

Scale-Net: A Hierarchical U-Net Framework for Cross-Scale Generalization in Multi-Task Vehicle Routing

Suyu Liu¹, Zhiguang Cao², Nan Yin^{3*}, Yew-Soon Ong^{1,4}

¹College of Computing and Data Science, Nanyang Technological University

²School of Computing and Information Systems, Singapore Management University

³Department of Computer Science and Engineering, Hong Kong University of Science and Technology

⁴Centre for Frontier AI Research, Institute of High Performance Computing, Agency for Science, Technology and Research
liusuyu97@gmail.com, zgcao@smu.edu.sg, yinnan8911@gmail.com, asysong@ntu.edu.sg

Abstract

Neural solvers for Vehicle Routing Problems (VRPs) have shown great advantages in solving various kinds of problem types. However, they also face critical challenges in generalizing from small-scale training to large-scale problems and in identifying the most salient topological information for decision-making. To mitigate these gaps, we introduce Scale-Net, a novel hierarchical framework that integrates a U-Net architecture into a unified, multi-task VRP solver. Scale-Net explicitly captures multi-scale structural patterns by processing a nested hierarchy of input graph instances. This enriched, coarse-to-fine representation is extracted by the encoder and fed directly into the decoder, empowering decoder module with superior topological awareness for routing decisions while simultaneously reducing computational overhead in the encoder. We conducted extensive experiments on 16 VRP variants with instances ranging from 50 to 5,000 nodes. The experimental results show that Scale-Net demonstrates significant performance gains over state-of-the-art baselines across in-distribution, zero-shot, and real-world settings.

Code — <https://github.com/lisyys19711/Scale-Net>

Introduction

The Vehicle Routing Problems (VRPs) and its variants are a group of fundamental challenges in the area of combinatorial optimization with a broad range of applications in related downstream settings such as transportation (Ge, Li, and Tuzhilin 2019; Zhou et al. 2023a), supply chain management (Zhang et al. 2023), and city logistics (Cattaruzza et al. 2017). These problems target at determining the optimal set of routes for a group of vehicles that provide delivery service for customers, while adhering to different kinds of constraints like capacity, time limits, backhaul etc. However, the inherent combinatorial complexity of VRPs makes finding optimal solutions computationally intractable. To mitigate this issue, several sophisticated heuristic algorithms have been proposed with methods like Lin-Kernighan-Helsgaun (LKH) (Lin and Kernighan 1973) and Hybrid Genetic Search (HGS) (Vidal 2022) standing as prominent examples. Although these solvers have achieved

remarkable success in finding near-optimal solutions, they are developed on handcrafted rules and problem-specific domain knowledge, which severely limits their ability to generalize to new or unseen constraints that often encountered in real-world scenarios.

In response to these limitations, deep learning has given rise to neural solvers in which cases solvers learn to construct solutions in an automatic manner. Typically, these models utilize an encoder-decoder architecture, where an encoder maps the problem instance into a high-dimensional feature space and an auto-regressive decoder sequentially builds the routes during which an attention mechanism (Vinyals, Fortunato, and Jaitly 2015; Vaswani et al. 2017) guides the decision-making process. Furthermore, there have emerged two frontiers in recent research that have earned great attentions: 1) creating unified foundation models capable of solving multiple VRP variants with a single architecture (Liu et al. 2024; Zhou et al. 2024; Berto et al. 2024; Huang et al. 2025; Li et al. 2025a) and 2) developing specialized large-scale solvers that can handle instances with thousands of nodes for a single task (Luo et al. 2023; Pan et al. 2025; Li et al. 2025b; Luo et al. 2025).

Despite these remarkable progresses, a critical gap remains at the intersection of these frontiers, thereby creating two fundamental challenges. Firstly, there is a *scaling-unification dilemma*: current existing state-of-the-art unified models (Kwon et al. 2020; Zhou et al. 2023a; Berto et al. 2024; Zhou et al. 2024) can handle diverse tasks but are trained on small graphs (typically fewer than 100 nodes) and fail to scale the training into thousands of nodes due to prohibitive computational costs. Conversely, existing large-scale solvers (Luo et al. 2023; Pan et al. 2025; Li et al. 2025b; Luo et al. 2025) are effective on thousand-node instances but are task-specific, lacking the flexibility to generalize across different constraints and tasks. Secondly, most architectures suffer from a *decoder information bottleneck*: decoders typically operate with a narrow view, using only the current node’s embedding and dynamic state information to make next step’s decision. They lack of the access into the rich, multi-scale structural context of the graph, which is crucial for making topological-aware routing decisions. Consequently, a single neural solver that is simultaneously unified, scalable with high-fidelity performance remains as an open and under-explored area.

*Corresponding author.

In this paper, we introduce Scale-Net, a novel framework which is designed to directly address these challenges. In specifics, we propose a unified, multi-task neural solver built on a single backbone that achieves both cross-task versatility and large-scale applicability. The core of our innovation is a hierarchical encoder that integrates U-Net modules (Ronneberger, Fischer, and Brox 2015; Gao and Ji 2019) to extract multi-scale structural patterns. By iteratively processing a nested hierarchy of sampled sub-graphs, Scale-Net constructs a coarse-to-fine topological representation. This enriched, multi-level information is then explicitly fed into the decoder, empowering it to make more precise and contextually-aware decisions. Furthermore, by calculating attention scores only on sampled nodes at each layer, our approach significantly reduces computational overhead, making it feasible to train and perform inference on large-scale VRP instances. We summarize our contributions as follows:

- **A Unified and Scalable VRP Solver:** We introduce Scale-Net, the first neural architecture that is designed to function as a single, unified solver for multiple VRP variants while also scaling effectively to large-scale instances involving thousands of nodes. This resolves the critical trade-off between cross-task generalization and large-scale applicability which previous models lack.
- **Hierarchical U-Net for Enriched Decoding:** We propose a novel U-Net-based encoder that constructs a multi-scale feature representation of the input graph. This technique resolves the information bottleneck problem in standard decoders by providing a rich, coarse-to-fine topological-aware context, leading to more precise routing decisions and lower computational overhead.
- **State-of-the-Art Performance Across Scales and Tasks:** We conduct extensive experiments on 16 VRP variants with problem sizes ranging from 50 to 5,000 nodes. The results validate that Scale-Net establishes a new state of the art, outperforming existing baselines in in-distribution, zero-shot, few-shot, large-scale generalization and real-world settings.

