

Towards Generalizable Neural Solvers for Vehicle Routing Problems via Ensemble with Transferrable Local Policy

Chengrui Gao^{1,2}, Haopu Shang^{1,2}, Ke Xue^{1,2}, Dong Li³ and Chao Qian^{1,2}

¹National Key Laboratory for Novel Software Technology, Nanjing University, China

²School of Artificial Intelligence, Nanjing University, China

³Huawei Noah's Ark Lab

{gaocr, shanghp, xuek}@lamda.nju.edu.cn, lidong106@huawei.com, qianc@nju.edu.cn

Abstract

Machine learning has been adapted to help solve NP-hard combinatorial optimization problems. One prevalent way is learning to construct solutions by deep neural networks, which has been receiving more and more attention due to the high efficiency and less requirement for expert knowledge. However, many neural construction methods for Vehicle Routing Problems (VRPs) focus on synthetic problem instances with specified node distributions and limited scales, leading to poor performance on real-world problems which usually involve complex and unknown node distributions together with large scales. To make neural VRP solvers more practical, we design an auxiliary policy that learns from the local transferable topological features, named local policy, and integrate it with a typical construction policy (which learns from the global information of VRP instances) to form an ensemble policy. With joint training, the aggregated policies perform cooperatively and complementarily to boost generalization. The experimental results on two well-known benchmarks, TSPLIB and CVRPLIB, of travelling salesman problem and capacitated VRP show that the ensemble policy significantly improves both cross-distribution and cross-scale generalization performance, and even performs well on real-world problems with several thousand nodes.

1 Introduction

Vehicle Routing Problems (VRPs) are a kind of canonical NP-hard combinatorial optimization problems, with broad applications in logistics [Konstantakopoulos *et al.*, 2022]. They have been widely studied for decades, and various heuristic methods are proposed, such as LKH [Helsgaun, 2000; Helsgaun, 2017], SISR [Christiaens and Vanden Berghe, 2020], and HGS [Vidal, 2022]. However, existing heuristic methods are highly domain-specific, require lots of expert knowledge, and suffer from low efficiency in solving large-scale instances, which limit their application.

Recently, machine learning has been adapted to solve combinatorial optimization problems [Bengio *et al.*, 2021;

Cappart *et al.*, 2023], which utilizes deep neural networks to extract features, learn heuristics, and generate solutions, named Neural Combinatorial Optimization (NCO) [Bello *et al.*, 2017]. Depending on the training scheme, NCO includes Supervised Learning (SL) [Joshi *et al.*, 2019; Chen *et al.*, 2023] and Reinforcement Learning (RL) methods [Kool *et al.*, 2019]. SL-based methods usually require numerous (nearly) optimal solutions as labels, which are hardly applicable if optimal solutions are expensive to obtain [Bello *et al.*, 2017]. RL-based methods train the model using objective function evaluations as rewards, which are effective in most scenarios [Zhang *et al.*, 2020; Song *et al.*, 2023].

According to the process of generating solutions, RL-based methods can be generally divided into improvement and construction methods. Improvement methods learn to iteratively improve the quality of solutions, e.g., by selecting appropriate operators [Lu *et al.*, 2019]. They have the advantage of exploiting human knowledge, but suffer from limited search efficiency [Ma *et al.*, 2023] and high latency in inference [Sun and Yang, 2023]. Construction methods learn to extend a partial solution sequentially until a complete solution is constructed. This procedure can be modeled as a Markov Decision Process (MDP) and solved by RL [Bello *et al.*, 2017; Khalil *et al.*, 2017]. Construction methods are very computationally efficient, which can solve hundreds of instances in a few seconds. However, most construction methods only work well on problems with simple node distributions and limited scales. Their performance deteriorates dramatically on cross-distribution and cross-scale problems that are common in practice [Joshi *et al.*, 2022; Jiang *et al.*, 2022].

Towards generalizing to cross-distribution and cross-scale problem instances, several methods have been proposed recently. Some of them focus on generalization to large-scale problems. They often utilize the divide-and-conquer framework that can extend small-scale models to generalizable ones [Fu *et al.*, 2021; Hou *et al.*, 2023; Cheng *et al.*, 2023]. Meanwhile, the inductive bias of symmetries [Drakulic *et al.*, 2023] and some novel neural architectures, such as heavy decoder [Luo *et al.*, 2023] and diffusion models [Sun and Yang, 2023], are found to be useful for cross-scale generalization. The other methods aim to address the cross-distribution issue. A simple solution is directly training the model on instances with various node distributions, which is effective but can be further improved by applying some advanced tech-

niques, such as distributionally robust optimization [Jiang *et al.*, 2022], knowledge distillation [Bi *et al.*, 2022], graph contrastive learning [Jiang *et al.*, 2023b], and meta-learning [Fu *et al.*, 2023; Manchanda *et al.*, 2022; Son *et al.*, 2023]. Note that the latest meta-learning method [Zhou *et al.*, 2023] considers an omni setting including both cross-distribution and cross-scale generalization.

In this paper, we also focus on the challenging omni-generalization setting to pursue better performance on real-world problems which often have complex and unknown node distributions together with large scales. Unlike previous methods that rely on the divide-and-conquer strategy or try to improve learning algorithms, we propose an ensemble method that integrates a policy of learning from global information of VRPs and a local policy with transferability, where their strengths are combined to boost generalization.

Considering that the general optimization objective of VRPs is to achieve the shortest routing length, the local neighborhood, containing adjacent nodes, plays an important role in the decision of moving from node to node. Meanwhile, local topological features have great potential to be transferable across both node distributions [Jiang *et al.*, 2023b] and scales. To leverage the properties of local information, we propose a novel local policy that restricts the state and action space to a small number of local neighbor nodes. Furthermore, we integrate the local policy with a typical construction policy (e.g., POMO [Kwon *et al.*, 2020]) called global policy, which learns from the global information of complete VRP instances, to form an ensemble policy. The local and global policies are jointly trained to perform cooperatively and complementarily, achieving a good generalization.

We conduct experiments on two typical VRPs: Travelling Salesman Problem (TSP) and Capacitated VRP (CVRP). The proposed ensemble policy is trained on synthetic small-scale problem instances with uniform node distribution. Instead of using test datasets with fixed node distribution and limited problem scale, we select TSPLIB [Reinelt, 1991] and CVRPLIB [Uchoa *et al.*, 2017] benchmarks to evaluate the generalization performance, which contain diverse instances with complex node distributions and large scales. The empirical results show that the ensemble policy significantly improves both cross-distribution and cross-scale generalization performance, and can perform better than state-of-the-art construction methods (e.g., BQ [Drakulic *et al.*, 2023] and LEHD [Luo *et al.*, 2023]) in most cases. Moreover, the proposed method even performs well on real-world problems [Arnold *et al.*, 2019] with several thousand nodes, while most construction methods can hardly solve such real-world problems directly. The ablation study also verifies the key role of the transferrable local policy for better generalization.

