
Learning to Handle Constraints in Routing Problems via a Construct-and-Refine Framework

Anonymous Author(s)

Affiliation

Address

email

Abstract

Neural solvers have achieved impressive progress on simple routing problems via data-driven training, but often struggle with complex constraints. We rethink the popular single-paradigm neural solvers and identify paradigm-inherent limitations: construction solvers suffer from inflexible stepwise feasibility, and improvement solvers easily get stuck in infeasible searches with long runtimes. However, these paradigms are naturally complementary: construction efficiently provides strong initial solutions that help improvement rapidly reach feasible, high-quality solutions. Motivated by this, we propose Construct-and-Refine (CaR), the first generic neural framework for efficient constraint handling, compatible with existing construction and improvement solvers. To promote synergistic paradigm integration, we introduce a joint training framework with bespoke losses to generate diverse, high-quality, (near)-feasible solutions that are refined by a light improvement process (e.g., only 10 steps down from 5k). We also present the first study of a shared encoder for cross-paradigm representation learning in handling complex constraints. Extensive experiments on hard-constrained TSPTW and CVRPBLTW demonstrate that CaR achieves superior feasibility, solution quality, and efficiency compared to both traditional and neural state-of-the-art solvers.

1 Introduction

Vehicle routing problems (VRPs) often involve complex real-world constraints [1, 2]. For decades, the Operations Research (OR) community has put remarkable efforts into designing hand-crafted heuristics to handle these constraints and approximate near-optimal solutions, leading to powerful solvers such as OR-Tools [3] and LKH3 [4]. Recently, neural combinatorial optimization (NCO) solvers [5–7] have enriched the landscape for solving VRPs, which learn deep models to solve VRPs in a data-driven manner, thereby reducing dependence on domain knowledge. Benefiting from GPU parallelism, they are often fast while maintaining competitive solution quality. In general, existing neural solvers fall into two paradigms: 1) *Construction solvers* (e.g. [8, 9]), which construct solutions node by node from scratch, excelling in efficiency but prone to local optima; 2) *Improvement solvers* (e.g. [10, 11]), which iteratively refine a complete solution, exploring a broader search space but facing longer runtime. However, when they come to handling complex VRP constraints, both paradigms face critical inherent limitations that significantly hinder both feasibility and optimality.

Construction solvers are inherently built upon a *stepwise feasibility satisfaction* process, where constraints must be enforced at each node selection step. While feasibility masking addresses this by excluding the invalid candidates and works well for simple cases like the Capacitated VRP (CVRP), it fails to handle more complex ones such as the Traveling Salesman Problem with Time Windows (TSPTW), where computing accurate masks itself is NP-hard thus intractable [12]. Recent attempts, such as lookahead strategies [13] and learnable approximated masking [12], alleviate this but come with high computational costs and still yield large infeasibility rates. Moreover, as discussed

in Section 4, even for complex VRPs where feasibility masking is tractable, such as CVRP with backhaul, duration limit, and time window constraints (CVRPBLTW), enforcing multiple constraints through such stepwise masking can overly restrict learning, thereby degrading model performance and making it more prone to local optima. Hence, achieving stepwise feasibility satisfaction while maintaining optimality remains a critical challenge in current construction solvers.

On the other hand, the performance of improvement solvers on complex-constrained VRPs remains largely unexplored. Akin to the construction methods, they struggle with accurate feasibility masking for complex local search operators (e.g., k -opt). While prior work [11] showed that learning-guided infeasible region exploration improves performance in simpler constrained CVRP, it lacks feasibility guarantees and remains potentially time-consuming for complex VRPs. In more complex VRPs like CVRPBLTW, these improvement solvers face prolonged searches in infeasible regions (see Table 1), further amplifying the inherent inefficiencies in exploring complex search spaces.

As a result, both solver paradigms struggle with constraint handling when used alone, often exhibiting inefficiencies, poor adaptability to complex constraints, or high infeasibility and suboptimality. However, they offer complementary strengths: improvement solvers excel at exploring local neighborhoods to repair infeasible or suboptimal solutions, while construction solvers provide strong initializations that allow improvement to rapidly reach feasible, high-quality solutions. This raises our research question: *can we unify the strengths of the two paradigms to handle constraints without sacrificing efficiency?*

Motivated by this, we present **Construct-and-Refine (CaR)**, the first generic neural framework that promotes efficient constraint handling through paradigm integration. Unlike existing methods that are often paradigm-limited, inefficient or lack generality across VRPs with varying constraint complexities, CaR unites the efficiency of construction with the flexibility of improvement, enabling efficient constraint handling for general complex VRPs where feasibility masking is either intractable (TSPTW) or overly restrictive (CVRPBLTW). To enhance the synergy between paradigms, CaR adopts an end-to-end joint training framework with bespoke loss functions that enable the construction module to generate diverse, high-quality, (near)-feasible initial solutions, which are then passed to a light refinement module, reducing the number of improvement steps required, e.g., from 5k to 10. Moreover, to further balance computational cost and enhance synergy, we take the first step toward exploring how a unified encoder across construction and refinement modules can enable more effective cross-paradigm representation learning. Lastly, we show that CaR is a generic framework compatible with existing state-of-the-art construction and improvement methods for efficient constraint handling.

Our contributions are: 1) We introduce Construct-and-Refine (CaR), the first neural framework that enables efficient constraint handling through paradigm integration, excelling in solving complex VRPs where existing neural solvers fail to solve effectively; 2) We explore cross-paradigm representation learning by unifying encoders across the two paradigms, yielding notable benefits for hard-constrained scenarios; 3) We conduct extensive experiments on representative hard-constrained TSPTW and CVRPBLTW, where masking is intractable or overly restrictive. Notably, CaR is the first generic neural Construct-and-Refine framework for efficient constraint handling in VRPs, achieving substantial reduction in infeasibility, especially where construction alone fails, while simultaneously improving solution quality and runtime efficiency via light refinement.

2 Preliminaries

VRP variants. VRP is defined over a directed graph $\mathcal{G} = \{\mathcal{V}, \mathcal{E}\}$, where \mathcal{V} contains n customer nodes $\{v_1, v_2, \dots, v_n\}$ (for both TSP and VRP variants) and a depot $\{v_0\}$ (for VRP only), and $e(v_i, \rightarrow v_j)$ (or e_{ij}) $\in \mathcal{E}$ represents an edge from node v_i to v_j ($i \neq j$) weighted by the 2D Euclidean distances. It aims to minimize the total cost of a solution while satisfying variant-specific constraints. This paper focuses on two representative complex VRPs: TSPTW with time window constraints, and CVRPBLTW with backhaul, duration limit, and time window constraints. See Appendix B for more details on the constraints. Neural VRP solvers usually handle these constraints by feasibility masking, which is tractable and effective in simpler cases like CVRP. However, they struggle with harder variants, e.g., CVRPBLTW with multiple constraints, or TSPTW, where masking is NP-hard since computing accurate masks requires evaluating future time-related feasibility for all possible actions. We further discuss the limitations of the existing feasibility masking in Appendix F.1.