Related Works

VRP Solvers

Solvers for Vehicle Routing Problems (VRPs) are typically grouped into three main classes. 1) Traditional Solvers, such as the Lin-Kernighan-Helsgaun (LKH) algorithm (Lin and Kernighan 1973), Hybrid Genetic Search (HGS) (Vidal 2022), and OR-Tools (Perron and Didier 2024), utilize established heuristic search techniques. Their reliance on expert-crafted rules and delicate domain knowledge, however, can restrict their adaptability to unseen problem variants. 2) Neural Solvers, originating from models like Pointer Networks (Vinyals, Fortunato, and Jaitly 2015), use deep learning to build solutions iteratively. The integration of self-attention (Vaswani et al. 2017; Kool, van Hoof, and Welling 2019; Kwon et al. 2020) dramatically boosted performances on both of the in-distribution and out-of distribution settings. Subsequent research has improved generalization through varied training (Kim, Park, and Park 2022; Zhou et al.

2023b), scaled to larger instances (Luo et al. 2023; Pan et al. 2023a; Ye et al. 2024; Cheng et al. 2023) and explored non-autoregressive decoding (Sun and Yang 2023; Xiao et al. 2024). A key limitation is their frequent dependence on auxiliary heuristics for complex instances, which can harm the efficiency. 3) Hybrid Solvers merge both paradigms, often using neural networks to enhance classical heuristics, like generating candidate solutions (Xin et al. 2021) or large-scale neighbour search (Hottung and Tierney 2020; Hottung, Kwon, and Tierney 2021; Hottung, Bhandari, and Tierney 2021). These approaches (Hottung and Tierney 2020; Hottung, Kwon, and Tierney 2021; Xin et al. 2021; Chalumeau et al. 2023; Ma, Cao, and Chee 2024; Chen et al. 2024) show significant success on combinatorial optimization problems, but their requirements for task-specific training often prevents them from generalizing across different VRP types. To address these generalization issues, recent efforts have shifted towards cross-task learning (Liu et al. 2024; Zhou et al. 2024; Berto et al. 2024; Li et al. 2025a). For instance, some methods use attribute composition to tackle diverse VRPs (Liu et al. 2024), while others use Mixture-of-Experts (MoE) models to balance performance and computational cost (Zhou et al. 2024).

However, many of these models are pre-trained on small-to-medium-scale data and they fail to capture fine-grained, multi-scale information from the inputs. Different from these existing works, our approach leverages the power of U-Nets (Ronneberger, Fischer, and Brox 2015) to capture critical information and lower down computational overhead at the same time, making it especially suitable for handling instances containing thousands of nodes.

Large-scale Neural Solvers for VRPs

Generalizing neural solvers into large-scale training or inference setting has attracted great attentions over the past years (Luo et al. 2023; Li et al. 2025b; Pan et al. 2025; Luo et al. 2025). The genre along this direction consists of adaptation, divide-and-conquer and local search methods. 1) For adaptation methods (Drakulic et al. 2023; Zhou et al. 2023a; Luo et al. 2023; Li et al. 2025b), neural solvers are often trained on instances with smaller node numbers (e.g., 100 nodes) and during the testing stage they are broad-casted into the settings with much larger node numbers. During training, different regularization methods (e.g., local reconstruction in (Li et al. 2025b)) or architectures (e.g., heavy decoder in (Luo et al. 2023)) are adopted in order to enhance generalization abilities when these models meet unseen large-scale instances. 2) The divide-and-conquer methods (Pan et al. 2023b; Ye et al. 2024; Zong et al. 2022; Hou et al. 2023; Luo et al. 2025; Pan et al. 2025) often takes a two stage workflow where in the first stage instances are split into smaller sub-graphs and once these sub-problems are solved, these partial solutions will merge into a complete solution. In (Pan et al. 2025), a neural policy module is designed for fabricating more nuanced boundaries between clusters. 3) Different from the above two categories, the family of local search methods (Hottung and Tierney 2020; Hottung, Kwon, and Tierney 2021; Huang et al. 2023) restricts the available choices of next node from the whole graph into

a few neighbors. For instance, in (Hottung, Bhandari, and Tierney 2021) a variational auto-encoder is utilized to map instances into hidden feature space and unconstrained continuous optimization is applied for searching. (Huang et al. 2023) further generalizes large neighbor searching into general integer programming problems.

Different from previous approaches that focus on training separate models for each task, we propose a unified solver. Our model is pre-trained across multiple VRPs with small-scale node numbers and then fine-tuned on instances with thousands of nodes, leading into a single, powerful architecture that handles diverse routing problems simultaneously.

Preliminaries

Definitions of VRPs

The VRPs take the format of fully connected graphs denoted by $G = (V, E)$ where $V = \{v_0, \dots, v_n\}$ is the node set consists of n customer nodes plus one depot node v_0 . The $E = \{e_{ij}, i, j = 0, \dots, n\}$ here is the edge set and in our case it consists of edges that link each pair of nodes without self-loops. Each edge e_{ij} is assigned with a cost denoted by c_{ij} . Depending on the type of tasks, we can further acquire two sets of features: *static* features and *dynamic* features. The *static* features of each node v_i consist of coordinates (x_i, y_i) , demand d_i and time window $[t_1^i, t_2^i]$. On the other hand, the *dynamic* features for each time step t consist of remaining capacity c_t , current time t_t , current trajectory length l_t and a binary indicator o_t for whether the current route is open. Note that the *static* features are primarily processed and saved by the encoder module while the *dynamic* features are fed into the decoder at each step to inform the policy.

The Training of Neural Solvers

Following the paradigm of previous works (Kwon et al. 2020; Liu et al. 2024; Zhou et al. 2024), we frame the VRP as a sequential decision-making problem and train our model using reinforcement learning. To be specific, for a generated trajectory τ over the graph G , we modularize the whole process as:

$$p_\theta(\tau|G) = \prod_{t=1}^T p_\theta(a_t|a_{t-1}, G), \quad (1)$$

where θ , T and a_t represent model parameter, total time step and actions taken at the t -th time step. The objective function is defined as:

$$\mathcal{L} = E_{\tau \sim p_\theta(\tau|G)}[R(\tau)], \quad (2)$$

where $R(\tau)$ is the length of trajectory τ . By virtue of reinforcement gradient (Williams 1992), the gradient of loss function takes the following format:

$$\nabla_\theta \mathcal{L} = \frac{1}{N} \sum_{i=1}^N (R(\tau^i) - b^i(G)) \nabla_\theta \log p_\theta(\tau^i|G), \quad (3)$$

where τ_i represents the i -th trajectory and $b^i(G)$ is the introduced baseline term for stabilizing the optimization stage. Note that for some other models like MVMoE(-L) (Zhou et al. 2024), extra regularization loss like expert balancing loss may also exist.