2 Background

2.1 Vehicle Routing Problems

VRPs [Dantzig and Ramser, 1959] are well-known NP-hard combinatorial optimization problems with wide applications [Konstantakopoulos *et al.*, 2022]. In this paper, we consider two typical VRPs, i.e., Euclidean TSP and CVRP. Given a CVRP instance I as shown in Figure 1, it contains




Figure 1: Example illustration of CVRP: A CVRP instance with 8 customer nodes and the sequence representation of its solution.

a depot node n_0 and N customer nodes $\{n_1, n_2, \dots, n_N\}$ to serve by multiple identical vehicles, where each node n_i has static Euclidean coordinate (x_{n_i}, y_{n_i}) and demand d_i , and the vehicles have identical capacity Q . The vehicles start from the depot node, visit all the customer nodes, and fulfill their demands. Every customer node must be visited exactly once. If a vehicle runs out of capacity, it has to return to the depot node. The objective of CVRP is to minimize the total traveling distance (Euclidean distance) of the vehicles, i.e.,

$$f(\tau) = \sum_i \sqrt{(x_{\tau_i} - x_{\tau_{i+1}})^2 + (y_{\tau_i} - y_{\tau_{i+1}})^2},$$

without violating the constraints of no repetitive visiting and limited capacity. As shown in Figure 1, an example solution of CVRP can be naturally represented by a sequence $\tau = \{\tau_0, \tau_1, \dots\}$ (where $\tau_i \in \{n_0, n_1, \dots, n_N\}$), which indicates the visiting order of nodes. The Euclidean TSP can be regarded as a simplified version of CVRP that only has one route and no capacity constraint.

2.2 Neural Construction Methods

Neural construction methods learn to construct the sequential solution in an autoregressive way [Vinyals *et al.*, 2015], i.e., learn the probability distribution $P(\tau)$ of generating a solution τ as

$$P(\tau) = \prod_t P(\tau_t | I, \{\tau_0, \tau_1, \dots, \tau_{t-1}\}).$$

This construction procedure can be naturally formulated as a Markov Decision Process (MDP) that takes the partial solution $\{\tau_0, \tau_1, \dots, \tau_{t-1}\}$ as state and the next node τ_t to visit as action. Following such MDP, RL can be used to solve VRPs [Bello *et al.*, 2017; Kool *et al.*, 2019]. In practice, the partial solution information contained in the state is always reduced to the current and start nodes.

Recent prevalent construction methods often employ Transformer [Vaswani *et al.*, 2017] without positional embedding to encode unordered nodes and decode solutions, including Attention Model (AM) [Kool *et al.*, 2019] and Policy Optimization with Multiple Optima (POMO) [Kwon *et al.*, 2020]. Their models have an encoder-decoder architecture. In brief, the encoder learns the node embeddings from the complete VRP graph that covers the global information of a problem instance, and the decoder computes the compatibility of query and keys, i.e.,

$$u_{\text{global}} = \frac{q^T \mathbf{K}}{\sqrt{d}}, \quad q \in \mathbb{R}^{d \times 1}, \quad \mathbf{K} \in \mathbb{R}^{d \times N},$$

as the score for selecting actions, where \mathbf{q} represents the query vector, \mathbf{K} represents the matrix of keys, N is the number of customer nodes, and d is the embedding dimension.

The query vector \mathbf{q} comes from the context embedding that contains the information of the current node and the loaded capacity, representing the state. The keys \mathbf{K} are the embeddings of candidate nodes that represent actions. Thus, the score u_{global} measures the compatibility of state and actions. To avoid violating constraints, u_{global} is masked as u_{masked} with the i -th dimension

$$u_{\text{masked}}^i = \begin{cases} C \cdot \tanh(u_{\text{global}}^i), & \text{if node } n_i \text{ is valid,} \\ -\infty, & \text{otherwise,} \end{cases}$$

where \tanh clips scores to be within $[-1, 1]$, C is a parameter to control the scale, and the scores of invalid actions are set as $-\infty$. The final action probability is computed as

$$\pi_{\text{global}} = \text{softmax}(\mathbf{u}_{\text{masked}}).$$

In each step, the decoder selects one node to grow the partial solution and transits to a new state. The decoder repeats until a complete solution is constructed. When the construction is done, the negative of the objective function is computed as the reward $R(\tau) = -f(\tau)$. A prevalent RL algorithm is REINFORCE [Williams, 1992] with greedy rollout baseline [Kool *et al.*, 2019] or shared baseline [Kwon *et al.*, 2020].

2.3 Generalization Issue

Most existing methods are trained on small-scale problem instances (where $N \leq 100$) generated from uniform node distribution, and tested on instances from the *same* distribution. However, real-world problems are often with complex and unknown node distributions and large scales. This has led to the less satisfactory generalization performance of existing neural methods, which is recognized as a bottleneck for applications [Joshi *et al.*, 2022; Bi *et al.*, 2022; Liu *et al.*, 2023]. Many related works have tried to alleviate the generalization issue by leveraging divide-and-conquer strategy [Hou *et al.*, 2023], meta learning [Manchanda *et al.*, 2022], symmetry of CO problems [Drakulic *et al.*, 2023], and heavy decoder architecture [Luo *et al.*, 2023], while most of them ignore either the cross-distribution or scale setting.

In this paper, we aim to address the generalization issue by ensemble with a transferrable local policy. Unlike a related ensemble method for NCO [Jiang *et al.*, 2023a] that employs the well-known Bootstrap method, we leverage different state and action spaces (i.e., local and global views) to generate diverse and complementary policies. To effectively verify the generalization performance, we select the challenging CVRPLIB [Uchoa *et al.*, 2017; Arnold *et al.*, 2019] and TSPLIB [Reinelt, 1991] as our test datasets since they cover diverse instances and many of them are from the real world, while previous synthetic test datasets with either fixed node distributions or limited problem scales can only cover a narrow subspace of the whole problem space.

3 Method

This section presents the proposed Ensemble of Local and Global policies (ELG), including transferrable local policy,

neural architecture, and ensemble strategy. Overall, the proposed ELG contains two base policies, i.e., a primary global policy and an auxiliary local policy, as shown in Figure 2. The primary global policy can be any popular neural combinatorial optimization model that learns from the global information of a complete VRP instance. Specifically, we choose POMO as the global policy, which has been introduced in Section 2.2. The process of constructing solutions of the local policy is similar to that of the global policy, but the state and action space is different, which focuses on the local neighborhood of the current node.

