and Majid Sarvi. Deepmdv: Global spatial
 661 matching for multi-depot vehicle routing problems. *International Conference on Advances in*
 662 *Geographic Information Systems (SIGSPATIAL)*, 2025.

663 Mohammadreza Nazari, Afshin Oroojlooy, Lawrence Snyder, and Martin Takáć. Reinforcement
 664 learning for solving the vehicle routing problem. *Advances in Neural Information Processing*
 665 *Systems (NeurIPS)*, 31, 2018.

666 Alex Nowak, David Folqué, and Joan Bruna. Divide and conquer networks. In *International Con-*
 667 *ference on Learning Representations (ICLR)*, 2018.

668 Ruizhong Qiu, Zhiqing Sun, and Yiming Yang. Dimes: A differentiable meta solver for combinator-
 669 ial optimization problems. *Advances in Neural Information Processing Systems (NeurIPS)*, 35:
 670 25531–25546, 2022.

671 Michael Schneider, Andreas Stenger, and Dominik Goeke. The electric vehicle-routing problem
 672 with time windows and recharging stations. *Transportation Science*, 48(4):500–520, 2014.

673 Michael Schneider, Andreas Stenger, and Julian Hof. An adaptive vns algorithm for vehicle routing
 674 problems with intermediate stops. *Or Spectrum*, 37:353–387, 2015.

675 Christian Szegedy, Sergey Ioffe, Vincent Vanhoucke, and Alexander Alemi. Inception-v4, inception-
 676 resnet and the impact of residual connections on learning. In *Proceedings of the AAAI Conference*
 677 *on Artificial Intelligence (AAAI)*, volume 31, 2017.

678 Emilia M Szumska, Rafał S Jurecki, and Rafał S Jurecki. Parameters influencing on electric vehicle
 679 range. *Energies*, 14(16):4821, 2021.

680 Paolo Toth and Daniele Vigo. A heuristic algorithm for the symmetric and asymmetric vehicle
 681 routing problems with backhauls. *European Journal of Operational Research*, 113(3):528–543,
 682 1999.

683 Thibaut Vidal. Hybrid genetic search for the cvrp: Open-source implementation and swap* neigh-
 684 borhood. *Computers & Operations Research*, 140:105643, 2022.

685 Daniele Vigo. A heuristic algorithm for the asymmetric capacitated vehicle routing problem. *Euro-*
 686 *pean Journal of Operational Research*, 89(1):108–126, 1996.

687 Mengqin Wang, Yanling Wei, Xueliang Huang, and Shan Gao. An end-to-end deep reinforcement
 688 learning framework for electric vehicle routing problem. *IEEE Internet of Things Journal*, 2024.

689 Yang Wang, Ya-Hui Jia, Wei-Neng Chen, and Yi Mei. Distance-aware attention reshaping for en-
 690 hancing generalization of neural solvers. *IEEE Transactions on Neural Networks and Learning*
 691 *Systems*, 2025.

692 Niels A. Wouda, Leon Lan, and Wouter Kool. PyVRP: a high-performance VRP solver package. *IN-*
 693 *FORMS Journal on Computing*, 2024. URL <https://doi.org/10.1287/ijoc.2023.0055>.

702 Yubin Xiao, Di Wang, Boyang Li, Mingzhao Wang, Xuan Wu, Changliang Zhou, and You Zhou.
703 Distilling autoregressive models to obtain high-performance non-autoregressive solvers for ve-
704 hicle routing problems with faster inference speed. In *Proceedings of the AAAI Conference on*
705 *Artificial Intelligence (AAAI)*, volume 38, pp. 20274–20283, 2024.

706 Liang Xin, Wen Song, Zhiguang Cao, and Jie Zhang. Multi-decoder attention model with embedding
707 glimpse for solving vehicle routing problems. In *Proceedings of the AAAI Conference on Artificial*
708 *Intelligence (AAAI)*, volume 35, pp. 12042–12049, 2021.

710 Haoran Ye, Jiarui Wang, Helan Liang, Zhiguang Cao, Yong Li, and Fanzhang Li. GLOP: Learning
711 global partition and local construction for solving large-scale routing problems in real-time. In
712 *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, volume 38, pp. 20284–
713 20292, 2024.

714 Manzil Zaheer, Guru Guruganesh, Kumar Avinava Dubey, Joshua Ainslie, Chris Alberti, Santiago
715 Ontanon, Philip Pham, Anirudh Ravula, Qifan Wang, Li Yang, et al. Big bird: Transformers for
716 longer sequences. *Advances in Neural Information Processing Systems (NeurIPS)*, 33:17283–
717 17297, 2020.

719 Zhi Zheng, Changliang Zhou, Tong Xialiang, Mingxuan Yuan, and Zhenkun Wang. Udc: A uni-
720 fied neural divide-and-conquer framework for large-scale combinatorial optimization problems.
721 *Advances in Neural Information Processing Systems (NeurIPS)*, 37:6081–6125, 2024.

722 Changliang Zhou, Xi Lin, Zhenkun Wang, Xialiang Tong, Mingxuan Yuan, and Qingfu Zhang.
723 Instance-conditioned adaptation for large-scale generalization of neural routing solver. *arXiv*
724 *preprint arXiv:2405.01906*, 2024.

725 Jianan Zhou, Yaxin Wu, Wen Song, Zhiguang Cao, and Jie Zhang. Towards omni-generalizable
726 neural methods for vehicle routing problems. In *International Conference on Machine Learning*
727 (*ICML*), pp. 42769–42789. PMLR, 2023.

729 Zefang Zong, Hansen Wang, Jingwei Wang, Meng Zheng, and Yong Li. Rbg: Hierarchically solving
730 large-scale routing problems in logistic systems via reinforcement learning. In *Proceedings of the*
731 *28th ACM Conference on Knowledge Discovery and Data Mining (SIGKDD)*, pp. 4648–4658,
732 2022.

733

734

735

736

737

738

739

740

741

742

743

744

745

746

747

748

749

750

751

752

753

754

755

756 **A PROBLEM FORMULATIONS OF RWVRPs**
 757

758 **A.1 VEHICLE ROUTING PROBLEM WITH REPLENISHMENT STOPS (VRP-RS)**
 759

760 The Vehicle Routing Problem with Replenishment Stops (VRP-RS) [Schneider et al., 2015] extends
 761 the traditional VRP by incorporating replenishment stops where vehicles can restock goods at in-
 762 termediate nodes to continue their service. This model accounts for real-world constraints such as
 763 maximum driving range limitations due to driver working hours or operational constraints.

764 Let \mathcal{V} define a set of vehicles, where each vehicle $w \in \mathcal{V}$ has a capacity Q_w . Each customer $n \in \mathcal{N}$
 765 has a demand q_n and $c \in \mathcal{C}$ shows a set of intermediate replenishment stops, which is a subset of
 766 $U = \{n_0, n_1, \dots, c_1, c_2\}$. The distance between any two nodes i and j , where $i, j \in U$, is shown as
 767 d_{ij} . Q defines the maximum vehicle capacity, while the maximum driving range is indicated by R .

768 **Decision variables:**
 769

- $x_{ij}^w \in \{0, 1\}$: A binary variable that takes the value 1 if vehicle w travels from node i to
 node j , and 0 otherwise.
- q_i^w : the load at node i delivered by vehicle w
- r_i^w : the remaining driving range of vehicle w after visiting node i

775 **Problem formulation:**

$$\text{Minimize: } \sum_{w \in \mathcal{V}} \sum_{i \in U} \sum_{j \in U} d_{ij} x_{ij}^w \quad (11)$$

778 Where,

780 All vehicles start and end at the depot:

$$\sum_{j \in U} x_{0j}^w = 1, \quad \sum_{i \in U} x_{i0}^w = 1, \quad \forall w \in \mathcal{V} \quad (12)$$

784 Every customer must be visited exactly once:

$$\sum_{w \in \mathcal{V}} \sum_{j \in U} x_{ij}^w = 1, \quad \forall i \in \mathcal{N} \quad (13)$$

788 At each node, the number of incoming and outgoing flows must match:

$$\sum_{i \in U} x_{ij}^w = \sum_{k \in U} x_{jk}^w, \quad \forall j \in U, \forall w \in \mathcal{V} \quad (14)$$

792 The capacity limit of each vehicle must be satisfied:

$$q_j^w \leq Q, \quad \forall j \in \mathcal{N}, \forall w \in \mathcal{V} \quad (15)$$

795 A vehicle's remaining capacity decreases after serving a customer:

$$q_j^w = q_i^w - q_j x_{ij}^w, \quad \forall i, j \in \mathcal{N}, \forall w \in \mathcal{V} \quad (16)$$

798 The total travel time of each vehicle must not exceed its maximum driving range:

$$0 \leq r_i^w \leq R, \quad \forall i \in U, \forall w \in \mathcal{V} \quad (17)$$