Methods

Our proposed Scale-Net, illustrated in Figure 1, is a unified solver built upon a hierarchical U-Net architecture designed to overcome the scaling and information-bottleneck challenges inherent in previous models. The framework consists of three key stages. First, a hierarchical encoder constructs a multi-scale representation by progressively down-sampling the input graph into a nested series of coarser sub-graphs. Second, a reconstructive decoder up-samples the representation, integrating high-resolution features from the encoder via skip connections. Finally, the policy decoder aggregates these multi-scale features to inform its auto-regressive decision-making process. The following sections provide a detailed description of each component.

Embedding Layer

Given a VRP instance $G = (V, E)$ consists of depot node v_0 and a group of customer nodes $\{v_i; i = 1, \dots, n\}$ where each node is assigned with a 2D coordinate (x_i, y_i) , demand d_i and time window $[t_1^i, t_2^i]$, we concatenate these inputs and embed them into high-dimensional feature space via a shared linear transform layer:

$$\mathbf{h}_i^{(0)} = \text{Embedding}([x_i, y_i, d_i, t_1^i, t_2^i]). \quad (4)$$

Note that the depot node v_0 does not have demand and time window so for this node we only embed its coordinate-level information. After this, we concatenate features between depot node and customer nodes to formalize a new hidden feature vector:

$$\mathbf{H}^{(0)} = [\mathbf{h}_0^{(0)}; \mathbf{h}_1^{(0)}; \dots; \mathbf{h}_n^{(0)}]. \quad (5)$$

This matrix $\mathbf{H}^{(0)}$ serves as the initial input to the first layer of our hierarchical U-Net encoder.

Pooling Layer of U-Net

Standard neural solvers, such as POMO-MTL (Liu et al. 2024) and MVMoE(-L) (Zhou et al. 2024), feed the initial node embeddings directly into a stack of Transformer layers. This approach, however, presents two significant challenges. Firstly, it creates an *information bottleneck problem*: the decoder module makes decisions based on a single level of feature representation, lacking access to a richer hierarchy of global and local structural information. Secondly, the $O(N^2)$ complexity of the self-attention mechanism makes this design computationally infeasible for large-scale problems with thousands of nodes.

To address both issues, we design a hierarchical encoder based on the U-Net architecture (Ronneberger, Fischer, and Brox 2015). Instead of processing the entire graph at every layer, our encoder iteratively down-samples the graph to build a multi-scale representation. This process is composed of a series of down-sampling blocks followed by a final bottleneck layer. Specifically, the encoder path consists of L down-sampling blocks. Let $\mathbf{H}^{(l)}$ be the matrix of node embeddings and $V^{(l)}$ be the set of active nodes at level l . Initially, at level l , $V^{(0)}$ contains all $n + 1$ nodes from the input graph. Each block performs two sequential operations: *feature extraction* and *graph pooling*.

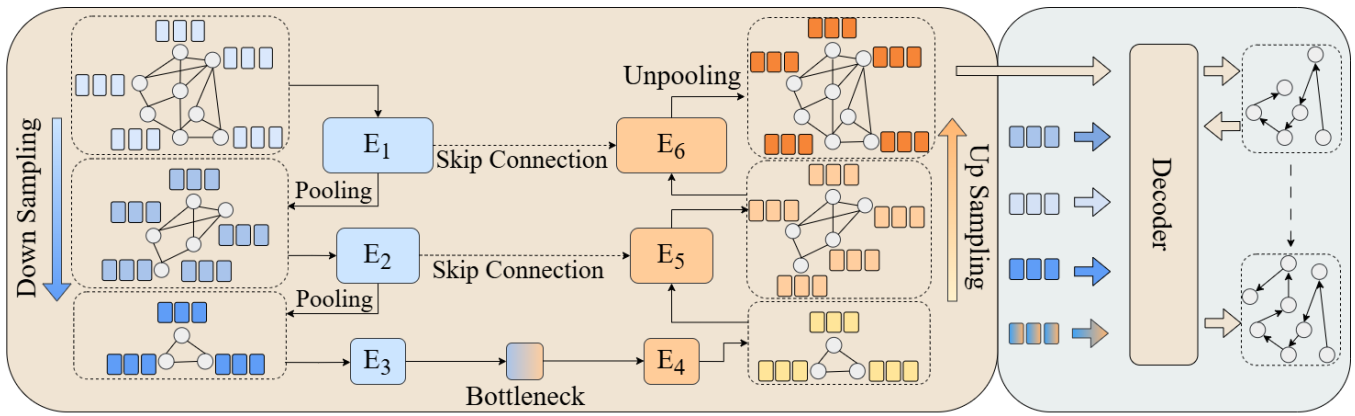


Figure 1: The architecture of our proposed Scale-Net. The encoder path (left) progressively down-samples the input graph using pooling layers to create a nested hierarchy of coarse-to-fine representations. After processing by a bottleneck layer, the decoder path (right) reconstructs the graph features, leveraging skip connections to integrate high-resolution information from the encoder. Finally, the main decoder aggregates these multi-scale features to auto-regressively construct the solution routes. E_i here denotes the i -th layer in the encoder module.

At a given level l , an attention layer operates exclusively on the active nodes $V^{(l)}$ and their features $\mathbf{H}^{(l)}$. This allows nodes within the current sub-graph to exchange information, producing a set of refined, context-aware embeddings $\hat{\mathbf{H}}^{(l)}$:

$$\hat{\mathbf{H}}^{(l)} = \text{Attention}(\mathbf{H}^{(l)}, V^{(l)}). \quad (6)$$

To create the next coarser level of the hierarchy, we perform a pooling operation. We first compute a scalar importance score z_i for each node $v_i \in V^{(l)}$ using a linear projection head on its refined embedding:

$$z_i = \text{Proj}(\hat{\mathbf{h}}_i^{(l)}). \quad (7)$$

We then determine the set of nodes $V^{(l+1)}$ based on the top- k_l scores where k_l is a pre-defined hyperparameter for the number of nodes at level $l+1$. The node features for the next level $\mathbf{H}^{(l+1)}$ are the corresponding embeddings from $\hat{\mathbf{H}}^{(l)}$. This process is repeated L times, yielding a nested hierarchy of graph representations $\{\mathbf{H}^{(0)}, \dots, \mathbf{H}^{(L)}\}$, each capturing the graph structure at a different scale. In addition, since the attention block only operates on the progressively smaller sets of active nodes, the computational overhead is reduced.