3.1 Transferrable Local Policy

When solving VRPs, a common observation is that the optimal action (i.e., next node) is usually included in a small local neighborhood of the current node [Helsgaun, 2000], and the patterns of such local neighborhoods have the potential to be transferrable across various node distributions [Jiang *et al.*, 2023b] and scales. Inspired by these facts, we propose to design a local MDP formulation that restricts the state and action space to the local neighborhood, aiming to achieve better generalization performance. The overall procedure of the proposed local policy is illustrated in Figure 2. Here we define the state and action space of local policy in details.

State. The state space is reduced to a small local neighborhood $\mathcal{N}_K(c_t)$, which contains the K nearest *valid* neighbor nodes of the current node c_t . The *valid* nodes are unvisited ones that meet constraints, e.g., capacity constraint in CVRP. Intuitively, the local state represents a subproblem that the policy needs to solve in the next few steps. By learning to solve local subproblems, the policy captures more intrinsic characteristics of VRPs, which are transferable across diverse problem instances. To better represent the features of the local state, we utilize polar coordinates (ρ_i, θ_i) centered at the current node c_t to indicate the location information of the neighbor nodes, where i is the index of node $n_i \in \mathcal{N}_K(c_t)$. The polar coordinates directly provide the relative distances (i.e., edge cost) to c_t , and we can normalize all the ρ_i into $[0, 1]$ by $\tilde{\rho}_i = \rho_i / \max\{\rho_i \mid n_i \in \mathcal{N}_K(c_t)\}$. Therefore, all the neighbor nodes are in a unit ball. Such local topological features are insensitive to the changes in node distribution and problem scale. The state of CVRP also contains the node demands $\{\tilde{d}_i \mid n_i \in \mathcal{N}_K(c_t)\}$, normalized by the remaining capacity Q_{remain} , i.e., $\tilde{d}_i = d_i / Q_{\text{remain}}$. Thus, Q_{remain} is no longer required to be included in the state.

Action. The local policy outputs the scores u_{local} for selecting the next node to visit from the local neighbors, where the score of node n_i can be formulated as

$$u_{\text{local}}^i = \begin{cases} (g_{\theta}(\mathbf{s}_t))_i, & \text{if } n_i \in \mathcal{N}_K(c_t), \\ 0, & \text{otherwise,} \end{cases}$$

where $\mathbf{s}_t = \{[\tilde{\rho}_i, \theta_i, \tilde{d}_i] \mid n_i \in \mathcal{N}_K(c_t)\}$ for CVRP (note that demands \tilde{d}_i are removed for TSP), and g_{θ} is a parameterized neural network, which will be described in Section 3.2.

The transition and reward design remains the same as the global policy [Kwon *et al.*, 2020]. Overall, the local policy learns from the local topological features and outputs the

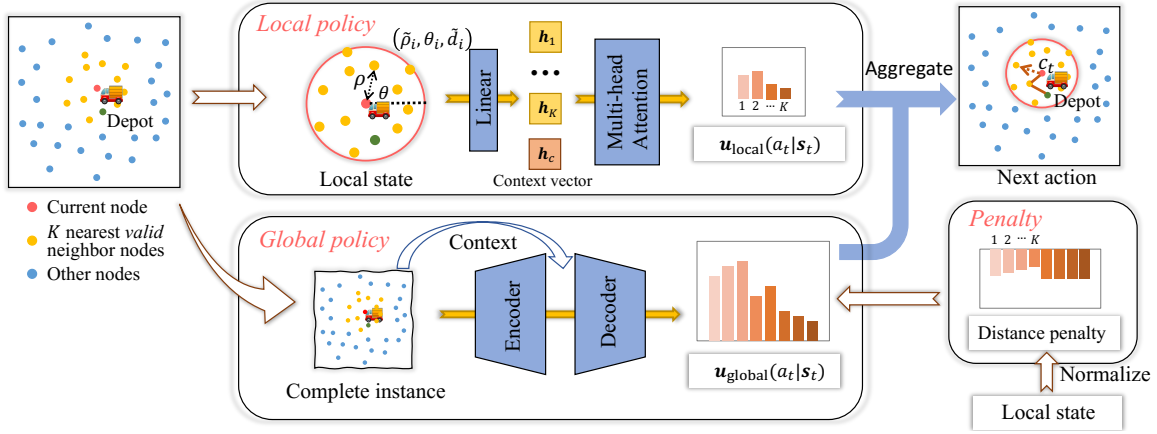


Figure 2: Framework of the ELG method, which aggregates a transferrable local policy that focuses on a small neighborhood of the current node and a global policy that learns from the information of the complete instance.

probability of selecting neighbor nodes. The advantages of the proposed local policy can be summarized as follows: 1) The local pattern is transferrable since it is insensitive to the shifts of node distribution and problem scale; 2) The local MDP gives a good inductive bias that the policy can focus on the local neighbor nodes which are probably more promising.

3.2 Neural Architecture

To parameterize the local policy, we design an attention-based neural network, which has been shown to be effective for solving VRPs [Kool *et al.*, 2019; Kwon *et al.*, 2020]. To ensure the efficiency of local policy, we simply use a lightweight architecture, which consists of one shared linear layer for node embedding and one multi-head attention layer [Vaswani *et al.*, 2017] for decoding action scores.

The node features of local neighbors $\{x_i = [\tilde{\rho}_i, \theta_i, \tilde{d}_i] \mid n_i \in \mathcal{N}_K(c_t)\}$ are firstly mapped into d_e -dimension embeddings using a shared linear layer with parameters $[W; \mathbf{b}]$. Unlike previous attention models [Kool *et al.*, 2019; Kwon *et al.*, 2020] that abandon positional encoding to acquire permutation-invariant node embeddings, we retain the positional encoding to incorporate distance ordering information into the embeddings, which can benefit the local node selection. Concretely, the neighbor nodes x_i are sorted in ascending order by relative distance ρ_i , and the ordering information is injected by the sinusoidal positional encoding [Vaswani *et al.*, 2017]. Thus, the node embedding is computed as

$$\mathbf{h}_i = W\mathbf{x}_i + \mathbf{b} + \text{PE}(\mathbf{x}_i),$$

where PE denotes the sinusoidal positional encoding.