801 The remaining driving range is updated after each traversal:

$$r_j^w = r_i^w - d_{ij} x_{ij}^w, \quad \forall i, j \in U, \forall w \in \mathcal{V} \quad (18)$$

804 Traveling from the current node to a customer and then to the depot must not exceed the vehicle's
 805 driving range:

$$r_i^w \geq d_{ij} + d_{j0}, \quad \forall i, j \in U, \forall w \in \mathcal{V} \quad (19)$$

807 Visiting a replenishment stop fully restores the vehicle's capacity:

$$q_j^w = Q, \quad \forall j \in \mathcal{C}, \forall w \in \mathcal{V} \text{ such that } \sum_{i \in U} x_{ij}^w = 1 \quad (20)$$

810 A.2 ELECTRIC VEHICLE ROUTING PROBLEM WITH CHARGING STATIONS (EVRP-CS)
811812 The Electric Vehicle Routing Problem with Charging Stations (EVRP-CS) (Koç et al., 2019) extends
813 the VRP by adding a fleet of EVs that have limited driving ranges due to battery constraints. There
814 is also set of charging stations along the routes to maintain operational feasibility.815 We use the same notation as VRP-RS for problem formulation, with the key difference being that
816 intermediate stops $c \in \mathcal{C}$ now represent charging stations rather than replenishment stops. EVs must
817 visit these charging stations to recharge their batteries and continue their routes.818 **Decision variables:**
819820 • $x_{ij}^w \in \{0, 1\}$: binary variable that equals 1 if vehicle w travels from i to j , and 0 otherwise
821 • q_i^w : the load at node i for vehicle w
822 • r_i^w : the remaining battery range of vehicle w at node i 824 **Problem formulation:**
825

826 Minimize:
$$\sum_{w \in \mathcal{V}} \sum_{i \in U} \sum_{j \in U} d_{ij} x_{ij}^w \quad (21)$$

827

828 Where,

829 All vehicles start and end at the depot:

830
$$\sum_{j \in U} x_{0j}^w = 1, \quad \sum_{i \in U} x_{i0}^w = 1, \quad \forall w \in \mathcal{V} \quad (22)$$

831
832

833 Every customer must be visited exactly once:

834
$$\sum_{w \in \mathcal{V}} \sum_{j \in U} x_{ij}^w = 1, \quad \forall i \in \mathcal{N} \quad (23)$$

835
836

837 At each node, the number of incoming and outgoing flows must match:

838
$$\sum_{i \in U} x_{ij}^w = \sum_{k \in U} x_{jk}^w, \quad \forall j \in U, \forall w \in \mathcal{V} \quad (24)$$

839
840

841 The capacity limit of each vehicle must be satisfied:

842
$$q_j^w \leq Q, \quad \forall j \in \mathcal{N}, \forall w \in \mathcal{V} \quad (25)$$

843
844

845 A vehicle's remaining capacity decreases after serving a customer:

846
$$q_j^w = q_i^w - q_j x_{ij}^w, \quad \forall i, j \in \mathcal{N}, \forall w \in \mathcal{V} \quad (26)$$

847
848

849 Each vehicle's battery has a minimum and maximum capacity:

850
$$0 \leq r_i^w \leq R, \quad \forall i \in U, \forall w \in \mathcal{V} \quad (27)$$

851
852

853 The level of remaining battery is updated after traversing:

854
$$r_j^w = r_i^w - d_{ij} x_{ij}^w, \quad \forall i, j \in U, \forall w \in \mathcal{V} \quad (28)$$

855
856

857 Traveling from the current node to a customer and then to the depot must not deplete the vehicle's
858 battery:
859

860
$$r_i^w \geq d_{ij} + d_{jc}, \quad \forall i, j \in U, \forall w \in \mathcal{V}, \forall c \in \mathcal{C} \quad (29)$$

861
862

863 Whenever an EV visits a charging station, its battery is fully recharged:

864
$$r_j^w = R, \quad \forall j \in C, \forall w \in \mathcal{V} \text{ such that } \sum_{i \in U} x_{ij}^w = 1 \quad (30)$$

864 A.3 VEHICLE ROUTING PROBLEM WITH TIME WINDOWS (VRPTW)
865866 The Vehicle Routing Problem with Time Windows (VRPTW) (Kallehauge et al. 2005) extends VRP
867 by incorporating time constraints at customer locations. Each customer must be served within a
868 specified time window, making the problem more realistic for applications such as delivery services,
869 waste collection, and appointment scheduling where timing is crucial.870 We use the same notation as VRP-RS for problem formulation, with the key difference being that
871 each demand q_i must be served within a time window $[e_i, l_i]$, where e_i is the earliest service time
872 and l_i is the latest service time. The distance between any two nodes i and j is denoted as d_{ij} , with
873 an associated travel time t_{ij} . Each customer i requires a service time s_i .874 **Decision variables:**
875876 • $x_{ij}^w \in \{0, 1\}$: binary variable that equals 1 if vehicle w travels from node i to node j , and 0
877 otherwise
878 • T_i^w : arrival time of vehicle w at node i
879 • q_i^w : cumulative load of vehicle w after serving node i
880881 **Problem formulation:**

882 Minimize:
$$\sum_{w \in \mathcal{V}} \sum_{i \in U} \sum_{j \in U} d_{ij} x_{ij}^w \quad (31)$$

883

884 Where,

885 All vehicles start and end at the depot:

886
$$\sum_{j \in \mathcal{N}} x_{0j}^w = 1, \quad \sum_{i \in \mathcal{N}} x_{i0}^w = 1, \quad \forall w \in \mathcal{V} \quad (32)$$

887

888 Every customer must be visited exactly once:

889
$$\sum_{w \in \mathcal{V}} \sum_{j \in U} x_{ij}^w = 1, \quad \forall i \in \mathcal{N} \quad (33)$$

890

891 At each node, the number of incoming and outgoing flows must match:

892
$$\sum_{i \in U} x_{ij}^w = \sum_{k \in U} x_{jk}^w, \quad \forall j \in \mathcal{N}, \forall w \in \mathcal{V} \quad (34)$$

893

894 The capacity limit of each vehicle must be satisfied:

895
$$q_j^w = q_i^w + q_j x_{ij}^w, \quad \forall i, j \in U, \forall w \in \mathcal{V} \quad (35)$$

896

897
$$q_i^w \leq Q, \quad \forall i \in U, \forall w \in \mathcal{V} \quad (36)$$

898

899 Time window constraints must be satisfied:

900
$$e_i \leq T_i^w \leq l_i, \quad \forall i \in U, \forall w \in \mathcal{V} \quad (37)$$

901

902 Time consistency constraints must be satisfied:

903
$$T_j^w \geq T_i^w + s_i + t_{ij} - M(1 - x_{ij}^w), \quad \forall i, j \in U, \forall w \in \mathcal{V} \quad (38)$$

904

905 where M is a sufficiently large constant.

906 Depot time window is between the range:

907
$$e_0 \leq T_0^w \leq l_0, \quad \forall w \in \mathcal{V} \quad (39)$$

908

909 Decision variables must satisfy non-negativity:

910
$$x_{ij}^w \in \{0, 1\}, \quad T_i^w \geq 0, \quad q_i^w \geq 0 \quad (40)$$

911

A.4 ASYMMETRIC VEHICLE ROUTING PROBLEM (AVRP)

The Asymmetric Vehicle Routing Problem (AVRP) (Toth & Vigo 1999) extends the VRP by allowing the travel cost or distance from node i to j to differ from that of traveling from j to i . This asymmetry reflects real-world urban transportation scenarios, where factors such as traffic congestion, one-way streets, and temporarily closed roads can cause travel times between two nodes to vary depending on the direction.

The key distinction of AVRP with other extensions is that the distance matrix is asymmetric: $d_{ij} \neq d_{ji}$ for some pairs (i, j) . This creates a directed graph $G = (U, A)$ where A is the set of directed arcs.