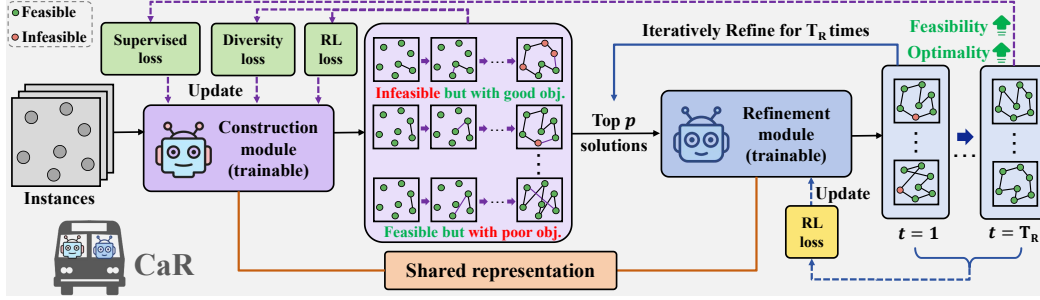


Figure 1: Overall framework of CaR, where the construction module provides *diverse* near-feasible and/or high-quality solutions for the refinement module to generate better solutions.

MDP formulations. We model construction and improvement as Constrained Markov Decision Processes (CMDPs), defined by the tuple $(\mathcal{S}, \mathcal{A}, \mathcal{P}, \mathcal{R}, \mathcal{C})$, where \mathcal{S} is the state space, \mathcal{A} the action space, \mathcal{P} the transition function, \mathcal{R} the reward, and \mathcal{C} the constraint function over m constraints. The objective is to learn a policy π_θ maximizing expected reward while satisfying constraints:

$$\max_{\theta} \mathcal{J}(\pi_\theta) = \mathbb{E}_{\tau \sim \pi_\theta} [\mathcal{R}(\tau | \mathcal{G})], \text{ s.t. } \pi_\theta \in \Pi_F, \Pi_F = \{\pi \in \Pi \mid \mathcal{J}_C(\pi_\theta) = \mathbf{0}^m\}, \quad (1)$$

where Π_F denotes the feasible policy set. *Construction* solvers build solutions sequentially, with state s_t encoding partial routes, vehicle status, and unvisited nodes; action a_t selects the next node; and the reward is the negative tour cost, $\mathcal{R}(\tau | \mathcal{G}) = -C(\tau)$. In contrast, *improvement* solvers refine complete solutions, where s_t includes the current and best-so-far solutions; actions apply local operators (e.g., k -opt, remove-and-reinsert); and the reward is $\mathcal{R}_t = \min[C(\tau_t^*) - C(\tau_{t+1}), 0]$ [14].

Relaxation of CMDP. Following [12], we adopt Lagrangian relaxation to facilitate training by penalizing constraint violations. The objective in Eq. (1) is augmented with a cost term:

$$C(\tau) = \sum_{e_{ij} \in \tau} \left[C_L(e_{ij}) + \sum_{\eta=1}^m C_V^\eta(e_{ij}) \right], \quad (2)$$

where C_L denotes the objective cost (e.g., tour length), and C_V captures violations of m constraints. For instance, if a vehicle arrives at node v_j later than the time window upper bound u_j , the violation cost is $C_V(e_{ij}) = t_j - u_j$. C_V also incorporates the count of violated nodes.

3 Methodology

As discussed above, single-paradigm methods have inherent limitations for constraint handling but can complement each other (see Appendix C for details). On one hand, improvement can rapidly repair construction’s infeasible or suboptimal solutions, due to its strength in exploring local neighborhoods; on the other hand, construction is highly efficient and can provide strong initializations that help improvement rapidly reach feasible, high-quality solutions, complementing improvement’s limitation in inefficiency. To further promote the synergic effects, we present our unified CaR framework for complex VRPs where masking is *intractable* (TSPTW) or *restrictive* (CVRPBLTW).

3.1 Joint training framework of CaR

CaR aims to transcend simple paradigm merging to handle constraints more efficiently than either paradigm alone. We develop a unified training framework to jointly optimize the two policies in CaR.

Relaxation of CMDP in CaR. Building on the success of relaxed CMDP formulation for construction [12], we extend it to refinement to better balance objective and feasibility. Unlike prior infinite-horizon approaches [11, 15], we adopt a fixed rollout limit T_R , treating steps equally to align with CaR’s lightweight and efficient refinement design.

Collaborative training framework. To integrate construction and refinement effectively, we design a joint training framework that optimizes both processes simultaneously in each gradient step, allowing them to co-evolve. As illustrated in Figure 1, for each batch of training instances \mathcal{G} , the construction module first generates a small set of diverse, high-quality initial solutions in parallel. These solutions

are then refined by a lightweight neural improvement process within T_R steps (i.e., $T_R = 10$ significantly fewer than the thousands of iterations in conventional improvement methods), enabling rapid enhancement of high-potential candidates. The refined outputs then supervise construction, promoting collaborative correction of infeasibility and sub-optimality.

Construction module loss. The policy π_θ^C is trained via REINFORCE [16], using the loss function:

$$\mathcal{L}_{\text{RL}}^C = \frac{1}{S} \sum_{i=1}^S \left[\left(\mathcal{R}(\tau_i) - \frac{1}{S} \sum_{j=1}^S \mathcal{R}(\tau_j) \right) \log \pi_\theta^C(\tau_i) \right], \quad (3)$$

where the solution probability is factorized as $\pi_\theta^C(\tau) = \prod_{t=1}^{|\tau|} \pi_\theta^C(e_t | \tau_{<t})$, with $\tau_{<t}$ denoting the partial solution prior to selecting edge e_t at step t . We employ a group baseline with diverse rollouts to reduce REINFORCE variance. For simpler variants like CVRP, S solutions are generated via POMO’s multi-start strategy [9], while for time-constrained variants (e.g., TSPTW), we sample S solutions to avoid infeasibility (see Appendix F.2). To compensate for reduced diversity due to the removal of the multi-start mechanism and to enhance the diversity of initial constructed solutions for refinement, we introduce an auxiliary entropy-based diversity loss:

$$\mathcal{L}_{\text{DIV}} = - \sum_{t=1}^{|\tau|} \pi_\theta^C(e_t | \tau_{<t}) \log \pi_\theta^C(e_t | \tau_{<t}), \quad (4)$$

which largely encourages policy exploration during RL training. To avoid inefficiency, we evaluate candidates using the cost in Eq. (2), and only feed the top p high-quality candidates to subsequent refinement. If the refinement module improves a constructed solution (indicated by $\mathbb{I} = 1$), the best-refined solution τ^* is used as a pseudo ground truth to supervise π_θ^C :

$$\mathcal{L}_{\text{SL}} = -\mathbb{I} \cdot \sum_{t=1}^{|\tau^*|} \log \pi_\theta^C(e_t^* | \tau_{<t}^*), \quad (5)$$

where \mathbb{I} indicates whether such refinement led to improvement. The final construction loss integrates three components, i.e., $\mathcal{L}(\theta^C) = \mathcal{L}_{\text{RL}}^C + \alpha_1 \mathcal{L}_{\text{DIV}} + \alpha_2 \mathcal{L}_{\text{SL}}$.