Given node embeddings $\{\mathbf{h}_i \mid n_i \in \mathcal{N}_K(c_t)\}$, a Multi-Head Attention (MHA) layer is executed to compute the node selection scores. The context embedding used for decoding typically relies on the embedding of the current node [Kwon *et al.*, 2020], while the current node in our local encoding is the original point of the polar coordination, and its features are all zeros. To address this issue, we introduce a learnable vector $\mathbf{h}_c \in \mathbb{R}^{d_e}$ to represent the context. We then aggregate the information of node embeddings to the context vector us-

ing an MHA layer as follows:

$$\begin{aligned} \mathbf{k}_i &= W^k \mathbf{h}_i, \mathbf{v}_i = W^v \mathbf{h}_i, \mathbf{q} = W^q \mathbf{h}_c, \\ \tilde{\mathbf{h}}_c &= \text{MHA}([\mathbf{k}_1, \dots, \mathbf{k}_K], [\mathbf{v}_1, \dots, \mathbf{v}_K], \mathbf{q}), \end{aligned}$$

where the keys \mathbf{k}_i and values \mathbf{v}_i are from the node embeddings, and the query \mathbf{q} is from the context vector. The detailed computation of MHA is provided in Appendix B.1. Finally, the compatibility of the aggregated context vector $\tilde{\mathbf{h}}_c$ and the node embeddings is computed as

$$u_{\text{local}}^i = \frac{\tilde{\mathbf{h}}_c^T \mathbf{h}_i}{\sqrt{d_e}}, \quad \forall n_i \in \mathcal{N}_K(c_t),$$

which determines the action scores for selecting nodes.

3.3 Ensemble of Global and Local Policies

We integrate the global policy (which is implemented as POMO [Kwon *et al.*, 2020] in the paper) and the auxiliary local policy to generate an ensemble policy, where the strong in-distribution learning capacity of the global policy and the out-of-distribution transferability of the local policy are combined to boost generalization. Before integration, we adjust the original global policy by adding a distance penalty for better generalization as follows.

Normalized distance penalty for global policy. Considering that most optimal actions are included in the local neighbor nodes, we utilize normalized distances to penalize the global policy. The distance penalty encourages the policy to prefer nearby nodes and be conservative to select remote nodes, which is practically helpful for generalization. Unlike previous methods that directly use the distance values as bias [Son *et al.*, 2023], we propose to normalize distances ρ_i into $[0, 1]$ by the largest ρ_i in $\mathcal{N}_K(c_t)$ and add a fixed penalty ξ ($\xi \geq 1$) for non-neighbor nodes, i.e.,

$$\tilde{u}_{\text{global}}^i = \begin{cases} u_{\text{global}}^i - \frac{\rho_i}{\max\{\rho_i \mid n_i \in \mathcal{N}_K(c_t)\}}, & \text{if } n_i \in \mathcal{N}_K(c_t), \\ u_{\text{global}}^i - \xi, & \text{otherwise,} \end{cases}$$

where $\mathcal{N}_K(c_t)$ contains the K nearest valid neighbor nodes of the current node c_t , and provides a good perspective for normalization, making the distance penalty more scalable.

To integrate the global and local policies, we first add the action scores computed from the two base policies, i.e., $\mathbf{u}_{\text{ens}} = \tilde{\mathbf{u}}_{\text{global}} + \mathbf{u}_{\text{local}}$. After that, the action probability of the ensemble policy π_{ens} is computed by

$$u_{\text{masked}}^i = \begin{cases} C \cdot \tanh(\tilde{u}_{\text{global}}^i + u_{\text{local}}^i), & \text{if node } n_i \text{ is valid,} \\ -\infty, & \text{otherwise,} \end{cases}$$

$$\pi_{\text{ens}} = \text{softmax}(\mathbf{u}_{\text{masked}}).$$

Instead of training each base policy independently, we use a joint training method to directly optimize π_{ens} , which encourages the global policy and the local policy to work cooperatively, so that their strengths can be well combined.

Training. In practice, we first pretrain the global policy with distance penalty for T_1 epochs before joint training, because the state and action space of the global policy is more complex than the local one. At the joint training stage, we use the policy gradient method to directly train the ensemble policy π_{ens} for T_2 epochs, which includes the trainable parameters $\tilde{\theta}$ from the global policy and θ from the local policy. Following POMO, we use multiple rollouts from different start nodes to get multiple trajectories in a single feed-forward, and utilize the REINFORCE [Williams, 1992] algorithm with shared baseline [Kwon *et al.*, 2020] to estimate the gradient of the expected return J . Specifically, the average reward of multiple rollouts is used as the REINFORCE baseline, and the gradient $\nabla_{\tilde{\theta}, \theta} J(\tilde{\theta}, \theta)$ is estimated by

$$\frac{1}{N \cdot B} \sum_{i=1}^B \sum_{j=1}^N \left(R_{i,j} - \frac{1}{N} \sum_{j=1}^N R_{i,j} \right) \nabla_{\tilde{\theta}, \theta} \log \pi_{\text{ens}}(\tau_{i,j}),$$

where N is the number of trajectories that is equal to the number of nodes, B is the batch size, and $R_{i,j}$ is the reward of the j -th trajectory $\tau_{i,j}$ on the i -th instance.

4 Experiments

We conduct experiments on two typical VRPs, i.e., Traveling Salesman Problem (TSP) and Capacitated VRP (CVRP). This section presents the experimental settings and results to evaluate the performance of our method and verify the efficacy of important components. Our code and data are available in <https://github.com/lamda-bbo/ELG>.

4.1 Experimental Settings

We first introduce the experimental settings, including baselines, datasets, and performance metrics.

Baselines. For non-learning heuristic methods, two highly optimized solvers LKH3 [Helsgaun, 2017] and HGS [Vidal, 2022] are compared and their results on CVRPLIB are taken from the paper of HGS [Vidal, 2022]. For learning-based methods, we collect them into several groups according to their main ideas, which are introduced as follows.

- **POMO-based methods:** POMO [Kwon *et al.*, 2020] is a widely-used backbone model for solving VRPs. We include POMO and its variant Sym-POMO [Kim *et al.*, 2022] in our experiments. As the training of both POMO

and Sym-POMO suffers from severe over-fitting, which results in degraded generalization performance, we early stop their training based on a validation dataset with larger scale ($N=100$ for training and $N=200$ for validation).