After the final down-sampling block, the features of the coarsest graph $\mathbf{H}^{(L)}$ are processed by a bottleneck module to produce a single, dense context vector:

$$\mathbf{H}_{\text{condensed}} = \text{Bottleneck}(\mathbf{H}^{(L)}), \quad (8)$$

which encapsulates a global summary of the problem instance, serves as the initial input to the up-sampling.

Unpooling Layer of U-Net

The up-sampling process begins with the condensed feature vector $\mathbf{H}_{\text{condensed}}$. Then it proceeds iteratively from the coarsest level $l = L$ to the finest level $l = 0$. In details, we propagate features from coarser level $\mathbf{H}^{(l+1)}$ to

nodes in $V^{(l)}$ and for nodes in $V^{(l)} \setminus V^{(l+1)}$, we fill in placeholder embeddings (in our case, they are zero-vectors.). To re-integrate high-resolution details, we utilize skip connections to propagate features from corresponding layer of the down-sampling path and then make a concatenation operation:

$$\mathbf{H}_{\text{final}}^{(l)} = \text{Attention}(\text{Concat}(\text{Unpool}(\mathbf{H}^{(l+1)}), \hat{\mathbf{H}}^{(l)})). \quad (9)$$

By repeating this process up to the original graph resolution, we obtain a final set of node embeddings, $\mathbf{H}_{\text{final}}^{(0)}$, where each node's feature vector is enriched with multi-scale context.

Fusion Layer

At each step t of the auto-regressive process, we construct a context vector \mathbf{c}_t that aggregates a comprehensive view of the current state and the problem structure at all scales via a concatenation operation with dynamic feature $\mathbf{h}_{\text{dynamic}}^{(t)}$ at step t :

$$\mathbf{c}_t = \text{Concat}(\mathbf{h}_{\text{current}}^{(t)}, \mathbf{h}_{\text{dynamic}}^{(t)}, \mathbf{H}^{(0)}, \dots, \mathbf{H}^{(L)}), \quad (10)$$

where $\mathbf{h}_{\text{current}}^{(t)}$ is the current node embedding extracted from $\mathbf{H}_{\text{final}}^{(0)}$. This enriched context vector is then used as the query in a final attention block to compute a probability distribution over the available nodes.

Experiments

This section presents a comprehensive empirical evaluation designed to validate the performance, scalability, and generalization capabilities of Scale-Net. To demonstrate its effectiveness as a unified, multi-task solver, we benchmark our framework on a diverse suite of 16 distinct VRP variants spanning a wide range of problem sizes, from small-scale (50-100 nodes) to large-scale (up to 5,000 nodes) instances. All experiments are conducted on a machine equipped with 48 CPUs and four NVIDIA RTX A6000 GPUs, each with 48 GB of CUDA memory.

Type	Model	$n = 1,000$			$n = 2,000$			$n = 3,000$		
		Obj	Gap	Time	Obj	Gap	Time	Obj	Gap	Time
CVRP	OR-Tools	32.218	0.000%	25m	43.692	0.000%	25m	72.186	0.000%	25m
	ReLD-MVMoE-L	34.965	8.607%	26s	50.628	15.945%	69s	83.588	15.946%	148s
	POMO-MTL-1k	32.443	0.773%	33s	45.764	4.808%	112s	72.431	0.417%	162s
	Scale-Net-1k	31.834	-1.090%	27s	45.204	3.541%	83s	79.871	10.693%	138s
	Scale-Net-2k	32.793	1.881%	25s	45.076	3.248%	63s	71.411	-0.989%	133s
	Scale-Net-3k	32.834	2.016%	26s	45.100	3.305%	65s	71.307	-1.136%	135s
VRPTW	OR-Tools	186.405	0.000%	25m	376.931	0.000%	25m	553.302	0.000%	25m
	ReLD-MVMoE-L	204.298	9.949%	35s	422.736	12.494%	88s	621.070	13.592%	202s
	POMO-MTL-1k	183.805	-1.313%	38s	373.877	-0.987%	147s	559.616	2.043%	215s
	Scale-Net-1k	182.960	-1.774%	34s	372.698	-1.283%	120s	554.821	1.120%	188s
	Scale-Net-2k	184.609	-0.806%	34s	369.244	-2.196%	82s	541.382	-1.431%	187s
	Scale-Net-3k	184.760	-0.724%	34s	369.235	-2.200%	85s	540.206	-1.652%	189s
OVRP	OR-Tools	27.576	0.000%	25m	39.184	0.000%	25m	195.178	0.000%	25m
	ReLD-MVMoE-L	30.036	9.028%	28s	44.998	14.861%	69s	215.191	9.492%	159s
	POMO-MTL-1k	29.028	5.351%	28s	42.709	9.024%	108s	222.294	13.262%	171s
	Scale-Net-1k	28.229	2.452%	25s	41.889	6.950%	82s	272.987	38.653%	158s
	Scale-Net-2k	29.034	5.377%	24s	41.789	6.694%	63s	238.612	21.457%	150s
	Scale-Net-3k	29.210	6.038%	25s	41.785	6.677%	64s	236.382	20.265%	150s
VRPL	OR-Tools	34.036	0.000%	25m	52.352	0.000%	25m	369.314	0.000%	25m
	ReLD-MVMoE-L	40.035	17.243%	27s	65.836	25.235%	72s	383.975	3.568%	172s
	POMO-MTL-1k	35.122	3.158%	30s	53.625	2.290%	115s	405.970	9.153%	186s
	Scale-Net-1k	33.983	-0.058%	27s	53.158	1.482%	93s	409.195	10.285%	166s
	Scale-Net-2k	35.188	3.472%	26s	52.455	0.192%	66s	408.520	10.119%	165s
	Scale-Net-3k	35.210	3.571%	27s	52.511	0.304%	68s	405.075	9.150%	165s
VRPB	OR-Tools	32.962	0.000%	25m	48.108	0.000%	25m	268.721	0.000%	25m
	ReLD-MVMoE-L	31.906	-3.155%	27s	47.877	-0.453%	70s	265.551	-1.420%	164s
	POMO-MTL-1k	30.640	-7.003%	29s	44.810	-6.835%	114s	262.150	-2.316%	176s
	Scale-Net-1k	29.846	-9.396%	26s	43.980	-8.545%	92s	263.376	-2.015%	153s
	Scale-Net-2k	30.638	-6.999%	25s	43.597	-9.341%	65s	268.681	-0.085%	154s
	Scale-Net-3k	30.622	-7.041%	26s	43.664	-9.202%	67s	266.544	-0.810%	155s
OVRPTW	OR-Tools	92.051	0.000%	25m	173.529	0.000%	25m	247.902	0.000%	25m
	ReLD-MVMoE-L	104.405	13.745%	33s	199.834	15.165%	84s	287.652	15.640%	195s
	POMO-MTL-1k	93.843	2.060%	36s	174.189	0.161%	139s	268.028	7.701%	205s
	Scale-Net-1k	93.184	1.364%	33s	173.419	-0.287%	105s	256.657	2.978%	175s
	Scale-Net-2k	95.067	3.422%	33s	172.364	-0.880%	79s	246.641	-1.059%	175s
	Scale-Net-3k	95.105	3.467%	33s	172.386	-0.860%	78s	245.782	-1.403%	174s
OVRPB	OR-Tools	28.984	0.000%	25m	42.311	0.000%	25m	65.908	0.000%	25m
	ReLD-MVMoE-L	29.738	2.759%	27s	45.306	7.184%	70s	64.443	-2.434%	157s
	POMO-MTL-1k	28.937	-0.006%	30s	43.138	2.064%	114s	57.616	-12.743%	172s
	Scale-Net-1k	28.221	-2.453%	26s	42.302	0.115%	94s	59.032	-10.588%	150s
	Scale-Net-MTL-2k	28.834	-0.340%	27s	41.807	-1.066%	65s	54.156	-17.968%	144s
	Scale-Net-MTL-3k	28.851	-0.286%	27s	41.810	-1.057%	65s	54.361	-17.662%	144s
OVRPL	OR-Tools	27.204	0.000%	25m	39.328	0.000%	25m	197.772	0.000%	25m
	ReLD-MVMoE-L	30.123	10.737%	27s	46.750	18.886%	71s	218.923	9.567%	169s
	POMO-MTL-1k	29.336	7.840%	29s	43.617	10.920%	114s	225.764	13.123%	187s
	Scale-Net-1k	28.409	4.453%	26s	42.869	9.029%	92s	278.210	38.963%	171s
	Scale-Net-2k	29.380	8.014%	27s	42.893	9.086%	65s	242.583	21.499%	163s
	Scale-Net-3k	29.414	8.137%	26s	42.941	9.210%	66s	239.431	19.825%	163s