Refinement module loss. The refinement policy π_θ^R iteratively improves solutions over T_R steps, with probability of the refined solution at step t is factorized as $\pi_\theta^R(\tau_t) = \prod_{\kappa=1}^{\mathbf{K}} \pi_\theta^R(a_\kappa | a_{<\kappa}, \tau_{t-1})$, where \mathbf{K} denotes the total number of sequential refinement moves/actions, with further details in Appendix D. The RL loss $\mathcal{L}_{\text{RL}}^R(t)$ for refinement is computed at each step t using the REINFORCE algorithm in Eq. (3), where S is replaced by p , since only p solutions are refined. The final refinement loss is defined as the average across all T_R steps: $\mathcal{L}(\theta^R) = \frac{1}{T_R} \sum_{t=1}^{T_R} \mathcal{L}_{\text{RL}}^R(t)$, encouraging each refinement step to contribute meaningfully and improving overall refinement efficiency.

Joint training loss. The joint training loss combines the above two losses, i.e., $\mathcal{L}(\theta) = \mathcal{L}(\theta^C) + \omega \mathcal{L}(\theta^R)$, where ω balances their scales. Such joint loss promotes information exchange between modules, enhancing synergy in collaboratively handling complex constraints.

3.2 Unified model architecture

To reduce overhead and enhance synergy, we take the first step in exploring shared encoders for more effective cross-paradigm representation learning for hard-constrained VRPs.

Encoder. Given an instance batch $\{\mathcal{G}_i\}_{i=1}^B$, both paradigms learn to obtain high-dimensional node embeddings h_i via encoders. For CVRPBLTW, each node v_i is represented by its coordinates, demand (i.e., linehaul or backhaul), time window, and duration limit, i.e., $f_i^n = \{x_i, y_i, q_i, l_i, u_i, \ell\}$. Unlike construction, refinement also incorporates solution features by encoding the sequential structure via positional information. To support both paradigms, we use a shared 6-layer Transformer encoder [9] with multi-head attention. Positional encoding, required only in refinement, is injected using cyclic positional encoding (CPE) via the synthesis-attention (Syn-Att) mechanism [17]. Afterwards, a multi-layer perceptron (MLP) fuses node-level attention scores a^n and solution-level scores a^s from positional embedding vectors through element-wise aggregation (see details in Figure 4).

Decoder. The decoder generates action probabilities from node representations, selecting the next node for construction or the modification for refinement. In CaR, we retain the original neural solver

Table 1: Results on constrained VRPs: best are **bolded**; best within 1 min are shaded to show solver efficiency.

	Method	#Params	Paradigm §	n=50				n=100			
				Obj. ↓	Gap ↓	Infbs% ↓	Time	Obj. ↓	Gap ↓	Infbs% ↓	Time
TSPTW	LKH3 (max trials = 100)	/	I	25.590	0.004%	11.88%	7m	46.625	0.103%	31.05%	27m
	LKH3 (max trials = 10000)	/	I	25.611	◇	0.12%	7h	46.858	◇	0.13%	1.4d
	OR-Tools†	/	I	25.763	-0.001%	65.72%	2.4h	46.424	0.026%	97.45%	12m
	Greedy-C	/	C	26.394	1.534%	72.55%	4.5s	51.945	9.651%	99.85%	11.4s
	POMO	1.25M	L2C	/	/	100.00%	4s	/	/	100.00%	14s
	POMO*	1.25M	L2C	26.222	1.635%	37.27%	4s	47.249	1.959%	38.22%	14s
	POMO* + PIP (greedy)	1.25M	L2C	25.657	0.177%	2.67%	7s	47.372	1.223%	6.96%	32s
	POMO* + PIP (sample 10)	1.25M	L2C	25.650	0.152%	1.87%	1m	47.294	1.030%	4.67%	5.2m
	NeuOpt-GIRE *‡ ($T = 1k$)	0.69M	L2I	25.627	0.061%	0.19%	2.3m	47.011	0.336%	0.13%	5.9m
	NeuOpt-GIRE *‡ ($T = 5k$)	0.69M	L2I	25.617	0.028%	0.02%	11.6m	46.913	0.123%	0.02%	30m
	CaR-POMO ($T_R = 5$)	1.64M	L2(C+I)	25.619	0.034%	0.02%	15s	47.278	1.065%	4.20%	36s
	CaR-POMO ($T_R = 20$)	1.64M	L2(C+I)	25.614	0.014%	0.01%	51s	47.001	0.406%	2.34%	2.1m
	CaR-PIP ($T_R = 5$)	1.64M	L2(C+I)	25.613	0.010%	0.02%	17s	47.000	0.315%	0.10%	58s
	CaR-PIP ($T_R = 20$)	1.64M	L2(C+I)	25.612	0.005%	0.00%	52s	46.923	0.146%	0.02%	2.4m
CVRPBLTW	OR-Tools (short)	/	I	14.890	1.402%	0.00%	10.4m	25.979	2.518%	0.00%	20.8m
	OR-Tools	/	I	14.677	◇	0.00%	1.7h	25.342	◇	0.00%	3.5h
	POMO	1.25M	L2C	15.999	9.169%	0.00%	2s	27.046	7.004%	0.00%	4s
	POMO+EAS+SGBS*	1.25M	L2C	15.156	3.263%	0.00%	10.3m	25.558	0.854%	0.00%	1h
	NeuOpt-GIRE *‡ ($T = 1k$)	0.69M	L2I	14.521	1.329%	33.80%	1.1m	24.832	4.290%	56.30%	2.9m
	NeuOpt-GIRE *‡ ($T = 5k$)	0.69M	L2I	14.201	-1.163%	27.30%	5.5m	24.237	-0.533%	41.20%	15m
	POMO*	1.25M	L2C	14.873	2.310%	0.00%	2s	24.592	-1.645%	0.00%	4s
	CaR (k -opt) ($T_R = 5$)	1.64M	L2(C+I)	14.872	2.271%	0.00%	3s	24.597	-1.674%	0.00%	5s
	CaR (k -opt) ($T_R = 20$)	1.64M	L2(C+I)	14.844	2.114%	0.00%	8s	24.585	-1.724%	0.00%	17s
	CaR (R&R) ($T_R = 5$)	1.72M	L2(C+I)	14.725	1.328%	0.00%	3s	24.552	-1.835%	0.00%	6s
	CaR (R&R) ($T_R = 20$)	1.72M	L2(C+I)	14.601	0.463%	0.00%	10s	24.400	-2.448%	0.00%	19s

§ The abbreviations refer to: **I** – Improvement; **L2C** – Learning to Construct; **L2I** – Learning to Improve.