- **Cross-distribution methods:** For those methods that improve cross-distribution generalization, we compare the state-of-the-art methods AMDKD [Bi *et al.*, 2022] and Omni-POMO [Zhou *et al.*, 2023], where Omni-POMO also considers cross-scale performance.
- **Cross-scale methods:** We compare three end-to-end learning methods that consider cross-scale generalization, Pointerformer [Jin *et al.*, 2023], LEHD [Luo *et al.*, 2023] and BQ [Drakulic *et al.*, 2023]. For BQ, we select the best BQ-transformer model for comparison. We also compare a novel divide-and-conquer method TAM [Hou *et al.*, 2023] on CVRPLIB Set-X ($N \geq 500$) and Set-XXL. Since the code of TAM has not been released, and their reported results are not aligned with ours, we provide the comparison with TAM in Appendix C.3 due to space limitation.
- **Diffusion model-based methods:** Recent works utilize the diffusion model to predict the solution distributions as heatmaps, showing great potential in generalization, especially on TSBLIB. We select DIFUSCO [Sun and Yang, 2023] and T2TCO [Li *et al.*, 2023] for comparison.

All the learning methods use greedy mode at the inference stage. POMO-based methods include $\times 8$ instance augmentation technique [Kwon *et al.*, 2020]. Our proposed method ELG is implemented by setting the global policy as POMO, denoted as ELG-POMO. Its detailed hyperparameter settings are provided in Appendix A due to space limitation. Note that the training of ELG-POMO is aligned with POMO; thus, the comparison between them could be seen as an ablation study of the proposed local policy.

Datasets. The training instances are randomly generated with uniform node distribution and the problem scale $N = 100$, which is a common setting in previous works [Kool *et al.*, 2019; Kwon *et al.*, 2020]. More details of training settings can be found in Appendix B.2. We consider the following test datasets to compare different methods.

- **TSPLIB and CVRPLIB Set X & XXL:** These benchmarks cover diverse problem instances with complex and unknown distributions together with varying and large scales, which can be used to evaluate the cross-distribution and cross-scale generalization in more realistic scenarios.
- **Cross-distribution small-scale instances:** Some cross-distribution methods (e.g., AMDKD [Bi *et al.*, 2022]) only focus on small-scale instances. Testing them on large-scale instances like CVRPLIB may overshadow their advantages. Thus, we conduct experiments under their settings, where small-scale instances with complex synthetic distributions including cluster, cluster mixture, explosion, and rotation [Bossek *et al.*, 2019] are used for evaluation.
- **Uniform small-scale instances:** The improvement on out-of-distribution (scale) generalization should not sacrifice the in-distribution (scale) performance. Therefore, we include the uniform small-scale dataset to examine the in-distribution (scale) performance.

Performance metrics. We report the average gap to Best Known Solutions (BKS), i.e., $\frac{1}{n} \sum_{i=1}^n \frac{f_i^* - f_i(\text{BKS})}{f_i(\text{BKS})}$, on each dataset to evaluate the generalization performance, where n is the number of instances, f_i^* and $f_i(\text{BKS})$ denote the objective value of the solution generated by each method and the BKS, respectively. We also report the average time of solving one instance on TSPLIB and CVRPLIB to measure the efficiency.

4.2 Results on TSPLIB and CVRPLIB

For fair comparisons, all the learning methods are trained on uniform instances with the problem scale $N = 100$, except that Omni-POMO has a stronger training setting that uses instances with complex distributions and varying scales [Zhou *et al.*, 2023]. In the following, we analyze the generalization performance of different methods in detail. The CVRPLIB Set-X and TSPLIB are divided into two subsets according to the problem scales, which are $N \leq 200$ and $200 < N \leq 1000$ (or $200 < N \leq 1002$). First, we examine the cross-distribution performance on the $N \in (0, 200]$ subset, and use the $N \in (200, 1000]$ (or $N \in (200, 1002]$) subset to compare the cross-distribution and cross-scale performance simultaneously. Then, the CVRPLIB Set-XXL is used to evaluate the performance on real-world instances with $N \in [3000, 7000]$. Finally, we compare our method with highly optimized solvers LKH3 and HGS.

Cross-distribution performance. Because the test distribution is complex and the problem scales of $N \leq 200$ are close to that of training instances, we can evaluate the cross-distribution performance by the $N \leq 200$ subsets of TSPLIB and CVRPLIB Set-X. As shown in the second columns of Tables 1 and 2, ELG-POMO outperforms all the methods on the $N \leq 200$ subset, and the gap reductions on CVRP and TSP are 0.53% and 0.50%, respectively, compared to the runner-up (Omni-POMO for CVRP, and BQ for TSP). These results clearly demonstrate the superiority of ELG-POMO in cross-distribution generalization, which will be further verified by the results in Table 4.

Cross-distribution and cross-scale performance. The subset of $N \in (200, 1000]$ (or $N \in (200, 1002]$) contains instances with both cross-distribution and cross-scale characteristics. ELG-POMO reduces the gap from 6.95% of Omni-POMO to 6.46% for CVRP, and reduces the gap from 7.47% of Omni-POMO to 5.90% for TSP. Considering that the Omni-POMO is trained on cross-distribution (uniform and Gaussian mixture) and cross-scale ($N \in [50, 200]$) instances while ELG-POMO only uses uniform instances with the problem scale $N = 100$ in the training, the above comparison clearly demonstrates the better cross-distribution and cross-scale performance of the proposed method. Note that ELG-POMO outperforms BQ and LEHD on CVRP, but lags behind them on TSP with $N \in (200, 1002]$. This may be attributed to their heavy decoder or decoder-only architectures, which can learn strong scale-invariant features [Luo *et al.*, 2023] to achieve good large-scale performance. We can also observe from Table 2 that ELG-POMO can surpass the diffusion mode-based methods, i.e., DIFUSCO and T2TCO, under the greedy inference setting.

Method	(0, 200]	(200, 1000]	Total	Time
LKH3	0.36%	1.18%	1.00%	16m
HGS	0.01%	0.13%	0.11%	16m
HGS (less time)	0.21%	1.28%	1.59%	10s
POMO	5.26%	11.82%	10.37%	0.80s
Sym-POMO	9.99%	27.09%	23.32%	0.87s
Omni-POMO	5.04%	6.95%	6.52%	0.75s
LEHD	11.11%	12.73%	12.25%	1.67s
BQ	10.60%	10.97%	10.89%	3.36s
ELG-POMO (Ours)	4.51%	6.46%	6.03%	1.90s

Table 1: Empirical results on CVRPLIB Set-X covering cross-distribution and cross-scale test instances. Note that Pointerformer, DIFUSCO, and T2TCO are not intended for CVRP.