Decision variables:

- $x_{ij}^w \in \{0, 1\}$: binary variable that equals 1 if vehicle w traverses arc (i, j)
- q_i^w : load of vehicle w after serving node i

Problem formulation:

$$\text{Minimize: } \sum_{w \in \mathcal{V}} \sum_{(i,j) \in A} d_{ij} x_{ij}^w \quad (41)$$

Where,

All vehicles start and end at the depot:

$$\sum_{j \in \mathcal{N}} x_{0j}^w = 1, \quad \sum_{i \in \mathcal{N}} x_{i0}^w = 1, \quad \forall w \in \mathcal{V} \quad (42)$$

Every customer visited exactly once:

$$\sum_{w \in \mathcal{V}} \sum_{i \in U, i \neq j} x_{ij}^w = 1, \quad \forall j \in \mathcal{N} \quad (43)$$

The capacity limit of each vehicle must be satisfied:

$$q_i^w = q_i^w + q_i x_{ij}^w, \quad \forall i, j \in U, \forall w \in \mathcal{V} \quad (44)$$

$$q_i^w \leq Q, \quad \forall i \in U, \forall w \in \mathcal{V} \quad (45)$$

Vehicle route continuity:

$$\sum_{i \in \mathcal{N}} x_{0j}^w \leq 1, \quad \sum_{i \in \mathcal{N}} x_{i0}^w \leq 1, \quad \forall w \in \mathcal{V} \quad (46)$$

$$\sum_{i \in U, i \neq j} x_{ij}^w = \sum_{k \in U, k \neq j} x_{jk}^w, \quad \forall j \in \mathcal{N}, \forall w \in \mathcal{V} \quad (47)$$

A.4.1 DATASET GENERATION

To generate asymmetric dataset instances, we implement a directional cost perturbation approach. We randomly select β (asymmetry scaling parameter) customers as origin nodes, then independently select β different customers as destination nodes. For each origin-destination pair, we introduce asymmetry by augmenting the forward travel distance by a random factor uniformly sampled from $[1, 1 + \gamma]$, where $\gamma = 0.2$ in our experiments, while keeping the reverse direction unchanged. This creates directional biases that mirror real-world scenarios such as one-way streets, traffic patterns, or elevation changes.

972
973

B PROACTIVE MASKING FUNCTION

974
975
976
977
978
979
980

A masking function is an essential component in approaches for VRP optimization to restrict the search space and prevent model from generating infeasible solutions. The definition of a masking function varies for each specific problem. In the EVRPCS, infeasibility is defined as selecting an already-visited customer, selecting a customer which violate cargo capacity or entails EV to be out of battery before reaching a CS. In the VRPRS, infeasibility occurs when a previously visited customer is visited again, a customer that violate cargo weight is open to select, or decisions that cause the vehicle to exceed its maximum allowable driving range before returning to the depot.

981
982
983
984
985
986

In standard VRP, at each decision step, we evaluate whether any of the remaining customers can be feasibly visited based on the vehicle’s current load and remaining capacity. However, VRPRS introduces additional complexity: a vehicle may lack sufficient driving range to reach any subsequent location after completing its current service, rendering the proposed solution infeasible. Similarly, in EVRPCS, the battery constraint adds another layer of feasibility checking beyond simple capacity constraints.

987
988
989
990
991
992

Previous approaches, such as those proposed by Wang et al. (2024), have introduced penalty functions to encourage the agent to autonomously explore the feasible domain and learn to generate valid solutions. While this penalty-based approach shows promise, it significantly expands the search space, resulting in longer convergence times during model training. Furthermore, even after complete training, such models may still produce infeasible solutions, compromising their practical reliability.

993
994
995
996
997
998

To address these limitations, we propose a proactive masking function that preemptively eliminates infeasible actions from the decision space. During the decoding step at time t , any customer or intermediate stop is masked if: (i) the customer has already been visited up to time $t - 1$, (ii) visiting the customer would cause the vehicle to exceed its capacity limit, or (iii) the vehicle’s remaining battery charge (for EVRPCS) or driving range (for VRPRS) is insufficient to reach the selected node and subsequently travel to the nearest charging station, replenishment stop, or depot.

999
1000
1001
1002
1003
1004
1005

This prevents the model from making locally feasible but globally infeasible decisions. While MT-POMO, MVMoE, and RouteFinder consider resource constraints for VRPL (vehicle routing with limited resources), they handle infeasibility through penalty functions during training rather than hard masking, meaning they can still generate infeasible solutions that require costly post-processing repair. In contrast, our proactive masking guarantees 100% feasibility by construction during inference, eliminating repair overhead and forcing the model to focus solely on viable routing decisions rather than relying on penalty-based correction.

1006
1007

B.1 PROACTIVE MASKING FUNCTION FOR EVRPCS

1008
1009
1010
1011
1012

For the EVRPCS, let $D = \{n_0\} \cup \mathcal{C}$ define the set containing both the depot and all charging stations, where \mathcal{C} represents the set of charging stations. Let \wedge denote the logical AND operator, $\phi \subset \mathcal{N}$ represent the set of visited customers, and $\mathcal{U}_t(j)$ indicate the masking function for visiting node j at time t when vehicle w is currently at node i . The masking function is defined as:

1013
1014
1015

$$\mathcal{U}_t(j) = \begin{cases} \text{False}, & j \notin \phi \wedge q_i^w + q_j \leq Q \wedge 1131 \\ \text{True}, & \text{Otherwise} \end{cases} \quad (48)$$

1016
1017
1018
1019

This formulation ensures that a node j is only selectable if: (1) it has not been visited, (2) serving it would not exceed vehicle capacity, and (3) the vehicle has sufficient battery to reach node j and then travel to the nearest charging facility or depot.

1020
1021

B.2 PROACTIVE MASKING FUNCTION FOR VRPRS

1022
1023
1024
1025

For the Vehicle Routing Problem with Replenishment Stops, let \mathcal{C} define the set of intermediate replenishment stops, n_0 denote the depot node, R represent the maximum driving range of the vehicle, and $\phi \subset \mathcal{N}$ show the set of visited customers. Let r_i^w denote the remaining driving range of vehicle w at node i . The lookahead masking function $\mathcal{U}_t(j)$ for VRPRS, which determines the feasibility of visiting node j at time t when vehicle w is at node i , is defined as:

$$1026 \quad 1027 \quad \mathcal{U}_t(j) = \begin{cases} \text{False,} & j \notin \phi \wedge q_i^w + q_j \leq Q \wedge r_i^w + d_{ij} + d_{j0} \leq R \\ 1028 \quad \text{True,} & \text{Otherwise} \end{cases} \quad 1029 \quad (49)$$

1030 This ensures that a customer can only be selected if visiting them and returning to the depot would
1031 not violate the driving range constraint.

1033 **C EXPERIMENTS ON VERY LARGE VRP PROBLEMS**

1034

1035

1036

1037 We evaluated SEAFormer’s scalability on the exceptionally large problem instances proposed
1038 by [Hou et al. (2022)], comprising 5,000 and 7,000 nodes in both VRP and RWVRP settings. The re-
1039 sults, presented in Table 4, show that SEAFormer consistently outperforms state-of-the-art methods
1040 across all RWVRP configurations and scales. The same holds for standard VRP, where SEAFormer
1041 surpasses UDC which is the leading divide-and-conquer approach for large-scale VRPs.

1042 Table 4: Experimental results on very large-scale VRP instances. The best overall performance is
1043 shown in bold, and the top learning-based method is shaded. We compare the best-performing base-
1044 lines that run without encountering out-of-memory errors in a reasonable time. Gaps are measured
1045 with respect to the best-performing approach.

1046

METHODS	5K CUSTOMERS			7K CUSTOMERS		
	OBJ.	G(%)	T	OBJ.	G(%)	T
LKH	175	26.7	4.3H	245	30.2	14H
TAM-LKH3	144.6	4.7	35M	196.9	4.67	1H
GLOP-LKH3	142.4	3.11	8M	191.2	1.64	10M
LEHD	140.7	1.88	3H	-	-	-
UDC ₂₅₀ ($\alpha = 1$)	139.0	0.65	15M	188.6	0.26	20M
SEAFormer	138.1	0.00	22M	188.1	0.00	34M

1053

1054 Table 5 presents results for very large-scale RWVRPs. SEAFormer demonstrates exceptional scal-
1055 ability on 5,000- and 7,000-customer instances, achieving the best solutions across all four problem
1056 variants while existing methods struggle. For VRPTW, it reduces objectives by 94.3% compared to
1057 MTPOMO on 5K instances and is the only method effectively solving 7K instances. In EVRPCS,
1058 SEAFormer outperforms specialized methods like EVGAT by 45% on 5K instances and maintains
1059 this advantage at 7K scale, with similar results for VRPRS. For AVRPs, it achieves a 5.57% im-
1060 provement over the state-of-the-art solver, highlighting SEAFormer’s strong generalizability across
1061 different problem types and scales.

1062

1063 Table 5: Performance of SEAFormer on a very large scale RWVRP instances. The best results are
1064 bolded while the best learning-based method is highlighted. Gaps are measured with respect to the
1065 best-performing approach.