† OR-Tools presolves before search; if it detects infeasibility, it terminates immediately, making runtime shorter than preset.

designs when applied to *one* construction and *one* improvement at a time. To validate generality, we experiment with two construction backbones, POMO [9] and PIP [12], and two refinement backbones, NeuOpt [11] and N2S [17]. To adapt improvement solvers for new variants (e.g., TSPTW, CVRPBLTW), we follow their original design and introduce variant-specific features, such as refinement history and node-level feasibility information (see Appendix D for details), to enhance constraint awareness. While we also explored a unified decoder, results in Figure 7 show degraded performance, suggesting that while a shared encoder benefits representation learning, separate decoders remain important for paradigm-specific optimization – an insight for future reference.

4 Experiments

We now evaluate our proposed Construct-and-Refine (CaR) framework in handling hard-constrained TSPTW and CVRPBLTW instances. Additional results are provided in Appendix F.

Experimental settings. Training instances are generated on the fly as in [12, 18] (see Appendix B). For TSPTW, we mainly focus on the hard variants in [12]. All experiments are conducted on problem sizes $n = 50/100$, following established benchmarks [8, 10]. Models are trained with 20,000 instances per epoch for 5,000 epochs with a batch size of 128 [18]. We set $T_R = 5$ during training. During inference, $8 \times$ augmentation [9] is used to construct initial solutions, followed by T_R -step refinement.

Baseline. We compare CaR with *SoTA* traditional solvers (LKH3 [4], OR-Tools [3], and Greedy heuristics) and neural solvers, including construction methods (POMO [9]+EAS [19]+SGBS [20], PIP [12]) and the improvement solver NeuOpt-GIRE [11]. To ensure fairness, we upgrade NeuOpt-GIRE and POMO-based solvers with our relaxed CMDP formulation (*) and solution-level features (§; see Appendix D). All models are trained to convergence with comparable training budgets.

Evaluation metrics. We evaluate performance using: 1) average solution length (Obj.), the mean length of best feasible solutions; 2) average optimality gap (Gap), the difference from (near-)optimal solutions found by top traditional solvers (LKH3 for TSPTW, OR-Tools for CVRPBLTW, HGS for CVRP, marked with ◇); 3) total inference time (Time) for solving 10,000 TSPTW or 1,000 CVRP/CVRPBLTW instances using single-GPU parallelization; and 4) average infeasible solution rate per instance (Infbs%) after construction and refinement.




Figure 2: Performance over time on TSPTW-100 (left) and CVRPBLTW-100 (right). For CVRPBLTW, POMO+EAS+SGBS and CaR always achieve 0% infeasibility due to feasibility masking. We show the SoTA method for each paradigm for comparison.

4.1 Model performance on constrained VRPs

TSPTW results. We first test CaR on TSPTW, where most neural solvers fail due to the lack of feasibility masking. To verify generality, we use POMO* and PIP as construction backbones, with CaR-POMO training $1.42\times$ (at $n=50$) and $1.65\times$ (at $n=100$) faster than CaR-PIP. As shown in Table 1, CaR-POMO consistently outperforms PIP in both quality and feasibility. On TSPTW-50, CaR reduces infeasibility to 0.00%, outperforming PIP (1.87%) and NeuOpt* (0.02%). On TSPTW-100, CaR lowers PIP’s 4.67% infeasibility to 0.02% and reduces the gap from 1.030% to 0.146%. While NeuOpt* improves with longer runtimes, CaR achieves competitive results with an $8\times$ speedup within a runtime budget of 10 minutes (Figure 2), which aligns with CaR’s aim of efficiency. Notably, CaR surpasses LKH3 and finds feasible solutions even when it fails, highlighting the strength of our cross-paradigm framework under complex constraints.

CVRPBLTW results. On complex CVRPBLTW, feasibility masking filters out over 60% of nodes, severely limiting the search space (Figure 6). Interestingly, removing these masks and applying the relaxed CMDP in POMO significantly improves performance (e.g., CVRPBLTW-100: from 7.004% to -1.645% in Table 1). Unlike TSPTW, NeuOpt* fails in CVRPBLTW with 27-56% infeasibility, while CaR guarantees feasibility as other construction solvers. We compare CaR with the best single-paradigm solvers in Figure 2. CaR achieves best area under the curve, indicating superior efficiency and effectiveness. We also validate CaR with k -opt and R&R, where R&R performs better (-2.448% vs. -1.724%) due to a finer-grained search better suited to multi-constraints variants.

4.2 Effects of the unified representation on handling complex VRPs

We explore the idea of learning unified representations for two paradigms, which by design ensures complementary knowledge sharing, enabling synergistic module functionality when one lacks capability. As shown in Table 2, CaR with a shared encoder and separate decoders outperforms the construction-only model on the hard-constrained *Medium* TSPTW-100 and *Hard* TSPTW-50 cases, with the unified variant achieving the best results. This suggests that CaR not only strengthens standalone construction but also enables cross-paradigm knowledge transfer through shared representations. See Appendix F for more ablation results and further analysis.

Table 2: Effects of unified encoder on TSPTW.

Method	TSPTW-50 Hard		TSPTW-100 Medium	
	Inf%b	Gap	Inf%b	Gap
Construction-only	37.270%	1.635%	0.120%	10.931%
Separate	0.675%	0.199%	0.000%	7.589%
Unified (Ours)	0.010%	0.014%	0.000%	5.815%

5 Conclusion

This paper proposes Construct-and-Refine (CaR), the first neural framework to handle constraints via paradigm integration. CaR jointly learns to construct diverse, high-quality solutions and refine them with a lightweight improvement module, enabling efficient constraint satisfaction. We also explore shared encoders for cross-paradigm representation learning. CaR performs strongly across constrained VRPs, offering key insights: 1) single-paradigm solvers struggle with complex constraints; 2) strict feasibility masking may hurt performance, while relaxed penalties help; 3) combining construction with light refinement efficiently recovers feasible, high-quality solutions; and 4) shared representations boost performance particularly for complex cases. Future work includes applying CaR to more VRPs, integrating diverse solvers, improving scalability, and developing foundation NCO models.