Table 1: In distribution (first 6 tasks) and Out-of-distribution (last 10 tasks) performances for large-scale instances ($n = 1,000, 2,000, 3,000$). The p-values of Scale-Net-1k, Scale-Net-2k and Scale-Net-3k with respect to POMO-MTL-1k are 0.0156, 0.0104 and 0.0249, respectively. These show statistical significance of performances improvements.

Type	Model	$n = 1,000$			$n = 2,000$			$n = 3,000$		
		Obj	Gap	Time	Obj	Gap	Time	Obj	Gap	Time
VRPBL	OR-Tools	34.408	0.000%	25m	53.321	0.000%	25m	268.489	0.000%	25m
	ReLD-MVMoE-L	41.137	19.478%	29s	66.097	23.723%	74s	261.721	-3.615%	178s
	POMO-MTL-1k	36.063	5.019%	31s	54.224	1.764%	120s	258.550	-4.453%	196s
	Scale-Net-1k	34.947	1.910%	28s	53.882	1.189%	96s	259.703	-4.175%	169s
	Scale-Net-2k	35.800	4.378%	28s	52.620	-1.122%	69s	264.974	-2.238%	169s
	Scale-Net-3k	35.791	4.374%	29s	52.600	-1.144%	69s	263.463	-2.748%	169s
VRPBTW	OR-Tools	192.325	0.000%	25m	362.567	0.000%	25m	547.855	0.000%	25m
	ReLD-MVMoE-L	216.243	12.847%	36s	415.563	15.013%	91s	627.378	16.401%	215s
	POMO-MTL-1k	194.645	1.340%	39s	366.118	0.849%	154s	558.619	3.156%	233s
	Scale-Net-1k	193.752	0.877%	37s	364.826	0.526%	113s	553.683	2.220%	202s
	Scale-Net-2k	195.582	1.927%	37s	361.998	-0.240%	84s	541.402	-0.120%	201s
	Scale-Net-3k	195.489	1.877%	36s	361.995	-0.245%	85s	540.258	-0.341%	201s
VRPLTW	OR-Tools	198.223	0.000%	25m	379.399	0.000%	25m	543.919	0.000%	25m
	ReLD-MVMoE-L	215.036	8.609%	36s	422.280	12.275%	92s	615.764	13.584%	216s
	POMO-MTL-1k	193.804	-2.429%	40s	373.665	-1.208%	154s	554.089	1.916%	236s
	Scale-Net-1k	193.042	-2.817%	36s	372.621	-1.463%	107s	549.739	1.070%	204s
	Scale-Net-2k	194.661	-1.901%	36s	369.078	-2.392%	86s	536.319	-1.484%	204s
	Scale-Net-3k	194.760	-1.840%	36s	369.152	-2.359%	85s	535.105	-1.720%	202s
OVRPBL	OR-Tools	26.660	0.000%	25m	40.169	0.000%	25m	160.436	0.000%	25m
	ReLD-MVMoE-L	30.413	14.106%	28s	47.580	18.510%	73s	166.424	3.296%	175s
	POMO-MTL-1k	29.552	10.873%	31s	44.543	10.954%	118s	159.925	-0.435%	192s
	Scale-Net-1k	28.688	7.663%	28s	43.605	8.646%	97s	179.165	11.448%	173s
	Scale-Net-2k	29.435	10.460%	28s	43.095	7.370%	68s	162.034	0.759%	168s
	Scale-Net-3k	29.432	10.446%	28s	43.140	7.483%	68s	160.285	-0.319%	167s
OVRPBTW	OR-Tools	95.491	0.000%	25m	170.087	0.000%	25m	250.815	0.000%	25m
	ReLD-MVMoE-L	111.234	16.561%	35s	199.877	17.765%	87s	292.355	16.885%	198s
	POMO-MTL-1k	100.867	5.529%	38s	174.790	2.785%	145s	263.822	5.370%	219s
	Scale-Net-1k	100.258	4.892%	34s	173.988	2.319%	111s	255.396	1.855%	187s
	Scale-Net-2k	102.357	7.114%	35s	173.361	1.963%	81s	247.751	-1.216%	187s
	Scale-Net-3k	102.304	7.062%	34s	173.434	2.008%	82s	246.973	-1.532%	186s
OVRPLTW	OR-Tools	95.031	0.000%	25m	182.703	0.000%	25m	253.948	0.000%	25m
	ReLD-MVMoE-L	105.920	11.559%	34s	199.618	9.700%	87s	292.133	14.876%	197s
	POMO-MTL-1k	95.365	0.315%	38s	173.861	-4.670%	144s	272.543	7.120%	221s
	Scale-Net-1k	94.810	-0.268%	34s	173.146	-5.071%	116s	261.412	2.583%	187s
	Scale-Net-2k	96.621	1.656%	34s	172.071	-5.644%	81s	251.330	-1.395%	187s
	Scale-Net-3k	96.668	1.715%	34s	172.080	-5.635%	83s	250.379	-1.772%	186s
VRPBLTW	OR-Tools	193.039	0.000%	25m	369.980	0.000%	25m	543.943	0.000%	25m
	ReLD-MVMoE-L	215.067	11.618%	38s	422.584	14.676%	15s	629.267	16.359%	228s
	POMO-MTL-1k	193.552	0.227%	41s	372.402	0.608%	160s	560.200	3.094%	249s
	Scale-Net-1k	192.592	-0.236%	38s	371.053	0.271%	110s	555.088	2.134%	217s
	Scale-Net-2k	194.385	0.774%	38s	367.860	-0.587%	89s	542.907	-0.184%	217s
	Scale-Net-3k	194.465	0.816%	38s	367.960	-0.545%	91s	541.626	-0.426%	216s
OVRPBLTW	OR-Tools	93.796	0.000%	25m	179.527	0.000%	25m	253.595	0.000%	25m
	ReLD-MVMoE-L	108.556	15.610%	36s	209.019	16.664%	89s	296.223	16.481%	210s
	POMO-MTL-1k	98.315	4.563%	39s	184.908	3.065%	151s	267.404	5.049%	336s
	Scale-Net-1k	97.693	3.901%	36s	183.042	1.989%	122s	259.156	1.657%	200s
	Scale-Net-2k	99.847	6.213%	36s	181.940	1.384%	84s	251.472	-1.382%	199s
	Scale-Net-3k	99.807	6.167%	36s	182.028	1.427%	87s	250.557	-1.756%	198s