1066

RWVRP	METHODS	5K CUSTOMERS			7K CUSTOMERS		
		OBJ.	G(%)	T	OBJ.	G(%)	T
VRPTW	POMO	997	67.2	28M	-	-	-
	MTPOMO	1158	94.3	30M	1656	110	54M
	SEAFormer	596	0.00	33M	786	0.00	65M
EVRPCS	EVRPRL	272.4	90.3	30M	-	-	-
	EVGAT	207.9	45.2	38M	296.8	50.5	75M
	SEAFormer	143.1	0.00	32M	197.1	0.00	65M
VRPRS	EVRPRL	193.2	64.9	30M	-	-	-
	EVGAT	167.3	42.8	38M	230.6	46.8	75M
	SEAFormer	117.1	0.00	32M	157.1	0.00	65M
AVRP	POMO	244	68	17M	-	-	-
	UDC ₂₅₀ ($\alpha = 1$)	154.8	7.05	14M	208.4	5.57	20M
	SEAFormer	144.6	0.00	21M	197.4	0.00	34M

1077

1078

1079

1080 **D EXPERIMENTS ON REAL-WORLD DATASETS**
10811082 **D.1 CVRPLIB DATASET**
10831084 On large-scale CVRPLib instances, SEAFormer demonstrates strong performance. Table 7 shows
1085 the gap of different methods relative to the best-known solution. When combined with SGBS, it
1086 surpasses state-of-the-art solutions; even without SGBS, SEAFormer is only 0.2% less effective on
1087 the largest instances, underscoring the robustness of the proposed approach.1088
1089 **Table 6: Gap to Best Known Solution on CVRPLib real-world Benchmark. The best overall perfor-**
1090 **mance is highlighted.**

Dataset	GLOP-LKH3	LEHD	UDC ₂₅₀	SEAFormer	SEAFormer-SGBS
Set-X(500, 1000)	16.8%	17.4%	7.1%	7.3%	6.8%
Set-XXL(1000, 10000)	19.1%	22.2%	13.2%	13.3%	-

1091 **D.2 EVRPCS**
10921093 We evaluate SEAFormer on the real-world EVRPCS benchmark dataset from Mavrovouniotis et al.
1094 (2020), which contains 24 instances with customer sizes ranging from 29 to 1,000 and 4–13 charging
1095 stations. We focus on the 16 larger instances with more than 100 customers to assess scalability on
1096 realistic problem sizes. Table 7 reports the gap to best-known solutions, showing that SEAFormer
1097 significantly outperforms existing learning-based methods. While EVGAT and EVRPRL reach gaps
1098 of 22.6% and 24.5%, respectively, SEAFormer with greedy decoding reduces the gap to 8.2%, and
1099 SEAFormer-SGBS further improves it to 6.5%.1100
1101 **Table 7: Gap to Best Known Solution on EVRPCS real-world Benchmark. The best overall perfor-**
1102 **mance is highlighted.**

Dataset	EVGAT	EVRPRL	SEAFormer	SEAFormer-SGBS
EVRPCS(100, 1000)	22.6%	24.5%	8.2%	6.5%

1111 **D.3 VRPTW**
11121113 We evaluate SEAFormer on the VRPTW benchmark (Gehring & Homberger), which includes 60
1114 problem instances with 1,000 customers. Table 8 shows that SEAFormer substantially outperforms
1115 existing neural baselines on these real-world large-scale instances. MTPOMO and POMO yield
1116 gaps of 56.1% and 47.3%, and the distance-aware DAR method reports a 21.4% gap. In contrast,
1117 SEAFormer with greedy decoding reduces the gap to 9.1%, and SEAFormer-SGBS attains an even
1118 smaller gap of 7.4%.1119
1120 **Table 8: Gap to Best Known Solution on VRPTW real-world Benchmark. The best overall perfor-**
1121 **mance is highlighted.**

Dataset	MTPOMO	POMO	DAR	SEAFormer	SEAFormer-SGBS
VRPTW(1000)	56.1%	47.3%	21.4%	9.1%	7.4%

1126 **E CROSS-DISTRIBUTION GENERALIZATION**
11271128 A key requirement for any neural combinatorial optimization (NCO) solver is the ability to gen-
1129 eralize beyond the data on which it was trained (Zheng et al., 2024). To evaluate this property,
1130 we assessed SEAFormer on two challenging out-of-distribution settings proposed by Zhou et al.
1131 (2023): the Rotation and Explosion distributions of the CVRP, each with 500 and 1,000 customers.
1132 As reported in Table 9 SEAFormer demonstrates consistent robustness and strong performance even
1133 when faced with large-scale instances that exhibit fundamentally different spatial structures.

1134 Table 9: Cross-distribution generalization on 128 instances from the dataset of Zhou et al. (2023).
1135 The best overall result is shown in bold, and the top learning-based result is highlighted.
1136

Method	Rotation			Explosion		
	Obj.↓	Gap	Time	Obj.↓	Gap	Time
HGS	32.97	0.00%	8h	32.87	0.00%	8h
POMO	64.76	96.4%	1m	58.17	76.9%	1m
Omni_VRP	35.9	8.8%	56.8m	35.65	8.45%	56.8m
ELG	37.31	13.16%	16.3m	36.53	8.1%	16.6m
UDC- $x_{250}(\alpha=1)$	35.14	6.58%	3.3m	35.11	6.81%	3.3m
SEAFormer	35.27	6.97%	1m	35.85	9.06%	1m
SEAFormer-SGBS	34.51	4.67%	100m	34.84	5.99%	100m

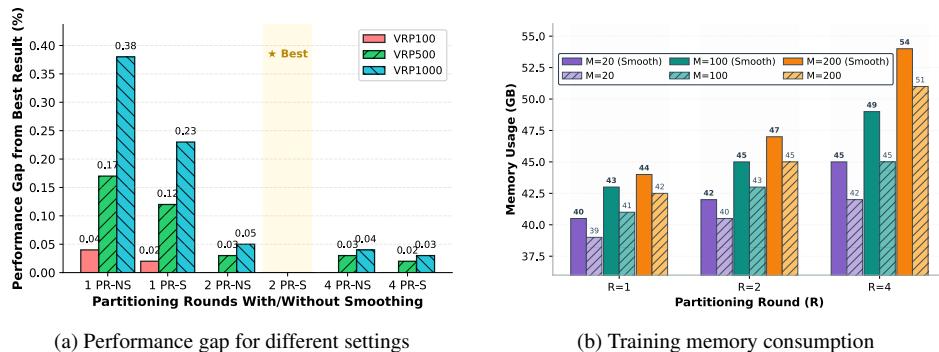
1144

1145 F ADDITIONAL ABLATION STUDY

1146 F.1 ABLATION ON CLUSTER SIZE AND PARTITIONING ROUNDS: IMPACT ON SOLUTION 1147 QUALITY AND MEMORY

1148 One of the main contributions of this paper is CPA with its partitioning rounds and smoothing
1149 techniques. Figure 5a demonstrates model accuracy under different partitioning and smoothing set-
1150 tings. The SEAFormer’s performance improves with increasing partitioning rounds. The accuracy
1151 gap between utilizing CPA with 1 partitioning round (PR) with and without smoothing is 0.15%
1152 for VRP1000, validating our smoothing approach’s effectiveness. Furthermore, the gap reduction
1153 from 0.38% to 0.05% demonstrates the power of our deterministic-yet-diverse partitioning strategy.
1154 Notably, performance across different partitioning rounds remains relatively consistent, highlight-
1155 ing SEAFormer’s architectural strength which achieves high-quality solutions by deterministically
1156 attending to small node groups.

1157 Performance gains come at the cost of higher memory usage. Figure 5b illustrates memory con-
1158 sumption during CVRP1k training with batch size 32 and pomo size 100. As anticipated, increasing
1159 partitioning rounds and enabling smoothing raise memory requirements, highlighting the trade-off
1160 between computational resources and solution quality.