References

- [1] Lu Duan, Yang Zhan, Haoyuan Hu, Yu Gong, Jiangwen Wei, Xiaodong Zhang, and Yinghui Xu. Efficiently solving the practical vehicle routing problem: A novel joint learning approach. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 3054–3063, 2020.
- [2] Lixia Wu, Haomin Wen, Haoyuan Hu, Xiaowei Mao, Yutong Xia, Ergang Shan, Jianbin Zhen, Junhong Lou, Yuxuan Liang, Liuqing Yang, et al. Lade: The first comprehensive last-mile delivery dataset from industry. *arXiv preprint arXiv:2306.10675*, 2023.
- [3] Vincent Furnon and Laurent Perron. Or-tools routing library, 2024. URL <https://developers.google.com/optimization/routing/>.
- [4] Keld Helsgaun. LKH-3 (version 3.0.7), 2017. URL <http://webhotel4.ruc.dk/~keld/research/LKH-3/>.
- [5] Federico Berto, Chuanbo Hua, Junyoung Park, Laurin Luttman, Yining Ma, Fanchen Bu, Jiarui Wang, Haoran Ye, Minsu Kim, Sanghyeok Choi, Nayeli Gast Zepeda, André Hottung, Jianan Zhou, Jieyi Bi, Yu Hu, Fei Liu, Hyeonah Kim, Jiwoo Son, Haeyeon Kim, Davide Angioni, Wouter Kool, Zhiguang Cao, Jie Zhang, Kijung Shin, Cathy Wu, Sungsoo Ahn, Guojie Song, Changhyun Kwon, Lin Xie, and Jinkyoo Park. RL4CO: an extensive reinforcement learning for combinatorial optimization benchmark. *arXiv preprint arXiv:2306.17100*, 2023.
- [6] Xuan Wu, Di Wang, Lijie Wen, Yubin Xiao, Chunguo Wu, Yuesong Wu, Chaoyu Yu, Douglas L Maskell, and You Zhou. Neural combinatorial optimization algorithms for solving vehicle routing problems: A comprehensive survey with perspectives. *arXiv preprint arXiv:2406.00415*, 2024.
- [7] Daniela Thyssens, Tim Dernerde, Jonas K Falkner, and Lars Schmidt-Thieme. Routing arena: A benchmark suite for neural routing solvers. *arXiv preprint arXiv:2310.04140*, 2023.
- [8] Wouter Kool, Herke van Hoof, and Max Welling. Attention, learn to solve routing problems! In *International Conference on Learning Representations*, 2018.
- [9] Yeong-Dae Kwon, Jinho Choo, Byoungjip Kim, Iljoo Yoon, Youngjune Gwon, and Seungjai Min. POMO: Policy optimization with multiple optima for reinforcement learning. In *Advances in Neural Information Processing Systems*, volume 33, pages 21188–21198, 2020.
- [10] Yaixin Wu, Wen Song, Zhiguang Cao, Jie Zhang, and Andrew Lim. Learning improvement heuristics for solving routing problems. *IEEE Transactions on Neural Networks and Learning Systems*, 33(9):5057–5069, 2021.
- [11] Yining Ma, Zhiguang Cao, and Yeow Meng Chee. Learning to search feasible and infeasible regions of routing problems with flexible neural k-opt. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023.
- [12] Jieyi Bi, Yining Ma, Jianan Zhou, Wen Song, Zhiguang Cao, Yaixin Wu, and Jie Zhang. Learning to handle complex constraints for vehicle routing problems. *Advances in Neural Information Processing Systems*, 2024.
- [13] Jingxiao Chen, Ziqin Gong, Minghuan Liu, Jun Wang, Yong Yu, and Weinan Zhang. Looking ahead to avoid being late: Solving hard-constrained traveling salesman problem. *arXiv preprint arXiv:2403.05318*, 2024.
- [14] Xinyun Chen and Yuandong Tian. Learning to perform local rewriting for combinatorial optimization. In *Advances in Neural Information Processing Systems*, volume 32, pages 6281–6292, 2019.
- [15] Qiaoyue Tang, Yangzhe Kong, Lemeng Pan, and Choonmeng Lee. Learning to solve soft-constrained vehicle routing problems with lagrangian relaxation. *arXiv preprint arXiv:2207.09860*, 2022.
- [16] Ronald J Williams. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning*, 8(3):229–256, 1992.
- [17] Yining Ma, Jingwen Li, Zhiguang Cao, Wen Song, Hongliang Guo, Yuejiao Gong, and Yeow Meng Chee. Efficient neural neighborhood search for pickup and delivery problems. In *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence*, pages 4776–4784, 7 2022.
- [18] Jianan Zhou, Zhiguang Cao, Yaixin Wu, Wen Song, Yining Ma, Jie Zhang, and Xu Chi. MVMoE: Multi-task vehicle routing solver with mixture-of-experts. In *International Conference on Machine Learning*, 2024.

- [19] André Hottung, Yeong-Dae Kwon, and Kevin Tierney. Efficient active search for combinatorial optimization problems. In *International Conference on Learning Representations*, 2022.
- [20] Jinho Choo, Yeong-Dae Kwon, Jihoon Kim, Jeongwoo Jae, André Hottung, Kevin Tierney, and Youngjune Gwon. Simulation-guided beam search for neural combinatorial optimization. In *Advances in Neural Information Processing Systems*, volume 35, pages 8760–8772, 2022.
- [21] Oriol Vinyals, Meire Fortunato, and Navdeep Jaitly. Pointer networks. In *Advances in Neural Information Processing Systems*, volume 28, pages 2692–2700, 2015.
- [22] Irwan Bello, Hieu Pham, Quoc V. Le, Mohammad Norouzi, and Samy Bengio. Neural combinatorial optimization with reinforcement learning. In *International Conference on Learning Representations Workshop Track*, 2017.
- [23] Mohammadreza Nazari, Afshin Oroojlooy, Martin Takáč, and Lawrence V Snyder. Reinforcement learning for solving the vehicle routing problem. In *Advances in Neural Information Processing Systems*, pages 9861–9871, 2018.
- [24] Haoran Sun, Katayoon Goshvadi, Azade Nova, Dale Schuurmans, and Hanjun Dai. Revisiting sampling for combinatorial optimization. In *International Conference on Machine Learning*, pages 32859–32874. PMLR, 2023.
- [25] Minsu Kim, Junyoung Park, and Jinkyoo Park. Sym-NCO: Leveraging symmetricity for neural combinatorial optimization. In *Advances in Neural Information Processing Systems*, 2022.
- [26] 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*, 2023.
- [27] Felix Chalumeau, Shikha Surana, Clément Bonnet, Nathan Grinsztajn, Arnau Pretorius, Alexandre Laterre, and Thomas D Barrett. Combinatorial optimization with policy adaptation using latent space search. In *Advances in Neural Information Processing Systems*, 2023.
- [28] Nathan Grinsztajn, Daniel Furelos-Blanco, Shikha Surana, Clément Bonnet, and Thomas D Barrett. Winner takes it all: Training performant RL populations for combinatorial optimization. In *Advances in Neural Information Processing Systems*, 2023.
- [29] 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*, 2023.
- [30] André Hottung, Mridul Mahajan, and Kevin Tierney. PolyNet: Learning diverse solution strategies for neural combinatorial optimization. In *International Conference on Learning Representations*, 2025.
- [31] Fu Luo, Xi Lin, Zhenkun Wang, Tong Xialiang, Mingxuan Yuan, and Qingfu Zhang. Self-improved learning for scalable neural combinatorial optimization. *arXiv preprint arXiv:2403.19561*, 2024.
- [32] Daisuke Kikuta, Hiroki Ikeuchi, Kengo Tajiri, and Yuusuke Nakano. Routeexplainer: An explanation framework for vehicle routing problem. In *Pacific-Asia Conference on Knowledge Discovery and Data Mining*, pages 30–42. Springer, 2024.
- [33] Zefang Zong, Hansen Wang, Jingwei Wang, Meng Zheng, and Yong Li. 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*, pages 4648–4658, 2022.
- [34] 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 AAAI Conference on Artificial Intelligence*, volume 37, pages 8132–8140, 2023.
- [35] 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 *International Conference on Learning Representations*, 2023.
- [36] James Fitzpatrick, Deepak Ajwani, and Paula Carroll. A scalable learning approach for the capacitated vehicle routing problem. *Computers & Operations Research*, 171:106787, 2024.
- [37] Simon Geisler, Johanna Sommer, Jan Schuchardt, Aleksandar Bojchevski, and Stephan Günnemann. Generalization of neural combinatorial solvers through the lens of adversarial robustness. In *International Conference on Learning Representations*, 2022.