Method	(0, 200]	(200, 1002]	Total	Time
LKH3	0.00%	0.00%	0.00%	24s
POMO	3.07%	13.35%	7.45%	0.41s
Sym-POMO	2.88%	15.35%	8.29%	0.34s
Omni-POMO	1.74%	7.47%	4.16%	0.34s
Pointerformer	2.31%	11.47%	6.32%	0.24s
LEHD	2.03%	3.12%	2.50%	1.28s
BQ	1.62%	2.39%	2.22%	2.85s
DIFUSCO	1.84%	10.83%	5.77%	30.68s
T2TCO	1.87%	9.72%	5.30%	30.82s
ELG-POMO (Ours)	1.12%	5.90%	3.08%	0.63s

Table 2: Empirical results on TSPLIB covering cross-distribution and cross-scale test instances.

Performance on real-world CVRP instances. We evaluate our method on four real-world instances, i.e., Antwerp1 (A1), Antwerp2 (A2), Leuven1 (L1), and Leuven2 (L2), with $N \in [3000, 7000]$ from CVRPLIB Set-XXL. As shown in Table 3, the ELG-POMO method performs the best on three instances and significantly reduces the optimality gap compared to other POMO-based methods.

Method	A1	A2	L1	L2
POMO	112.27%	159.22%	75.30%	78.16%
Sym-POMO	79.72%	179.55%	170.44%	179.55%
Omni-POMO	42.52%	48.59%	22.79%	60.39%
LEHD	18.90%	26.40%	14.04%	26.30%
BQ	11.21%	15.02%	13.27%	24.00%
ELG-POMO (Ours)	10.70%	17.69%	10.77%	21.80%

Table 3: Empirical results on four real-world instances from CVRPLIB Set-XXL with $N \in [3000, 7000]$.

Comparison to non-learning heuristic solvers. We can observe from Table 1 that the time consumption of ELG-POMO is smaller than LKH3 and HGS (1.90s vs. 16m). However, its optimality gap is worse, even when HGS uses less time, implying that further improvement needs to be done for neural construction methods in the future.

Runtime. As shown in Tables 1 and 2, ELG-POMO costs a lower runtime than BQ and diffusion model-based methods. Compared to the pure POMO method, introducing an additional local policy increases the time cost, but is acceptable.

4.3 Results on Synthetic Distributions

We also conduct experiments on small-scale instances generated from synthetic distributions. First, we will compare ELG-POMO with AMDKD on small-scale instances from four complex distributions to further verify the superiority of ELG-POMO in cross-distribution generalization. Second, we will show that the proposed ELG-POMO does not worsen the in-distribution (scale) performance when improving the out-of-distribution (scale) performance. In this subsection, we report the gap of average objective values following the common setting [Kool *et al.*, 2019; Bi *et al.*, 2022].

Cross-distribution performance. AMDKD is the state-of-the-art method for cross-distribution generalization [Bi *et al.*, 2022]. As it is not designed for cross-scale generalization, we present the comparison with AMDKD here. For fairness, we retrain the POMO and ELG-POMO (denoted as POMO* and ELG*) using three exemplar distributions: uniform, cluster, and cluster mixture, which are aligned with AMDKD. Then we select four difficult distributions to evaluate them. The results on TSP are shown in Table 4, where ELG* outperforms AMDKD on three distributions with $N = 100$ and four distributions with $N = 200$. The gap reductions on four distributions with $N = 200$ are roughly 1.0%–1.9%, which are very significant. Similar experimental results on CVRP are presented in Appendix C.4 due to space limitation.

Distributions	Scale	POMO*	AMDKD	ELG*
Cluster	$N = 100$	0.45%	0.35%	0.36%
Mixture	$N = 100$	0.49%	0.41%	0.37%
Explosion	$N = 100$	0.29%	0.30%	0.23%
Rotation	$N = 100$	0.38%	0.32%	0.31%
Cluster	$N = 200$	2.63%	3.01%	1.84%
Mixture	$N = 200$	2.45%	3.50%	1.62%
Explosion	$N = 200$	2.18%	2.63%	1.43%
Rotation	$N = 200$	2.33%	2.54%	1.58%

Table 4: Empirical results on TSP instances with four difficult synthetic distributions.

In-distribution test performance. We train and test all the methods on uniform instances with $N = 100$. The results are shown in Table 5, where we use the solutions generated by Concorde [Applegate *et al.*, 2006] and LKH3 as the BKS for TSP and CVRP, respectively, to compute gaps. For TSP100, the proposed ELG-POMO achieves competitive performance to state-of-the-art methods for TSP (Pointerformer and T2TCO) with 0.04%–0.06% difference in the gap. For CVRP100, ELG-POMO surpasses all the compared methods. Note that although BQ and LEHD improve the cross-distribution and cross-scale generalization as shown in Tables 1 to 3, their in-distribution test results get worse than POMO, except for BQ’s result on TSP100. On the contrary, besides the clear advantage of cross-distribution and cross-scale generalization, the proposed ELG-POMO can improve the in-distribution performance to some extent, i.e., the introduction of transferrable local policy also benefits the generalization to unseen in-distribution (scale) instances.

Method	TSP100	CVRP100
Concorde	0.00%	–
LKH3	0.00%	0.00%
HGS	–	-0.52%
POMO	0.40%	1.34%
Sym-POMO	0.41%	1.43%
Pointerformer	0.16%	–
LEHD	0.57%	3.64%
BQ	0.35%	2.76%
DIFUSCO	0.24%	–
T2TCO	0.18%	–
ELG-POMO (Ours)	0.22%	1.22%

Table 5: Empirical results on uniform instances with $N = 100$, where “–” means that the method is not intended for that problem.

4.4 Ablation Studies

We conduct ablation studies on three components associated with the local policy: normalized distance penalty, positional encoding, and polar coordinates. The results show that these components really contribute to improving performance, and the proposed ELG method still achieves good generalization performance even without these components. Detailed results are provided in Appendix C.1 due to space limitation.

We also examine the influence of two new hyperparameters introduced by the proposed method: the fixed penalty ξ and the local neighborhood size K . The results show that when ξ and K are within reasonable ranges, they do not affect the performance significantly. Detailed results are provided in Appendix C.2 due to space limitation.

5 Conclusion

The generalization issue of NCO methods has prevented them from real-world applications. In this paper, we propose to ensemble a global policy and a transferrable local policy for VRPs, where the global policy learns from the global information of a complete VRP graph, resulting in its strong in-distribution learning capacity, and the local policy learns from the local topological features which are insensitive to the variation of node distribution and problem scale, leading to its good transferability. The combination of these two complementary policies leads to better generalization, which is verified by the extensive comparison with state-of-the-art methods on two well-known benchmarks (i.e., TSPLIB and CVRPLIB) covering diverse problem instances. A limitation of the proposed method is that it costs more time to solve very large-scale instances than divide-and-conquer methods as shown in Appendix C.3. In future work, we will try to further improve the latency and extend the proposed method to more CO problems, such as the pickup-and-delivery problem.