1163
1164 Figure 5: (a) Performance gap of SEAFormer with different CPA configurations. (b) Training mem-
1165
1166
1167
1168
1169
1170
1171
1172
1173
1174
1175
1176
1177
1178
1179
1180
1181
1182
1183
1184
1185
1186
1187
1188
1189
1190
1191
1192
1193
1194
1195
1196
1197
1198
1199
1200
1201
1202
1203
1204
1205
1206
1207
1208
1209
1210
1211
1212
1213
1214
1215
1216
1217
1218
1219
1220
1221
1222
1223
1224
1225
1226
1227
1228
1229
1230
1231
1232
1233
1234
1235
1236
1237
1238
1239
1240
1241
1242
1243
1244
1245
1246
1247
1248
1249
1250
1251
1252
1253
1254
1255
1256
1257
1258
1259
1260
1261
1262
1263
1264
1265
1266
1267
1268
1269
1270
1271
1272
1273
1274
1275
1276
1277
1278
1279
1280
1281
1282
1283
1284
1285
1286
1287
1288
1289
1290
1291
1292
1293
1294
1295
1296
1297
1298
1299
1300
1301
1302
1303
1304
1305
1306
1307
1308
1309
1310
1311
1312
1313
1314
1315
1316
1317
1318
1319
1320
1321
1322
1323
1324
1325
1326
1327
1328
1329
1330
1331
1332
1333
1334
1335
1336
1337
1338
1339
1340
1341
1342
1343
1344
1345
1346
1347
1348
1349
1350
1351
1352
1353
1354
1355
1356
1357
1358
1359
1360
1361
1362
1363
1364
1365
1366
1367
1368
1369
1370
1371
1372
1373
1374
1375
1376
1377
1378
1379
1380
1381
1382
1383
1384
1385
1386
1387
1388
1389
1390
1391
1392
1393
1394
1395
1396
1397
1398
1399
1400
1401
1402
1403
1404
1405
1406
1407
1408
1409
1410
1411
1412
1413
1414
1415
1416
1417
1418
1419
1420
1421
1422
1423
1424
1425
1426
1427
1428
1429
1430
1431
1432
1433
1434
1435
1436
1437
1438
1439
1440
1441
1442
1443
1444
1445
1446
1447
1448
1449
1450
1451
1452
1453
1454
1455
1456
1457
1458
1459
1460
1461
1462
1463
1464
1465
1466
1467
1468
1469
1470
1471
1472
1473
1474
1475
1476
1477
1478
1479
1480
1481
1482
1483
1484
1485
1486
1487
1488
1489
1490
1491
1492
1493
1494
1495
1496
1497
1498
1499
1500
1501
1502
1503
1504
1505
1506
1507
1508
1509
1510
1511
1512
1513
1514
1515
1516
1517
1518
1519
1520
1521
1522
1523
1524
1525
1526
1527
1528
1529
1530
1531
1532
1533
1534
1535
1536
1537
1538
1539
1540
1541
1542
1543
1544
1545
1546
1547
1548
1549
1550
1551
1552
1553
1554
1555
1556
1557
1558
1559
1560
1561
1562
1563
1564
1565
1566
1567
1568
1569
1570
1571
1572
1573
1574
1575
1576
1577
1578
1579
1580
1581
1582
1583
1584
1585
1586
1587
1588
1589
1590
1591
1592
1593
1594
1595
1596
1597
1598
1599
1600
1601
1602
1603
1604
1605
1606
1607
1608
1609
1610
1611
1612
1613
1614
1615
1616
1617
1618
1619
1620
1621
1622
1623
1624
1625
1626
1627
1628
1629
1630
1631
1632
1633
1634
1635
1636
1637
1638
1639
1640
1641
1642
1643
1644
1645
1646
1647
1648
1649
1650
1651
1652
1653
1654
1655
1656
1657
1658
1659
1660
1661
1662
1663
1664
1665
1666
1667
1668
1669
1670
1671
1672
1673
1674
1675
1676
1677
1678
1679
1680
1681
1682
1683
1684
1685
1686
1687
1688
1689
1690
1691
1692
1693
1694
1695
1696
1697
1698
1699
1700
1701
1702
1703
1704
1705
1706
1707
1708
1709
1710
1711
1712
1713
1714
1715
1716
1717
1718
1719
1720
1721
1722
1723
1724
1725
1726
1727
1728
1729
1730
1731
1732
1733
1734
1735
1736
1737
1738
1739
1740
1741
1742
1743
1744
1745
1746
1747
1748
1749
1750
1751
1752
1753
1754
1755
1756
1757
1758
1759
1760
1761
1762
1763
1764
1765
1766
1767
1768
1769
1770
1771
1772
1773
1774
1775
1776
1777
1778
1779
1780
1781
1782
1783
1784
1785
1786
1787
1788
1789
1790
1791
1792
1793
1794
1795
1796
1797
1798
1799
1800
1801
1802
1803
1804
1805
1806
1807
1808
1809
1810
1811
1812
1813
1814
1815
1816
1817
1818
1819
1820
1821
1822
1823
1824
1825
1826
1827
1828
1829
1830
1831
1832
1833
1834
1835
1836
1837
1838
1839
1840
1841
1842
1843
1844
1845
1846
1847
1848
1849
1850
1851
1852
1853
1854
1855
1856
1857
1858
1859
1860
1861
1862
1863
1864
1865
1866
1867
1868
1869
1870
1871
1872
1873
1874
1875
1876
1877
1878
1879
1880
1881
1882
1883
1884
1885
1886
1887
1888
1889
1890
1891
1892
1893
1894
1895
1896
1897
1898
1899
1900
1901
1902
1903
1904
1905
1906
1907
1908
1909
1910
1911
1912
1913
1914
1915
1916
1917
1918
1919
1920
1921
1922
1923
1924
1925
1926
1927
1928
1929
1930
1931
1932
1933
1934
1935
1936
1937
1938
1939
1940
1941
1942
1943
1944
1945
1946
1947
1948
1949
1950
1951
1952
1953
1954
1955
1956
1957
1958
1959
1960
1961
1962
1963
1964
1965
1966
1967
1968
1969
1970
1971
1972
1973
1974
1975
1976
1977
1978
1979
1980
1981
1982
1983
1984
1985
1986
1987
1988
1989
1990
1991
1992
1993
1994
1995
1996
1997
1998
1999
2000
2001
2002
2003
2004
2005
2006
2007
2008
2009
2010
2011
2012
2013
2014
2015
2016
2017
2018
2019
2020
2021
2022
2023
2024
2025
2026
2027
2028
2029
2030
2031
2032
2033
2034
2035
2036
2037
2038
2039
2040
2041
2042
2043
2044
2045
2046
2047
2048
2049
2050
2051
2052
2053
2054
2055
2056
2057
2058
2059
2060
2061
2062
2063
2064
2065
2066
2067
2068
2069
2070
2071
2072
2073
2074
2075
2076
2077
2078
2079
2080
2081
2082
2083
2084
2085
2086
2087
2088
2089
2090
2091
2092
2093
2094
2095
2096
2097
2098
2099
2100
2101
2102
2103
2104
2105
2106
2107
2108
2109
2110
2111
2112
2113
2114
2115
2116
2117
2118
2119
2120
2121
2122
2123
2124
2125
2126
2127
2128
2129
2130
2131
2132
2133
2134
2135
2136
2137
2138
2139
2140
2141
2142
2143
2144
2145
2146
2147
2148
2149
2150
2151
2152
2153
2154
2155
2156
2157
2158
2159
2160
2161
2162
2163
2164
2165
2166
2167
2168
2169
2170
2171
2172
2173
2174
2175
2176
2177
2178
2179
2180
2181
2182
2183
2184
2185
2186
2187
2188
2189
2190
2191
2192
2193
2194
2195
2196
2197
2198
2199
2200
2201
2202
2203
2204
2205
2206
2207
2208
2209
2210
2211
2212
2213
2214
2215
2216
2217
2218
2219
2220
2221
2222
2223
2224
2225
2226
2227
2228
2229
2230
2231
2232
2233
2234
2235
2236
2237
2238
2239
2240
2241
2242
2243
2244
2245
2246
2247
2248
2249
2250
2251
2252
2253
2254
2255
2256
2257
2258
2259
2260
2261
2262
2263
2264
2265
2266
2267
2268
2269
2270
2271
2272
2273
2274
2275
2276
2277
2278
2279
2280
2281
2282
2283
2284
2285
2286
2287
2288
2289
2290
2291
2292
2293
2294
2295
2296
2297
2298
2299
2300
2301
2302
2303
2304
2305
2306
2307
2308
2309
2310
2311
2312
2313
2314
2315
2316
2317
2318
2319
2320
2321
2322
2323
2324
2325
2326
2327
2328
2329
2330
2331
2332
2333
2334
2335
2336
2337
2338
2339
2340
2341
2342
2343
2344
2345
2346
2347
2348
2349
2350
2351
2352
2353
2354
2355
2356
2357
2358
2359
2360
2361
2362
2363
2364
2365
2366
2367
2368
2369
2370
2371
2372
2373
2374
2375
2376
2377
2378
2379
2380
2381
2382
2383
2384
2385
2386
2387
2388
2389
2390
2391
2392
2393
2394
2395
2396
2397
2398
2399
2400
2401
2402
2403
2404
2405
2406
2407
2408
2409
2410
2411
2412
2413
2414
2415
2416
2417
2418
2419
2420
2421
2422
2423
2424
2425
2426
2427
2428
2429
2430
2431
2432
2433
2434
2435
2436
2437
2438
2439
2440
2441
2442
2443
2444
2445
2446
2447
2448
2449
2450
2451
2452
2453
2454
2455
2456
2457
2458
2459
2460
2461
2462
2463
2464
2465
2466
2467
2468
2469
2470
2471
2472
2473
2474
2475
2476
2477
2478
2479
2480
2481
2482
2483
2484
2485
2486
2487
2488
2489
2490
2491
2492
2493
2494
2495
2496
2497
2498
2499
2500
2501
2502
2503
2504
2505
2506
2507
2508
2509
2510
2511
2512
2513
2514
2515
2516
2517
2518
2519
2520
2521
2522
2523
2524
2525
2526
2527
2528
2529
2530
2531
2532
2533
2534
2535
2536
2537
2538
2539
2540
2541
2542
2543
2544
2545
2546
2547
2548
2549
2550
2551
2552
2553
2554
2555
2556
2557
2558
2559
2560
2561
2562
2563
2564
2565
2566
2567
2568
2569
2570
2571
2572
2573
2574
2575
2576
2577
2578
2579
2580
2581
2582
2583
2584
2585
2586
2587
2588
2589
2590
2591
2592
2593
2594
2595
2596
2597
2598
2599
2600
2601
2602
2603
2604
2605
2606
2607
2608
2609
2610
2611
2612
2613
2614
2615
2616
2617
2618
2619
2620
2621
2622
2623
2624
2625
2626
2627
2628
2629
2630
2631
2632
2633
2634
2635
2636

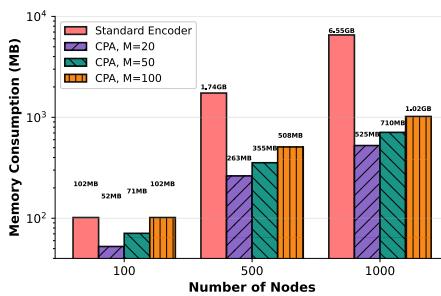


Figure 6: Logarithmic training memory consumption of CPA encoder on a 1,000-node VRP with batch size 32 and pomo size 100 under different configurations.

cluster sizes (20 and 50). In this evaluation, we vary the number of partitioning rounds (PR) and compare settings with smoothing (S) and without smoothing (NS).