Table 2: In distribution (first 6 tasks) and Out-of-distribution (last 10 tasks) performances for large-scale instances ($n = 1,000, 2,000, 3,000$). The p-values of Scale-Net-1k, Scale-Net-2k and Scale-Net-3k with respect to POMO-MTL-1k are 0.0156, 0.0104 and 0.0249, respectively. These show statistical significance of performances improvements.

Set Ups

Our empirical validation is designed to test Scale-Net’s performance across a wide spectrum of tasks and scales. We use a two-stage training pipeline (pre-training and fine-tuning) followed by a series of evaluation scenarios. In specifics, our training process begins with pre-training on small-scale instances (100 nodes) and then fine-tuning on progressively larger-scale ones. We use the Adam optimizer for all stages. During pre-training, a mixture of 6 VRP variants is utilized. After pre-training, we sequentially fine-tune the model on 1,000, 2,000, and 3,000-node instances using the same 6 VRPs variants. The U-Net pooling schedules are $\{75, 50, 25\}$ for 100-node instances and $\{500, 250, 125\}$ for larger-scale instances. On the other hand, we evaluate the pre-trained model on standard benchmarks, including in-distribution (6 tasks) and zero-shot (10 tasks) generalization tests. We also test on real-world datasets: Set-X from CVRPLIB (Uchoa et al. 2017) and the Solomon benchmark (Solomon 1987) for VRPTW instances. For more details regarding to training and testing, please refer to appendix.

Baselines

The baselines used in our study fall into two categories: traditional heuristic solvers and neural solvers. For traditional solvers, we rely on HGS (Vidal 2022), LKH3 (Helsgaun 2017) and OR-Tools (Perron and Didier 2024). For neural solvers, we adopt POMO-MTL (Liu et al. 2024), MVMoE(-L) (Zhou et al. 2024) and ReLD-MVMoE-L (Huang et al. 2025). Due to page limit, we move detailed descriptions into the appendix part.

Results

Results for Large-Size Settings After the pre-training stage on 100 node size, we further fine-tune the pre-trained model on 1,000, 2,000 and 3,000 nodes with 6 VRP tasks. Then, we test its performances from 1,000 to 5,000 nodes on all of the 16 VRP tasks. As shown in Table. 1 and Table. 2, the Scale-Net-1k model is dominant, achieving the best performance on all 16 VRP tasks. To ensure a fair comparison, we also fine-tune the original POMO-MTL on 1,000-node instances (denoted by POMO-MTL-1k). The superior results of our model validate that the hierarchical U-Net provides benefits for large-scale problems under both in-distribution and out-of-distribution testing scenarios. Note that fine-tuning baselines like POMO-MTL on 2k/3k nodes will produce out-of-memory problem so we don’t include results for them.

For $n = 2,000$ and $n = 3,000$, models continue to exhibit their advantages, achieving best performances on 16 out of 16 and 10 out of 16 tasks, respectively. We also observe that a model fine-tuned on larger-scale instances can sometimes outperform a model tuned on a smaller scale. For example, on several 2,000-node tasks (e.g., OVRP, VRPTW, OVRPTW), the Scale-Net-3k model achieves better solutions and lower gaps than the Scale-Net-2k. This suggests that training on more challenging, larger instances can sometimes acquire a more generalizable capability. Conversely, we notice that a model fine-tuned on a larger-scale setting (e.g., 3,000 nodes) can exhibit degraded performances when

tested on smaller-scales (e.g., 1,000 nodes). This is similar to the so called catastrophic forgetting phenomena (Wang et al. 2024) in continual learning where the model adapts its parameters to the new data distribution at the expense of its performance on the original distribution. Addressing this phenomenon remains an important direction for future work.

Besides these, we also include experimental results on large-scale instances in Set-X (Uchoa et al. 2017) which consists of problem instances with node sizes ranging from 500 to 1,000. As shown in Table. 7, Appendix, Scale-Net significantly surpasses all other methods, validating its practical effectiveness on real-world problem sets. Note that in order to produce a fair and comprehensive comparison, we also include the results of POMO-MTL-1k.