Acknowledgments

The authors want to thank the anonymous reviewers for their helpful comments and suggestions. This work was supported by the National Science and Technology Major Project (2022ZD0116600) and National Science Foundation

of China (62276124). Chao Qian is the corresponding author. Due to space limitation, some experimental settings and results are provided in Appendix. The version with Appendix is available at <https://arxiv.org/abs/2308.14104>.

References

- [Applegate *et al.*, 2006] David Applegate, Robert Bixby, Vasek Chvatal, and William Cook. Concorde TSP solver. <http://www.math.uwaterloo.ca/tsp/concorde/m>, 2006.
- [Arnold *et al.*, 2019] Florian Arnold, Michel Gendreau, and Kenneth Sörensen. Efficiently solving very large-scale routing problems. *Computers & Operations Research*, 107:32–42, 2019.
- [Bello *et al.*, 2017] Irwan Bello, Hieu Pham, Quoc V Le, Mohammad Norouzi, and Samy Bengio. Neural combinatorial optimization with reinforcement learning. In *Proceedings of the 5th ICLR*, Toulon, France, 2017.
- [Bengio *et al.*, 2021] Yoshua Bengio, Andrea Lodi, and Antoine Prouvost. Machine learning for combinatorial optimization: a methodological tour d’horizon. *European Journal of Operational Research*, 290(2):405–421, 2021.
- [Bi *et al.*, 2022] Jieyi Bi, Yining Ma, Jiahai Wang, Zhiguang Cao, Jinbiao Chen, Yuan Sun, and Yeow Meng Chee. Learning generalizable models for vehicle routing problems via knowledge distillation. In *Advances in Neural Information Processing Systems 35 (NeurIPS)*, pages 31226–31238, New Orleans, LA, 2022.
- [Bossek *et al.*, 2019] Jakob Bossek, Pascal Kerschke, Aneta Neumann, Markus Wagner, Frank Neumann, and Heike Trautmann. Evolving diverse TSP instances by means of novel and creative mutation operators. In *Proceedings of the 15th ACM/SIGEVO Conference on Foundations of Genetic Algorithms (FOGA)*, pages 58–71, Potsdam, Germany, 2019.
- [Cappart *et al.*, 2023] Quentin Cappart, Didier Chételat, Elias B Khalil, Andrea Lodi, Christopher Morris, and Petar Veličković. Combinatorial optimization and reasoning with graph neural networks. *Journal of Machine Learning Research*, 24(130):1–61, 2023.
- [Chen *et al.*, 2023] Wenjie Chen, Shengcai Liu, Yew-Soon Ong, and Ke Tang. Neural influence estimator: Towards real-time solutions to influence blocking maximization. *arXiv:2308.14012*, 2023.
- [Cheng *et al.*, 2023] Hanni Cheng, Haosi Zheng, Ya Cong, Weihao Jiang, and Shiliang Pu. Select and optimize: Learning to solve large-scale TSP instances. In *Proceedings of The 26th International Conference on Artificial Intelligence and Statistics (AISTATS)*, pages 1219–1231, Valencia, Spain, 2023.
- [Christiaens and Vanden Berghe, 2020] Jan Christiaens and Greet Vanden Berghe. Slack induction by string removals for vehicle routing problems. *Transportation Science*, 54(2):417–433, 2020.
- [Dantzig and Ramser, 1959] George B Dantzig and John H Ramser. The truck dispatching problem. *Management Science*, 6(1):80–91, 1959.
- [Drakulic *et al.*, 2023] Darko Drakulic, Sofia Michel, Florian Mai, Arnaud Sors, and Jean-Marc Andreoli. BQ-NCO: Bisimulation quotienting for generalizable neural combinatorial optimization. In *Advances in Neural Information Processing Systems 36 (NeurIPS)*, New Orleans, LA, 2023.
- [Fu *et al.*, 2021] Zhang-Hua Fu, Kai-Bin Qiu, and Hongyuan Zha. Generalize a small pre-trained model to arbitrarily large TSP instances. In *Proceedings of the 35th AAAI*, pages 7474–7482, Virtual, 2021.
- [Fu *et al.*, 2023] Qiming Fu, Zhechao Wang, Nengwei Fang, Bin Xing, Xiao Zhang, and Jianping Chen. Maml2: meta reinforcement learning via meta-learning for task categories. *Frontiers of Computer Science*, 17(4):174325, 2023.
- [Helsgaun, 2000] Keld Helsgaun. An effective implementation of the lin–kernighan traveling salesman heuristic. *European Journal of Operational Research*, 126(1):106–130, 2000.
- [Helsgaun, 2017] Keld Helsgaun. An extension of the lin–kernighan–helsgaun TSP solver for constrained traveling salesman and vehicle routing problems. Technical report, 2017.
- [Hou *et al.*, 2023] Qingchun Hou, Jingwei Yang, Yiqiang Su, Xiaoqing Wang, and Yuming Deng. Generalize learned heuristics to solve large-scale vehicle routing problems in real-time. In *Proceedings of the 11th ICLR*, Kigali, Rwanda, 2023.
- [Jiang *et al.*, 2022] Yuan Jiang, Yaoxin Wu, Zhiguang Cao, and Jie Zhang. Learning to solve routing problems via distributionally robust optimization. In *Proceedings of the 36th AAAI*, pages 9786–9794, Virtual, 2022.
- [Jiang *et al.*, 2023a] Yuan Jiang, Zhiguang Cao, Yaoxin Wu, Wen Song, and Jie Zhang. Ensemble-based deep reinforcement learning for vehicle routing problems under distribution shift. In *Advances in Neural Information Processing Systems 36 (NeurIPS)*, pages 53112–53125, New Orleans, LA, 2023.
- [Jiang *et al.*, 2023b] Yuan Jiang, Zhiguang Cao, Yaoxin Wu, and Jie Zhang. Multi-view graph contrastive learning for solving vehicle routing problems. In *Proceedings of the 39th Conference on Uncertainty in Artificial Intelligence (UAI)*, pages 984–994, pittsburgh, PA, 2023.
- [Jin *et al.*, 2023] Yan Jin, Yuandong Ding, Xuanhao Pan, Kun He, Li Zhao, Tao Qin, Lei Song, and Jiang Bian. Pointerformer: Deep reinforced multi-pointer transformer for the traveling salesman problem. In *Proceedings of the 37th AAAI*, pages 8132–8140, Washington, DC, 2023.
- [Joshi *et al.*, 2019] Chaitanya K. Joshi, Thomas Laurent, and Xavier Bresson. An efficient graph convolutional network technique for the travelling salesman problem. *arXiv:1906.01227*, 2019.
- [Joshi *et al.*, 2022] Chaitanya K. Joshi, Quentin Cappart, Louis-Martin Rousseau, and Thomas Laurent. Learning the travelling salesperson problem requires rethinking generalization. *Constraints*, 27(1-2):70–98, 2022.