The baseline POMO method achieves a 1.19% gap after 1000 epochs. In contrast, our CPA variants consistently deliver stronger performance, with several notable trends. First, the smoothing technique plays a critical role: comparing one partitioning round with and without smoothing reveals gap reductions of 0.51% and 0.45% at epoch 1000 for cluster sizes 20 and 50, respectively. Second, increasing the number of partitioning rounds leads to substantial improvements. For example, moving from one to four rounds with smoothing decreases the final gap from 1.19% to 0.69% for cluster size 20, and from 1.00% to 0.56% for cluster size 50.

We further observe that larger cluster sizes consistently outperform smaller ones across all settings. In particular, cluster size 50 with four partitioning rounds and smoothing achieves the best result, reaching a 0.56% gap. Beyond final performance, CPA also exhibits faster convergence: with four partitioning rounds and cluster size 50, the gap at epoch 100 is already 2.46%, outperforming POMO’s performance at epoch 500. This combination of accelerated convergence and superior asymptotic performance highlights the importance of both multiple partitioning rounds and well-chosen cluster sizes within the CPA framework.

Table 10: Performance gap relative to best achieved objective on VRP100 across training epochs with different CPA configurations (PR: Partitioning Rounds, S: Smoothing, NS: No Smoothing)

Method	Cluster size	epoch		Cluster size	epoch		
		100	500		100	500	1000
POMO	-	4.16%	1.89%	1.19%	-	-	-
1 PR-NS	20	3.97%	2.39%	1.7%	50	3.65%	2.01%
1 PR-S	20	3.65%	2.14%	1.19%	50	3.53%	1.57%
2 PR-S	20	3.4%	1.51%	0.88%	50	3.15%	1.38%
4 PR-S	20	2.58%	1.19%	0.69%	50	2.46%	1.07%

F.3 ABLATION ON EDGE MODULE EFFECT ON AVR

Table 11 reports the performance gap of SEAFormer without the edge module relative to the full SEAFormer. Across AVR instances of varying sizes, the gap ranges from 2.4% to 3.4%, highlighting the significant contribution of the edge module to SEAFormer’s performance in asymmetric routing settings.

Table 11: Gap between SEAFormer without Edge Embedding (SEAFormer-WOE) and the full SEAFormer on AVR.

Method	AVRP100	AVRP500	AVRP1000
SEAFormer-WOE	2.4%	2.7%	3.4%

1242
1243

F.4 ABLATION ON PARAMETER K IN EAM: IMPACT ON SOLUTION QUALITY AND MEMORY

1244
1245
1246
1247
1248
1249
1250
1251
1252
1253
1254
1255
1256
1257

We evaluate the sensitivity of SEAFormer to the number of nearest neighbors K used in the EAM, testing $K \in \{20, 50, 100, 200, 300\}$ on VRP-1000 and VRPTW-1000 benchmarks. Table 12 shows that performance improves monotonically as K increases, but with strongly diminishing returns beyond $K=50$. For VRP-1000, increasing K from 20 to 50 yields 0.34% improvement (from 43.90 to 43.75), while further increasing to 300 provides only an additional 0.32% gain (from 43.75 to 43.61). Similarly, on VRPTW-1000, the improvement from $K = 20$ to $K = 50$ is 0.39% (from 150.28 to 149.70), compared to just 0.29% from $K = 50$ to $K = 300$ (from 149.70 to 149.26). Moreover, memory consumption increases approximately linearly with K , from 2,054 MB at $K = 20$ to 3,326 MB at $K = 300$ (a 62% increase), primarily due to the larger edge embedding matrix stored during inference. These results demonstrate that $K = 50$ captures the majority of relevant edge information, while achieving favorable memory efficiency (only 2,255 MB, a modest 10% increase over $K = 20$). Larger K values provide marginal quality improvements ($<0.3\%$) at significant memory and computational cost. We therefore recommend $K = 50$ as the default configuration, balancing solution quality and memory efficiency.

1258
1259
1260

Table 12: Impact of edge module nearest neighbors K on solution quality and memory consumption across problem variants. Objective values and memory usage are averaged over 100 test instances with 1,000 customers, where each instance is solved individually.

1261
1262
1263
1264
1265

K	VRP-1000	VRPTW-1000	Memory (MB)
20	43.90	150.28	2,054
50	43.75	149.70	2,255
100	43.67	149.59	2,495
200	43.62	149.32	2,870
300	43.61	149.26	3,326

1266
1267

F.5 ABLATION ON CPA AND EAM: INFLUENCE ON MEMORY

1268
1269
1270
1271
1272
1273
1274
1275
1276
1277

We evaluate the training memory consumption of each SEAFormer component on VRP-1000 (batch size 32, POMO size 100), averaging results over 10 independent training runs. Our experiment results shown in Table 13 reveals that full SEAFormer with $R=1$, $M=50$ consumes 66 GB GPU memory, while SEAFormer-NoCPA (using standard $O(n^2)$ encoder instead of CPA) requires 75 GB, and SEAFormer-NoEAM (CPA only with $R=2$, $M=50$) consumes 60 GB. These results demonstrate that CPA provides substantial memory reduction (66 GB vs. 75 GB, a 12% saving), while the edge-aware module adds modest overhead (66 GB vs. 60 GB, a 10% increase), confirming that CPA’s $O(n)$ complexity is the primary enabler of large-scale training, with the edge module providing critical performance benefits (Figure 4, Table 11) at acceptable memory cost.

1278
1279
1280
1281

Table 13: Impact of proposed architectural components on SEAFormer training memory consumption across different configurations. The batch, POMO, and problem size is set to 32, 100, and 1000, respectively.

1282
1283

Configuration	SEAFormer	SEAFormer-NoCPA	SEAFormer-NOEAM
Memory	66 GB	75 GB	60 GB

1284

1285

G INTEGRATING STANDARD VRP METHODS ON RWVRPs

1286

As discussed earlier, EVRPCS and VRPRS impose state-dependent feasibility conditions such as battery level and remaining travel time, which are absent from standard VRP formulations. As a result, most existing neural or heuristic VRP solvers cannot be directly applied to EVRP-type problems, and many do not support EVRPCS or VRPRS at all.

1287
1288
1289
1290
1291

To ensure a fair comparison, we have applied a consistent procedure to generate valid EVRPCS solutions from existing VRP methods. Specifically, we first run UDC and HGS as standard VRP solvers, ignoring battery constraints. We then identify any route where a vehicle would run out of battery and apply a simple repair mechanism (UDC-R / HGS-R). For each vehicle in the UDC or HGS solution, and at each customer, we check whether it can reach the next customer and subsequently a charging

1296 station. If both are feasible, we follow the planned route; otherwise, we redirect the vehicle to the
 1297 nearest charging station before continuing. Please note that, since HGS does not natively support
 1298 EVRPCS, we apply this repair strategy to make it compatible with the EVRPCS setting.
 1299

1300 Table 14: Infeasibility rate and gap to SEAFormer of using VRP solutions on EVRPCS instances.
 1301 The infeasibility rate indicates the percentage of solutions that are invalid due to battery depletion
 1302 of vehicle during delivery.