- [38] Pei Xiao, Zizhen Zhang, Jinbiao Chen, Jiahai Wang, and Zhenzhen Zhang. Neural combinatorial optimization for robust routing problem with uncertain travel times. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024.
- [39] Zeyang Zhang, Ziwei Zhang, Xin Wang, and Wenwu Zhu. Learning to solve travelling salesman problem with hardness-adaptive curriculum. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2022.
- [40] 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*, 2022.
- [41] Yuan Jiang, Yaoxin Wu, Zhiguang Cao, and Jie Zhang. Learning to solve routing problems via distributionally robust optimization. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2022.
- [42] Jianan Zhou, Yaoxin Wu, Wen Song, Zhiguang Cao, and Jie Zhang. Towards omni-generalizable neural methods for vehicle routing problems. In *International Conference on Machine Learning*, pages 42769–42789. PMLR, 2023.
- [43] Chengrui Gao, Haopu Shang, Ke Xue, Dong Li, and Chao Qian. Towards generalizable neural solvers for vehicle routing problems via ensemble with transferrable local policy. In *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence*, 2024.
- [44] Han Fang, Zhihao Song, Paul Weng, and Yutong Ban. Invit: A generalizable routing problem solver with invariant nested view transformer. In *Forty-first International Conference on Machine Learning*, 2024.
- [45] Han Lu, Zenan Li, Runzhong Wang, Qibing Ren, Xijun Li, Mingxuan Yuan, Jia Zeng, Xiaokang Yang, and Junchi Yan. ROCO: A general framework for evaluating robustness of combinatorial optimization solvers on graphs. In *International Conference on Learning Representations*, 2023.
- [46] Chenguang Wang and Tianshu Yu. Efficient training of multi-task combinatorial neural solver with multi-armed bandits. *arXiv preprint arXiv:2305.06361*, 2023.
- [47] Fei Liu, Xi Lin, Zhenkun Wang, Qingfu Zhang, Tong Xialiang, and Mingxuan Yuan. 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*, pages 1898–1908, 2024.
- [48] Federico Berto, Chuanbo Hua, Nayeli Gast Zepeda, André Hottung, Niels Wouda, Leon Lan, Junyoung Park, Kevin Tierney, and Jinkyoo Park. Routefinder: Towards foundation models for vehicle routing problems. *arXiv preprint arXiv:2406.15007*, 2024.
- [49] Zhuoyi Lin, Yaoxin Wu, Bangjian Zhou, Zhiguang Cao, Wen Song, Yingqian Zhang, and Senthilnath Jayavelu. Cross-problem learning for solving vehicle routing problems. *arXiv preprint arXiv:2404.11677*, 2024.
- [50] Chaitanya K Joshi, Thomas Laurent, and Xavier Bresson. An efficient graph convolutional network technique for the travelling salesman problem. *arXiv preprint arXiv:1906.01227*, 2019.
- [51] Benjamin Hudson, Qingbiao Li, Matthew Malencia, and Amanda Prorok. Graph neural network guided local search for the traveling salesperson problem. In *International Conference on Learning Representations*, 2022.
- [52] Ruizhong Qiu, Zhiqing Sun, and Yiming Yang. Dimes: A differentiable meta solver for combinatorial optimization problems. *Advances in Neural Information Processing Systems*, 35:25531–25546, 2022.
- [53] Zhiqing Sun and Yiming Yang. Difusco: Graph-based diffusion solvers for combinatorial optimization. In *Advances in Neural Information Processing Systems*, 2023.
- [54] Yimeng Min, Yiwei Bai, and Carla P Gomes. Unsupervised learning for solving the travelling salesman problem. *Advances in Neural Information Processing Systems*, 2023.
- [55] Kexiong Yu, Hang Zhao, Yuhang Huang, Renjiao Yi, Kai Xu, and Chenyang Zhu. Disco: Efficient diffusion solver for large-scale combinatorial optimization problems. *arXiv preprint arXiv:2406.19705*, 2024.
- [56] Paulo da Costa, Jason Rhuggenaath, Yingqian Zhang, and Alp Eren Akçay. Learning 2-opt heuristics for the traveling salesman problem via deep reinforcement learning. In *Asian Conference on Machine Learning*, pages 465–480, 2020.