Sensitivity Analysis for Pooling Sizes Since the pooling size scheduler is a vital module in the U-Net, we also analyze the model’s sensitivity with respect to different choices of pooling schedule. Our findings in Table. 8, Appendix show that Scale-Net is highly robust, as different chosen schedules (in our experiment, we adopt $\{200, 100, 50\}$, $\{300, 150, 75\}$ and $\{500, 250, 125\}$) produce nearly identical results with a low standard deviation of 0.114. Furthermore, we found that performances can be further enhanced with an adaptive schedule tailored to each instance’s size. In Table. 9, Appendix, we can observe that after selecting specific pooling size for each node scale, this adaptive approach reduces the overall optimality gap on CVRPLIB from 23.45% to 21.41%, thereby indicating that the current fixed pooling strategy can be further optimized in terms of a more dynamic and fine-grained manner.

Cross-Size Generalization Analysis To further validate the effectiveness of our proposed framework, we also conduct experiments on out-of-distribution testings on node size. In details, we test models on $n = 4,000$ and $n = 5,000$ across 16 VRPs problems. From results in Table. 10, Appendix, our models consistently achieve lower objective values and provide much faster inference speed.

Conclusions

We introduced Scale-Net, a unified neural solver for addressing the critical challenges of scalability and cross-task generalization in Vehicle Routing Problems. By integrating a hierarchical U-Net module into the encoder, our architecture is able to capture rich, multi-scale features from the input graph while alleviating the computational burden of standard attention mechanisms in the encoder module. Our experiments, spanning 16 VRP variants and problem sizes from 50 to 5,000 nodes, validate that Scale-Net achieves state-of-the-art performance. However, this work contains limitations: 1) To manage computational costs, we limited the POMO size, which may harm performances. 2) While our model demonstrates effective training on instances up to 3,000 nodes, extending this capability to the 10,000-node scale and beyond remains a significant challenge. 3) The current U-Net pooling schedule is manually defined. Future works can develop a learnable mechanism that can adaptively choose pooling strategy based on each instance.

Acknowledgements

This research/project is supported by the National Research Foundation, Singapore under its AI Singapore Programme (AISG Award No: AISG3-RP-2022-031), the MTI under its AI Centre of Excellence for Manufacturing (AIMfg) (Award W25MCMF014), and the College of Computing and Data Science, Nanyang Technological University. We also want to express our sincere thanks to all reviewers and area chair for their constructive engagement and valuable suggestions during the rebuttal stage.

References

- Berto, F.; Hua, C.; Zepeda, N. G.; Hottung, A.; Wouda, N.; Lan, L.; Tierney, K.; and Park, J. 2024. RouteFinder: Towards Foundation Models for Vehicle Routing Problems. In *ICML 2024 Workshop on Foundation Models in the Wild*.
- Cattaruzza, D.; Absi, N.; Feillet, D.; and González-Feliu, J. 2017. Vehicle routing problems for city logistics. *EURO Journal on Transportation and Logistics*, 6(1): 51–79.
- Chalumeau, F.; Surana, S.; Bonnet, C.; Grinsztajn, N.; Pretorius, A.; Laterre, A.; and Barrett, T. 2023. Combinatorial optimization with policy adaptation using latent space search. *Proceedings of the 37th Advances in Neural Information Processing Systems (NeurIPS)*.
- Chen, J.; Wang, J.; Zhang, Z.; Cao, Z.; Ye, T.; and Chen, S. 2024. Efficient meta neural heuristic for multi-objective combinatorial optimization. *Proceedings of the 38th Advances in Neural Information Processing Systems (NeurIPS)*.
- Cheng, H.; Zheng, H.; Cong, Y.; Jiang, W.; and Pu, S. 2023. Select and optimize: Learning to solve large-scale tsp instances. In *Proceedings of the 26th International Conference on Artificial Intelligence and Statistics (AISTATS)*, 1219–1231.
- Drakulic, D.; Michel, S.; Mai, F.; Sors, A.; and Andreoli, J.-M. 2023. Bq-nc0: Bisimulation quotienting for efficient neural combinatorial optimization. *Advances in Neural Information Processing Systems*, 36: 77416–77429.
- Gao, H.; and Ji, S. 2019. Graph u-nets. In *international conference on machine learning*, 2083–2092. PMLR.
- Ge, Y.; Li, H.; and Tuzhilin, A. 2019. Route recommendations for intelligent transportation services. *IEEE Transactions on Knowledge and Data Engineering*, 33(3): 1169–1182.
- Helsgaun, K. 2017. An extension of the Lin-Kernighan-Helsgaun TSP solver for constrained traveling salesman and vehicle routing problems. *Roskilde: Roskilde University*, 12: 966–980.
- Hottung, A.; Bhandari, B.; and Tierney, K. 2021. Learning a Latent Search Space for Routing Problems using Variational Autoencoders. In *International Conference on Learning Representations*.
- Hottung, A.; Kwon, Y.-D.; and Tierney, K. 2021. Efficient Active Search for Combinatorial Optimization Problems. In *Proceedings of the 9th International Conference on Learning Representations (ICLR)*.
- Hottung, A.; and Tierney, K. 2020. Neural large neighborhood search for the capacitated vehicle routing problem. In *Proceedings of the 25th European Conference on Artificial Intelligence*, volume 313, 443–450. IOS Press.
- Hou, Q.; Yang, J.; Su, Y.; Wang, X.; and Deng, Y. 2023. Generalize learned heuristics to solve large-scale vehicle routing problems in real-time. In *The Eleventh International Conference on Learning Representations*.
- Huang, T.; Ferber, A. M.; Tian, Y.; Dilkina, B.; and Steiner, B. 2023. Searching large neighborhoods for integer linear programs with contrastive learning. In *International conference on machine learning*, 13869–13890. PMLR.
- Huang, Z.; Zhou, J.; Cao, Z.; and Xu, Y. 2025. Rethinking Light Decoder-based Solvers for Vehicle Routing Problems. In *International Conference on Learning Representations*.
- Kim, M.; Park, J.; and Park, J. 2022. Sym-nc0: Leveraging symmetry for neural combinatorial optimization. *Proceedings of the 36th Advances in Neural Information Processing Systems (NeurIPS)*.
- Kool, W.; van Hoof, H.; and Welling, M. 2019. Attention, Learn to Solve Routing Problems! In *Proceedings of the 7th International Conference on Learning Representations (ICLR)*.
- Kwon, Y.-D.; Choo, J.; Kim, B.; Yoon, I.; Gwon, Y.; and Min, S. 2020. Pomo: Policy optimization with multiple optima for reinforcement learning. *Proceedings of the 34th Advances in Neural Information Processing Systems (NeurIPS)*.
- Li, H.; Liu, F.; Zheng, Z.; Zhang, Y.; and Wang, Z. 2025a. CaDA: Cross-Problem Routing Solver with Constraint-Aware Dual-Attention. In *Forty-second International Conference on Machine Learning*.
- Li, K.; Liu, F.; Wang, Z.; and Zhang, Q. 2025b. Destroy and Repair Using Hyper-Graphs for Routing. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, 18341–18349.
- Lin, S.; and Kernighan, B. W. 1973. An effective heuristic algorithm for the traveling-salesman problem. *Operations research*, 21(2): 498–516.
- Liu, F.; Lin, X.; Wang, Z.; Zhang, Q.; Xialiang, T.; and Yuan, M. 2024. Multi-task learning for routing problem with cross-problem zero-shot generalization. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (SIGKDD)*, 1898–1908.
- Luo, F.; Lin, X.; Liu, F.; Zhang, Q.; and Wang, Z. 2023. Neural combinatorial optimization with heavy decoder: Toward large scale generalization. *Proceedings of the 37th Advances in Neural Information Processing Systems (NeurIPS)*.
- Luo, F.; Lin, X.; Wu, Y.; Wang, Z.; Xialiang, T.; Yuan, M.; and Zhang, Q. 2025. Boosting Neural Combinatorial Optimization for Large-Scale Vehicle Routing Problems. In *The Thirteenth International Conference on Learning Representations*.
- Ma, Y.; Cao, Z.; and Chee, Y. M. 2024. Learning to search feasible and infeasible regions of routing problems with flexible neural k-opt. *Proceedings of the 38th Advances in Neural Information Processing Systems (NeurIPS)*.