Method	Size	Infeasibility rate	Objective	Gap vs. SEAFormer
HGS	100	27%	-	-
HGS-R	100	0% (Repaired)	16.38	0.03%
HGS	1000	100%	-	-
HGS-R	1000	0% (Repaired)	51.8	13%
UDC	1000	100%	-	-
UDC-R	1000	0% (Repaired)	50.9	11%

1310 As shown in Table 14, VRP solvers with repair strategies yield sub-optimal solutions because they
 1311 rely on greedy local insertion of charging stops, which often delays charging until it is too late and
 1312 forces vehicles to take long detours for recharging before they can resume deliveries. An optimal
 1313 charging decision, however, requires global awareness of the entire route, not just local context.
 1314 Since VRP solvers optimize primarily for distance, they cannot account for battery-dependent fea-
 1315 sibility when choosing the next node. When charging stations are inserted after route construction,
 1316 the spatial coherence of the route is disrupted. For example, a VRP route $A \rightarrow B \rightarrow C$ may become
 1317 $A \rightarrow B \rightarrow$ charging $\rightarrow C$, while the optimal EVRPCS solution would be $A \rightarrow$ charging $\rightarrow B \rightarrow C$.
 1318

H OPTIONAL NODES AND THEIR EMBEDDINGS

1321 SEAFormer treats optional nodes differently from customer nodes for three key reasons. First, the
 1322 number of optional nodes can change across problem instances, and mixing them with customers
 1323 could make the customer embeddings unstable. Second, because CPA fixes the number of customers
 1324 per cluster, adding variable optional nodes would require extra padding, which wastes resources.
 1325 Third, CPA groups nearby nodes together, so if optional nodes were embedded with customers,
 1326 some clusters might miss important optional node information, which could hurt solution quality.

1327 For these reasons, we apply self-attention among all optional nodes to capture their inter-
 1328 relationships. These embeddings are then integrated with customer node representations through
 1329 a Gumbel-Softmax [Jang et al., 2016] and learnable fusion mechanism. During training, we imple-
 1330 ment a temperature annealing schedule: τ initializes at 1.0 to encourage exploration and decays by
 1331 a factor of 0.99 per epoch until reaching 0.2. Through this mechanism, customer embeddings learn
 1332 to initially explore all optional facilities without bias, then progressively concentrate on the most
 1333 relevant service nodes as determined by spatial proximity and emerging routing patterns.

I CPA PARTITIONING ROUND VISUALIZATION

1337 Figure 7 demonstrates how CPA employs 4 rounds of partitioning with boundary smoothing to gen-
 1338 erate diverse spatial patterns. This approach preserves the global perspective of problem instances
 1339 while achieving significant reductions in memory usage and computational cost.

J ADDITIONAL RELATED WORK

J.1 CLASSICAL APPROACHES AND THEIR LIMITATIONS

1345 Classical (meta-)heuristic methods for RWVRPs follow a construction-destruction-improvement
 1346 paradigm. Examples include variable neighborhood search with Tabu search for load-dependent
 1347 energy consumption [Schneider et al., 2014], ant colony optimization with look-ahead for EVR-
 1348 PCS [Mavrovouniotis et al., 2018], simulated annealing for partial recharging [Felipe et al., 2014],
 1349 and adaptive large neighborhood search with load-aware power estimation [Kanchanla & Ramadu-
 rai, 2018]. Although these approaches yield high-quality solutions and handle complex constraints,

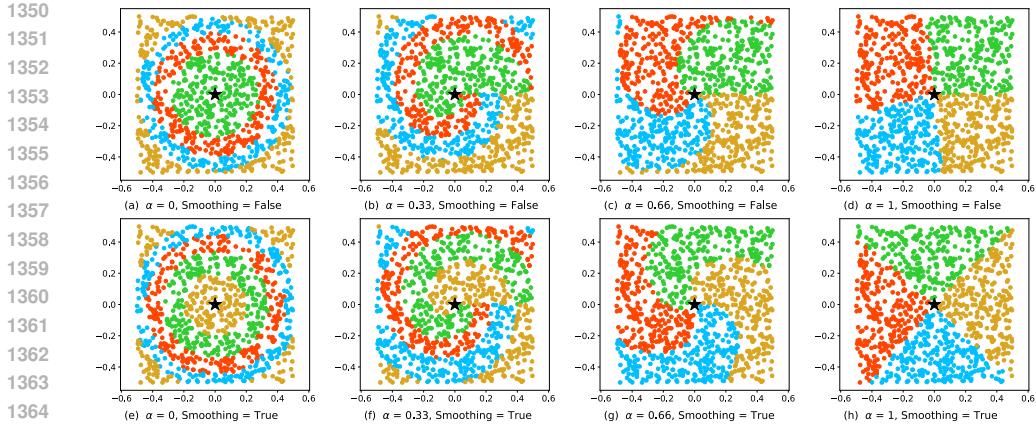


Figure 7: Visualization of Clustered Proximity Attention with $\alpha = [0, 0.33, 0.66, 1]$, where nodes are partitioned into $K=4$ clusters using polar coordinates centered at the depot (black star). Top row: Without boundary smoothing, hard boundaries can separate nearby nodes. Bottom row: With smoothing, cluster boundaries shift to maintain local neighborhoods. Colors indicate cluster assignment; attention is computed only within same-colored groups plus depot. The mixing parameter α interpolates between radial ($\alpha=0$) and angular ($\alpha=1$) clustering, capturing different routing patterns.

they face two main drawbacks: (i) solution time grows exponentially with problem size (Hou et al. 2022), and (ii) they struggle to generalize, as solutions cannot leverage patterns learned from previous instances. Such limitations make them unsuitable for real-time logistics, where solutions must be produced in seconds rather than hours.

J.2 LEARNING-BASED APPROACHES

The intersection of machine learning and combinatorial optimization has yielded approaches that directly produce high-quality VRP solutions without iterative refinement. These methods fall into two architectural paradigms: (1) Autoregressive models that build solutions sequentially, adding one decision at a time (Khalil et al. 2017; Kool et al. 2018; Kwon et al. 2020; Hou et al. 2022; Luo et al. 2025a, 2023; Zheng et al. 2024; Berto et al. 2025; Nasehi et al. 2025), and (2) Non-autoregressive models that generate complete solutions simultaneously, typically through learned heatmaps (Nowak et al. 2018; Kool et al. 2022; Ye et al. 2024; Xiao et al. 2024). Training paradigms include supervised learning from optimal solutions or reinforcement learning to directly optimize solution quality (Joshi et al. 2019).

Early work by Bello et al. (2016) pioneered learned heuristics using pointer networks trained with actor-critic methods. Nazari et al. (2018) enhanced this framework by incorporating attention mechanisms into the encoder. The field advanced significantly when Kool et al. (2018) applied Transformers to routing problems, demonstrating strong performance across TSP and VRP tasks. Subsequent improvements include POMO (Kwon et al. 2020), which introduced multiple rollouts for better exploration, and multi-decoder architectures (Xin et al. 2021) for enhanced solution refinement. Despite these advances, scaling to large problem instances remains computationally prohibitive.

Divide-and-conquer strategies have emerged as the dominant approach for large-scale instances (Nowak et al. 2018; Li et al. 2021; Zong et al. 2022; Fu et al. 2021; Zheng et al. 2024; Nasehi et al. 2025). TAM (Hou et al. 2022) employs two-stage decomposition followed by LKH3 (Helsgaun 2017) for sub-problem resolution. GLOP (Ye et al. 2024) integrates global partitioning through non-autoregressive models with local autoregressive construction. UDC (Zheng et al. 2024) addresses sub-optimal partitioning through robust training procedures, combining GNN-based decomposition with specialized sub-problem solvers.

Recently, heavy decoder architectures have emerged as another approach for achieving high-quality results. Luo et al. (2023) proposed shifting computational complexity from the encoder to the de-

1404 coder, introducing a partial reconstruction approach to enhance model accuracy while maintaining
 1405 reasonable training times through supervised learning. [Luo et al. \(2025a\)](#) extended this with
 1406 a boosted heavy decoder variant that restructures attention computation through two intermediate
 1407 nodes, reducing attention complexity to $O(2n)$. This efficiency gain enables training on signifi-
 1408 cantly larger instances by computing attention from each node to two pivot nodes, then from these
 1409 pivots to all other nodes, rather than computing full pairwise attention. [Huang et al. \(2025\)](#) provide
 1410 insights into models with light encoders and identify their weaknesses. They propose RELD, incor-
 1411 porating simple modifications such as adding identity mapping and a feed-forward layer to enhance
 1412 the decoder’s capacity.

1413
 1414 **J.3 LOCAL ATTENTION MECHANISM**

1415 Existing local-attention ([Fang et al. \(2024\)](#), [Gao et al. \(2024\)](#), [Wang et al. \(2025\)](#), [Li et al. \(2025\)](#), [Zhou](#)
 1416 [et al. \(2024\)](#)) approaches operate primarily in the decoder and modify attention weights based on
 1417 distance, aiming to speed up decoding and improve generalization in neural VRP solutions. While
 1418 these techniques can improve scalability on simple VRPs, modifying decoding scores introduces
 1419 a strong locality bias that may degrade solution quality for complex RWVRPs, such as VRPTW,
 1420 EVRPCS, VRPRS, and AVRPs, where feasible actions often depend on long-range decisions and
 1421 state-dependent constraints. In contrast, CPA preserves the decoder structure entirely and instead
 1422 introduces a pattern-based inductive bias in the encoder, enabling efficient training at scale without
 1423 sacrificing solution quality.