- [Khalil *et al.*, 2017] Elias B. Khalil, Hanjun Dai, Yuyu Zhang, Bistra Dilikina, and Le Song. Learning combinatorial optimization algorithms over graphs. In *Advances in Neural Information Processing Systems 30 (NeurIPS)*, pages 6348–6358, Long Beach, CA, 2017.
- [Kim *et al.*, 2022] Minsu Kim, Junyoung Park, and Jinkyoo Park. Sym-NCO: Leveraging symmetricity for neural combinatorial optimization. In *Advances in Neural Information Processing Systems 35 (NeurIPS)*, pages 1936–1949, New Orleans, LA, 2022.
- [Konstantakopoulos *et al.*, 2022] Grigorios D Konstantakopoulos, Sotiris P Gayialis, and Evripidis P Kechagias. Vehicle routing problem and related algorithms for logistics distribution: A literature review and classification. *Operational Research*, 22(3):2033–2062, 2022.
- [Kool *et al.*, 2019] Wouter Kool, Herke van Hoof, and Max Welling. Attention, learn to solve routing problems! In *Proceedings of the 7th ICLR*, New Orleans, LA, 2019.
- [Kwon *et al.*, 2020] Yeong-Dae Kwon, Jinho Choo, Byoungjip Kim, Iljoo Yoon, Youngjune Gwon, and Seungjai Min. POMO: Policy optimization with multiple optima for reinforcement learning. In *Advances in Neural Information Processing Systems 33 (NeurIPS)*, pages 21188–21198, Virtual, 2020.
- [Li *et al.*, 2023] Yang Li, Jinpei Guo, Runzhong Wang, and Junchi Yan. From distribution learning in training to gradient search in testing for combinatorial optimization. In *Advances in Neural Information Processing Systems 36 (NeurIPS)*, New Orleans, LA, 2023.
- [Liu *et al.*, 2023] Shengcai Liu, Yu Zhang, Ke Tang, and Xin Yao. How good is neural combinatorial optimization? A systematic evaluation on the traveling salesman problem. *IEEE Computational Intelligence Magazine*, 18(3):14–28, 2023.
- [Lu *et al.*, 2019] Hao Lu, Xingwen Zhang, and Shuang Yang. A learning-based iterative method for solving vehicle routing problems. In *Proceedings of the 7th ICLR*, New Orleans, LA, 2019.
- [Luo *et al.*, 2023] Fu Luo, Xi Lin, Fei Liu, Qingfu Zhang, and Zhenkun Wang. Neural combinatorial optimization with heavy decoder: Toward large scale generalization. In *Advances in Neural Information Processing Systems 36 (NeurIPS)*, New Orleans, LA, 2023.
- [Ma *et al.*, 2023] Yining Ma, Zhiguang Cao, and Yeow Meng Chee. Learning to search feasible and infeasible regions of routing problems with flexible neural k-opt. In *Advances in Neural Information Processing Systems 36 (NeurIPS)*, New Orleans, LA, 2023.
- [Manchanda *et al.*, 2022] Sahil Manchanda, Sofia Michel, Darko Drakulic, and Jean-Marc Andreoli. On the generalization of neural combinatorial optimization heuristics. In *Proceedings of the Joint European Conference on Machine Learning and Knowledge Discovery in Databases (ECML PKDD)*, pages 426–442, Grenoble, France, 2022.
- [Reinelt, 1991] Gerhard Reinelt. TSPLIB - A traveling salesman problem library. *ORSA Journal on Computing*, 3(4):376–384, 1991.
- [Son *et al.*, 2023] Jiwoo Son, Minsu Kim, Hyeonah Kim, and Jinkyoo Park. Meta-sage: Scale meta-learning scheduled adaptation with guided exploration for mitigating scale shift on combinatorial optimization. In *Proceedings of the 40th ICML*, Honolulu, Hawaii, 2023.
- [Song *et al.*, 2023] Wen Song, Nan Mi, Qiqiang Li, Jing Zhuang, and Zhiguang Cao. Stochastic economic lot scheduling via self-attention based deep reinforcement learning. *IEEE Transactions on Automation Science and Engineering*, pages 1–12, 2023.
- [Sun and Yang, 2023] Zhiqing Sun and Yiming Yang. DIFUSCO: Graph-based diffusion solvers for combinatorial optimization. In *Advances in Neural Information Processing Systems 36 (NeurIPS)*, New Orleans, LA, 2023.
- [Uchoa *et al.*, 2017] Eduardo Uchoa, Diego Pecin, Artur Pessoa, Marcus Poggi, Thibaut Vidal, and Anand Subramanian. New benchmark instances for the capacitated vehicle routing problem. *European Journal of Operational Research*, 257(3):845–858, 2017.
- [Vaswani *et al.*, 2017] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in Neural Information Processing Systems 30 (NeurIPS)*, pages 5998–6008, Long Beach, CA, 2017.
- [Vidal, 2022] Thibaut Vidal. Hybrid genetic search for the CVRP: Open-source implementation and swap* neighborhood. *Computers & Operations Research*, 140:105643, 2022.
- [Vinyals *et al.*, 2015] Oriol Vinyals, Meire Fortunato, and Navdeep Jaitly. Pointer networks. In *Advances in Neural Information Processing Systems 28 (NeurIPS)*, pages 2692–2700, Montreal, Canada, 2015.
- [Williams, 1992] Ronald J Williams. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine Learning*, 8(3):229–256, 1992.
- [Zhang *et al.*, 2020] Cong Zhang, Wen Song, Zhiguang Cao, Jie Zhang, Puay Siew Tan, and Xu Chi. Learning to dispatch for job shop scheduling via deep reinforcement learning. In *Advances in Neural Information Processing Systems 33 (NeurIPS)*, pages 1621–1632, Vancouver, Canada, 2020.
- [Zhou *et al.*, 2023] Jianan Zhou, Yaoxin Wu, Wen Song, Zhiguang Cao, and Jie Zhang. Towards omnigeneralizable neural methods for vehicle routing problems. In *Proceedings of the 40th ICML*, pages 42769–42789, Honolulu, HI, 2023.