- 385 [57] Yining Ma, Jingwen Li, Zhiguang Cao, Wen Song, Le Zhang, Zhenghua Chen, and Jing Tang. Learning
386 to iteratively solve routing problems with dual-aspect collaborative transformer. In *Advances in Neural*
387 *Information Processing Systems*, volume 34, pages 11096–11107, 2021.
- 388 [58] André Hottung and Kevin Tierney. Neural large neighborhood search for routing problems. *Artificial*
389 *Intelligence*, page 103786, 2022.
- 390 [59] Minjun Kim, Junyoung Park, and Jinkyoo Park. Learning to cross exchange to solve min-max vehicle
391 routing problems. In *The Eleventh International Conference on Learning Representations*, 2023.
- 392 [60] Minsu Kim, Jinkyoo Park, and joungho kim. Learning collaborative policies to solve np-hard routing
393 problems. In *Advances in Neural Information Processing Systems*, volume 34, pages 10418–10430, 2021.
- 394 [61] Haoran Ye, Jiarui Wang, Helan Liang, Zhiguang Cao, Yong Li, and Fanzhang Li. Glop: Learning global
395 partition and local construction for solving large-scale routing problems in real-time. In *Proceedings of the*
396 *AAAI Conference on Artificial Intelligence*, 2024.
- 397 [62] Zhi Zheng, Changliang Zhou, Tong Xialiang, Mingxuan Yuan, and Zhenkun Wang. Udc: A unified neural
398 divide-and-conquer framework for large-scale combinatorial optimization problems. *Advances in Neural*
399 *Information Processing Systems*, 2024.
- 400 [63] Detian Kong, Yining Ma, Zhiguang Cao, Tianshu Yu, and jianhua Xiao. Efficient neural collaborative
401 search for pickup and delivery problems. *IEEE Transactions on Pattern Analysis and Machine Intelligence*,
402 2024.
- 403 [64] Rodrigo Ferreira Da Silva and Sebastián Urrutia. A general vns heuristic for the traveling salesman problem
404 with time windows. *Discrete Optimization*, 7(4):203–211, 2010.
- 405 [65] Sirui Li, Zhongxia Yan, and Cathy Wu. Learning to delegate for large-scale vehicle routing. *Advances in*
406 *Neural Information Processing Systems*, 34:26198–26211, 2021.

Appendix

407	A Detailed literature review
408	A.1 Neural VRP solvers
409	A.2 Constraint handling for VRPs
410	A.3 Hybrid VRP solvers
411	B Problem selection and data generation
412	B.1 Traveling Salesman Problem with Time Window (TSPTW)
413	B.2 Capacitated VRP with Backhaul, Duration Limit, and Time Window Constraints (CVRPBLTW)
414	
415	C Limitations of individual paradigms for complex constraint handling
416	D Network architecture
417	D.1 Encoder
418	D.2 Decoder
419	E Experimental settings
420	F Additional experiments and discussion
421	F.1 Discussion of the existing feasibility masking
422	F.2 Discussion of the multi-start strategy and the diversification
423	F.3 Discussion of the unified model architecture
424	F.4 Additional results for the pattern of CaR’s refinement
425	F.5 Effects of joint training framework
426	F.6 Effects of \mathcal{L}_{DIV} and \mathcal{L}_{SL}
427	F.7 Efficiency under different refinement steps T_R
428	F.8 Sensitivity to random seeds and statistical significance

429 A Detailed literature review

430 A.1 Neural VRP solvers

431 Generally, neural VRP solvers can be broadly categorized into two paradigms: construction solvers
432 and improvement solvers.

433 1) *Construction-based solvers* learn to construct solutions from scratch in an end-to-end fashion.
434 Vinyals et al. [21] introduced the Pointer Network (PtrNet), leveraging Recurrent Neural Networks
435 (RNN) to solve the Traveling Salesman Problem (TSP) in a supervised manner. Building on this,
436 Bello et al. [22] explored reinforcement learning (RL) training for PtrNet. Nazari et al. [23] extended
437 the approach to solve CVRP in an autoregressive (AR) way. Among AR solvers, the Attention
438 Model (AM) [8] stands out as a milestone to solve multiple VRPs. This was further advanced by the
439 Policy Optimization with Multiple Optima (POMO) [9], which leverages diverse rollouts inspired
440 by the symmetry properties of VRP solutions. Subsequently, numerous studies have advanced AR
441 solvers in various perspectives, such as inference strategies [19, 20, 24], training paradigms [25–31],
442 interpretability [32], scalability [33–36], robustness [37, 38], benchmarking [7], and generalization
443 over different distributions [39–41], scales [42–44], and constraints [18, 45–49]. Beyond AR solvers,
444 another line of research predicts heatmaps in a non-autoregressive (NAR) manner to represent edge
445 probabilities for optimal solutions [50–55]. With the learned heatmap, these solvers can greatly
446 reduce the search space. Despite showing better scalability, NAR solvers often depend on post-search
447 procedures, which can be either time-consuming or ineffective in handling VRP constraints, even for
448 the simple cases such as CVRP.

449 2) *Improvement-based solvers* learn to iteratively improve initial solutions, drawing inspiration from
450 traditional (meta-)heuristics such as k-opt (e.g., 2-opt [10, 56, 57], extended to flexible k-opt [11]),

451 ruin-and-repair [17, 58], and crossover [59]. In general, improvement-based methods can achieve
 452 near-optimal solutions given prolonged search time, whereas construction-based methods typically
 453 offer a more efficient trade-off between performance and runtime.

454 A.2 Constraint handling for VRPs

455 Classic neural methods handle VRP constraints using feasibility masking to exclude invalid actions [9,
 456 57]. While effective for simpler VRPs (e.g., CVRP), such masking fails on more complex VRPs.
 457 Firstly, calculating masking itself may be NP-hard [12], e.g., for TSPTW and TSPDL. While recent
 458 methods have explored Lagrangian-based problem reformulation [15], lookahead strategies [13], and
 459 learnable approximated masks to prevent infeasibility [12], they may incur high computational costs
 460 and still yield high infeasibility rates. Secondly, current methods struggle with multi-constraint VRPs,
 461 e.g., CVRPBLTW, where construction faces overly restrictive masks and improvement drifts into
 462 infeasible regions (see Section 4). Moreover, no feasibility handling scheme effectively addresses
 463 both types of the above challenges. Hence, an efficient, general, and effective constraint-handling
 464 framework for neural VRP solvers remains absent.

465 A.3 Hybrid VRP solvers

466 Recent research has explored hybridizing construction-based approaches with search mechanisms to
 467 enhance combinatorial optimization solvers. For instance, active search [19] and beam search [20]
 468 have been used to augment construction policies. However, these methods are not well-suited for
 469 hard-constrained vehicle routing problems (VRPs), where feasibility is critical. Other strategies, such
 470 as collaborative policies [60], random reconstruction [29], and neural divide-and-conquer frameworks
 471 (e.g., GLOP [61] and UDC [62]), have shown strong performance on CVRP but often struggle
 472 to maintain feasibility when solutions are decomposed or reconstructed, even with penalty-based
 473 guidance.

474 Beyond post-construction search, hybrid methods that integrate construction with intensive improve-
 475 ment phases have been proposed. RL4CO [5] combines pretrained POMO [9] and NeuOpt [11]
 476 at inference time, while NCS [63] employs a shared critic to jointly train the two modules. These
 477 approaches, however, depend on separately optimized components and often require extensive
 478 improvement steps (e.g., 5,000 iterations), limiting their efficiency and scalability.

479 Our approach distinguishes itself in three critical ways: 1) it adopts a joint training framework
 480 with unified representation learning, unlike prior works that treat construction and improvement as
 481 separate [5] or loosely coupled [63] components; 2) it achieves efficient and precise refinement in as
 482 few as 10 steps, rather than relying on prolonged improvement; and 3) it is the first to systematically
 483 explore the synergy between construction and improvement for effective constraint handling in
 484 complex VRPs.