1424
 1425 **K HYPERPARAMETERS**

1426 Following previous work ([Kwon et al. \(2020\)](#)), we sample depot and customer coordinates uniformly
 1427 from $[0, 1]^2$ space, with customer demands drawn uniformly from $\{1, \dots, 10\}$. Vehicle capacities are
 1428 set to 50, 100, and 200 units per [Zhou et al. \(2023\)](#). For each epoch, we generate 10,000 training
 1429 instances on-the-fly. We train models for 2,000 epochs on 100-customer instances, 200 epochs
 1430 on 500-customer instances, and 100 epochs on 1,000-customer instances for each problem variant.
 1431 Training employs batch size 64 with POMO size 100 ([Kwon et al. \(2020\)](#)). Our architecture uses
 1432 a 6-layer encoder with 8 attention heads, embedding dimension 128, and feedforward dimension
 1433 512. The learning rate remains fixed at 10^{-4} . We leverage Flash Attention ([Dao et al. \(2022\)](#)) in our
 1434 multi-head attention mechanisms. The edge module uses 32-dimensional embeddings, computing
 1435 edge representations only between each node and its 50 nearest neighbors as a fixed constraint.
 1436 All remaining hyperparameters, training algorithms, and loss functions follow POMO ([Kwon et al.](#)
 1437 [2020](#)) specifications.

1438 **VRPTW.** we follow the time window generation procedure from [Liu et al. \(2024\)](#). Service times
 1439 and time window lengths are uniformly sampled from $[0.15, 0.2]$, representing normalized time
 1440 units relative to a planning horizon of $T=4.6$.

1441 **EVRPCS.** we randomly select 4 charging stations, and each EV can travel up to 2 units before
 1442 needing to recharge.

1443 **VRPRS.** we set 5 replenishment stops, with each vehicle able to travel a maximum of 4 units.

1444 **AVRP.** the asymmetry scaling parameter β is set to 50, 200, and 400 for 100-, 500-, and 1,000-
 1445 customer problems, respectively, while the directional bias factor γ is fixed at 0.2 for all instance
 1446 sizes.

1447
 1448 **L MODEL PARAMETERS AND TRAINING TIME**

1449 We provide a comprehensive comparison of model parameters and training costs across SEAFormer
 1450 and state-of-the-art baselines in Table 15. SEAFormer-Full contains 1.36M trainable parameters,
 1451 with the edge-aware module contributing only 95K parameters (7.5% increase over SEAFormer-
 1452 NoEAM’s 1.27M), demonstrating its lightweight design. Compared to baselines, SEAFormer main-
 1453 tains competitive parameter efficiency while requiring substantially less training time: it uses similar
 1454 capacity to UDC (1.51M) but trains in 87 hours versus UDC’s 140+ hours (excluding partitioner pre-
 1455 training), and it achieves superior performance with only 21% of EFormer’s parameters (6.39M) and

1458 29% of UniteFormer’s parameters (4.74M) while training 5x faster (87H vs. 390-440H). Moreover,
 1459 the edge-aware module adds modest computational overhead (+10% epoch time: 55s vs. 50s for
 1460 100-node instances).

1461

1462 Table 15: Comparison of trainable parameters across neural VRP methods. SEAFormer achieves
 1463 competitive parameter efficiency while outperforming heavier architectures.

1464

Method	Parameters	Total Epochs	Epoch Size	Epoch time	Total training time
SEAFormer-NoEAM	1.27M	2300	10,000	50s, 9M, 20 M	80 H
SEAFormer-Full	1.36M	2300	10,000	55s, 10M, 21M	87H
UDC	1.51M	200	1000	42 M	140 H + Partitioner training time
UniteFormer	4.74M	1010	100,000	23 M	390 H
EFormer	6.39M	1010	100,000	25 M	440 H

1465

1466

M BASELINES, CODES AND LICENSES

1471 As noted earlier, many existing methods cannot be directly applied to all RWVRP variants. For
 1472 EVRPCS, we use two learning-based approaches from the literature, following the implementation
 1473 details provided in their original papers. To our knowledge, no learning-based solutions exist for
 1474 VRPRS, so we adapted the EVRPCS methods, given their similarity in implementation and ap-
 1475 proach, to VRPRS and retrained them. For AVRPs, we use baseline methods to generate solutions
 1476 and then recompute route costs using the asymmetric distance matrix. Reported runtimes reflect
 1477 only model execution, excluding post-processing or cost recalculation. We keep all hyperparam-
 1478 eters at the defaults specified by the original authors. We also include recent concurrent works from
 1479 arXiv, as well.

1480 **OR-Tools.** We configure OR-Tools using the PATH_CHEAPEST_ARC strategy to generate initial
 1481 solutions and apply GUIDED_LOCAL_SEARCH as the local search metaheuristic. The runtime is
 1482 adapted to problem size, selecting from 30, 180, or 360 seconds per instances. The results reported
 1483 for OR-Tools are not the optimal, as the process was stopped once the time limit was reached.

1484 **HGS.** We use the HGS (Helsgaun [2017]) implementation in PyVRP (Wouda et al. [2024]) version
 1485 0.8.2, setting the neighborhood size to 50, minimum population to 50, and generation size to 100.
 1486 Runtime is adjusted by problem size, choosing 20, 120, or 240 seconds per instance. All other
 1487 parameters remain at default. The results reported for HGS are not the optimal, as the process was
 1488 stopped once the time limit was reached.

1489 The available codes and their licenses for used in this work are listed in Table 16.

1490

1491

Table 16: Resources for TSP and Optimization Problems

1492

1493

1494

Resource	Link	License
LKH3 (Helsgaun 2017)	http://webhotel4.ruc.dk/keld/research/LKH-3/	Academic research use
HGS (Vidal 2022)	https://github.com/chkwon/PyHygese	MIT License
PyVRP (Wouda et al. 2024)	https://github.com/PyVRP/PyVRP	MIT License
POMO (Kwon et al. 2020)	https://github.com/ya-kwon/POMO	MIT License
LEHD (Luo et al. 2023)	https://github.com/CIAM-Group/NCO_code/tree/main/single_objective/LEHD	MIT License
BLEHD (Luo et al. 2025a)	https://github.com/CIAM-Group/SIL	MIT License
ELG (Gao et al. 2024)	https://github.com/lamda-bbo/ELG	MIT License
GLOP (Ye et al. 2024)	https://github.com/henry-yeh/GLOP	MIT License
DeepMDV (Inasem et al. 2025)	https://github.com/SaeedNB/DeepMDV	MIT License
RELD (Huang et al. 2025)	https://github.com/zizweileonhuang/reld-nco	-
EFormer (Meng et al. 2025)	https://github.com/Regina921/EFormer	MIT License
UnitFormer (Meng et al. 2025)	https://github.com/Regina921/UnitFormer	MIT License
UDC (Zheng et al. 2024)	https://github.com/CIAM-Group/NCO_code/tree/main/single_objective/UDC-Large-scale-CO-master	MIT License

1507

1508

1509

N MDVRP EXTENSIBILITY

1510

1511

While our current implementation uses polar coordinates centered on a single depot, CPA can be
 naturally extended to Multi-depot problems (MDVRP) through several straightforward adaptations:

1512 **i) Unified depot representation:** Designate one depot as the primary reference point and embed
1513 all customers and secondary depots based on their spatial locations relative to this main depot. This
1514 approach requires minimal modification to our current implementation and treats additional depots
1515 as special customer nodes with supply capabilities, just like the VRPRS variant with replenishment
1516 stops. **ii) Depot-specific partitioning:** Compute polar coordinates relative to each depot indepen-
1517 dently and apply CPA within each depot’s customer assignment. In this variant, instead of using
1518 multiple partitioning rounds with different α values, each round computes customer embeddings
1519 with respect to a different depot, creating depot-specific spatial perspectives that naturally capture
1520 which customers are best served by which depot while maintaining the same computational com-
1521 plexity, where now corresponds to the number of depots. The key insight is that CPA exploits radial
1522 spatial structure around central points; whether that center is a single depot or multiple depots does
1523 not change the underlying geometric patterns that we show in Section 3.
1524

O USE OF LARGE LANGUAGE MODELS

1525 We use LLM assistance for grammar and presentation improvements. The original text and ideas
1526 remain the authors’ work, and they take full responsibility for the content.
1527
1528
1529
1530
1531
1532
1533
1534
1535
1536
1537
1538
1539
1540
1541
1542
1543
1544
1545
1546
1547
1548
1549
1550
1551
1552
1553
1554
1555
1556
1557
1558
1559
1560
1561
1562
1563
1564
1565