- Pan, X.; Jin, Y.; Ding, Y.; Feng, M.; Zhao, L.; Song, L.; and Bian, J. 2023a. H-tsp: Hierarchically solving the large-scale traveling salesman problem. In *Proceedings of the 37th AAAI Conference on Artificial Intelligence (AAAI)*, volume 37, 9345–9353.
- Pan, X.; Jin, Y.; Ding, Y.; Feng, M.; Zhao, L.; Song, L.; and Bian, J. 2023b. H-tsp: Hierarchically solving the large-scale traveling salesman problem. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, 9345–9353.
- Pan, Y.; Liu, R.; Chen, Y.; Cao, Z.; and Lin, F. 2025. Hierarchical Learning-based Graph Partition for Large-scale Vehicle Routing Problems. In *Proceedings of the 24th International Conference on Autonomous Agents and Multiagent Systems*, 1604–1612.
- Perron, L.; and Didier, F. 2024. CP-SAT.
- Ronneberger, O.; Fischer, P.; and Brox, T. 2015. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, 234–241. Springer.
- Solomon, M. M. 1987. Algorithms for the vehicle routing and scheduling problems with time window constraints. *Operations research*, 35(2): 254–265.
- Sun, Z.; and Yang, Y. 2023. Difusco: Graph-based diffusion solvers for combinatorial optimization. *Proceedings of the 37th Advances in Neural Information Processing Systems (NeurIPS)*.
- Uchoa, E.; Pecin, D.; Pessoa, A.; Poggi, M.; Vidal, T.; and Subramanian, A. 2017. New benchmark instances for the capacitated vehicle routing problem. *European Journal of Operational Research*, 257(3): 845–858.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. *Proceedings of the 31st Advances in Neural Information Processing Systems (NeurIPS)*.
- Vidal, T. 2022. Hybrid genetic search for the CVRP: Open-source implementation and SWAP* neighborhood. *Computers & Operations Research*, 140: 105643.
- Vinyals, O.; Fortunato, M.; and Jaitly, N. 2015. Pointer networks. *Proceedings of the 29th Advances in Neural Information Processing Systems (NeurIPS)*.
- Wang, L.; Zhang, X.; Su, H.; and Zhu, J. 2024. A comprehensive survey of continual learning: Theory, method and application. *IEEE transactions on pattern analysis and machine intelligence*, 46(8): 5362–5383.
- Williams, R. J. 1992. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning*, 8: 229–256.
- Xiao, Y.; Wang, D.; Li, B.; Wang, M.; Wu, X.; Zhou, C.; and Zhou, Y. 2024. Distilling autoregressive models to obtain high-performance non-autoregressive solvers for vehicle routing problems with faster inference speed. In *Proceedings of the 38th AAAI Conference on Artificial Intelligence (AAAI)*, volume 38, 20274–20283.
- Xin, L.; Song, W.; Cao, Z.; and Zhang, J. 2021. Neurolkh: Combining deep learning model with lin-kernighan-helsgaun heuristic for solving the traveling salesman problem. *Proceedings of the 35th Advances in Neural Information Processing Systems (NeurIPS)*.
- Ye, H.; Wang, J.; Liang, H.; Cao, Z.; Li, Y.; and Li, F. 2024. Glop: Learning global partition and local construction for solving large-scale routing problems in real-time. In *Proceedings of the 38th AAAI Conference on Artificial Intelligence (AAAI)*, volume 38, 20284–20292.
- Zhang, C.; Wu, Y.; Ma, Y.; Song, W.; Le, Z.; Cao, Z.; and Zhang, J. 2023. A review on learning to solve combinatorial optimisation problems in manufacturing. *IET Collaborative Intelligent Manufacturing*, 5(1): e12072.
- Zhou, J.; Cao, Z.; Wu, Y.; Song, W.; Ma, Y.; Zhang, J.; and Xu, C. 2024. MVMoE: Multi-Task Vehicle Routing Solver with Mixture-of-Experts. In *Proceedings of the 41th International Conference on Machine Learning (ICML)*, 61804–61824.
- Zhou, J.; Wu, Y.; Cao, Z.; Song, W.; Zhang, J.; and Chen, Z. 2023a. Learning large neighborhood search for vehicle routing in airport ground handling. *IEEE Transactions on Knowledge and Data Engineering*, 35(9): 9769–9782.
- Zhou, J.; Wu, Y.; Song, W.; Cao, Z.; and Zhang, J. 2023b. Towards omni-generalizable neural methods for vehicle routing problems. In *Proceedings of the 40th International Conference on Machine Learning (ICML)*, 42769–42789.
- Zong, Z.; Wang, H.; Wang, J.; Zheng, M.; and Li, Y. 2022. Rbg: Hierarchically solving large-scale routing problems in logistic systems via reinforcement learning. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 4648–4658.