485 Among the most relevant hybrid baselines, LCP [60], RL4CO [5], and NCS [63] are not directly
 486 applicable to complex constrained VRPs due to significant limitations: LCP’s divide-and-conquer
 487 strategy impairs constraint handling; RL4CO lacks joint optimization; and NCS fails to generate fea-
 488 sible moves or predict value functions accurately under hard constraints, leading to 100% infeasibility
 489 on TSPTW in our preliminary evaluations.

490 B Problem selection and data generation

491 As a single-task solver focused on constraint handling, *our work aligns with prior single-task studies*
 492 *that typically evaluate 2–3 representative VRPs*. Specifically, we consider three representative
 493 VRPs: complex constrained VRPs where feasibility masking is NP-hard thus intractable (TSPTW) or
 494 tractable but ineffective (CVRPBLTW), and simpler constraints (CVRP) where feasibility masking
 495 is tractable and effective. Below we introduce the detailed data generation process. Following the
 496 convention, all the node coordinates are generated under a uniform distribution, i.e., $x_i, y_i \sim \mathcal{U}[0, 1]$.

497 B.1 Traveling Salesman Problem with Time Window (TSPTW)

498 We primarily follow the settings of a recent study [12], which introduced TSPTW with three difficulty
 499 levels: Easy [13], Medium, and Hard. Given our focus on complex constrained VRPs, we report

results on the *Hard* variant in the main tables, consistent with the benchmark dataset in [64]. Except for the experiment in Table 2, where we use the Medium variant to show how CaR’s unified encoder works as constraint hardness changes, all other TSPTW experiments default to the Hard setting.

NP-hardness of computing feasibility masking on TSPTW. During solution construction, neural solvers typically apply feasibility masking to exclude actions that violate constraints. In TSPTW, for instance, a node v_j is masked out if its arrival time t_j exceeds its time window end u_j , i.e., $t_j > u_j$. However, as highlighted in PIP [12], feasibility in TSPTW is not purely local: selecting a node affects the current time, which in turn influences the feasibility of all future selections due to interdependent time window constraints. For example, a node with a late time window might be locally feasible, but choosing it can delay the tour such that earlier nodes become unreachable, leading to irreversible infeasibility. Ensuring global feasibility thus requires evaluating whether a current decision allows for any feasible completions, i.e., a process that involves simulating all future possibilities. This renders feasibility masking itself NP-hard, compounding the inherent difficulty of solving TSPTW.

TSPTW Hard. We first generate a random permutation τ of the node set and sequentially assign time windows to ensure the existence of a feasible solution for each instance. The time windows are drawn from a uniform distribution, where the lower and upper bounds are given by $l_i \sim \mathcal{U}[C_L(\tau'_i) - \eta, C_L(\tau'_i)]$, $u_i \sim \mathcal{U}[C_L(\tau'_i), C_L(\tau'_i) + \eta]$, where $C_L(\tau'_i)$ denotes the cumulative tour length of the partial sequence τ' up to step i , and η controls the time window width. To adaptively scale time windows with problem size, we set $\eta = n$, offering greater flexibility than the fixed value of 50 used in [12]. Time windows are then normalized following [8] to facilitate neural network training. Since TSP solutions form a Hamiltonian cycle, it is equivalent to setting any node as the starting point. Thus, we designate a specific starting node for TSPTW by redefining its upper bound u_0 as $u_0 = \max(u_i + C_L(e(v_i, v_0))), i \in [1, n]$.

TSPTW Medium. We generate the lower and upper bounds of the time windows following a uniform distribution: $l_i \sim \mathcal{U}[0, T_N]$, where T_N estimates the expected tour length for the given problem scale (e.g., $T_{20} \approx 10.9$ [13]). The upper bound u_i is derived from l_i as $u_i \sim l_i + T_N \cdot \mathcal{U}[0.1, 0.2]$. While this data generation rule does not guarantee instance feasibility, preliminary results show that the time windows overlap significantly, leading to a high feasibility rate in the generated instances.

For all the TSPTW instances, we normalize all l_i and u_i by u_0 to ensure their values fall within $[0, 1]$. Training data is generated on the fly, while inference uses the dataset from [12] for fair comparison. As noted in [12], all test instances are verified to be feasible, either empirically or theoretically.

B.2 Capacitated VRP with Backhaul, Duration Limit, and Time Window Constraints (CVRPBLTW)

CVRPBLTW follows the setting in [18, 47, 65] and includes four key constraints: 1) *Capacity (C)*. Each node’s demand q_i is sampled from $\mathcal{U}(1, \dots, 9)$, and the vehicle’s capacity Q varies by problem scale, with $Q^{50} = 40$ and $Q^{100} = 50$. 2) *Backhaul (B)*. Node demands are sampled from a discrete uniform distribution $\{1, \dots, 9\}$, with 20% of customers randomly designated as backhauls. Routes include both linehauls and backhauls without strict precedence constraints. 3) *Duration Limit (L)*. The maximum route length is set to $\ell = 3$, ensuring feasible solutions in the unit square space $\mathcal{U}(0, 1)$. 4) *Time Window (TW)*. The depot time window is defined as $[l_0, u_0] = [0, 3]$, and each customer node has a service time of $s_i = 0.2$. The time window for node v_i is computed as follows: the center is sampled as $\gamma_i \sim \mathcal{U}(l_0 + C_L(e(v_0, v_i)), u_0 - C_L(e(v_i, v_0)) - s_i)$, where $C_L(e(v_0, v_i)) = C_L(e(v_i, v_0))$ represents the travel time between the depot v_0 and node v_i . The half-width is drawn from $w_i \sim \mathcal{U}(\frac{s_i}{2}, \frac{u_0}{3}) = \mathcal{U}(0.1, 1)$. Finally, the time window is set as $[l_i, u_i] = [\max(l_0, \gamma_i - w_i), \min(u_0, \gamma_i + w_i)]$. Training data is generated on the fly, while inference uses the 1k-instance dataset from [18] for fair comparison.

C Limitations of individual paradigms for complex constraint handling

Figure 3 depicts the illustrative search behaviors of neural solvers on two hard-constrained VRPs: i) CVRPBLTW (a–c) where feasibility masking for construction is tractable but overly restrictive, leading to a fragmented feasible space due to multiple constraints; and ii) TSPTW (d–f), where computing feasibility masks is NP-hard and thus intractable [12], yielding a limited feasible region.

Both construction and improvement solvers may struggle with these hard constraints. In CVRPBLTW, overly restrictive masking limits the exploration of construction in the fragmented feasible space, causing early convergence to local optima (Figure 3(a)). This aligns with our findings in Section 4: removing hard masking for CVRPBLTW improves performance (POMO* vs. POMO in Table 1) but still yields infeasible and suboptimal solutions (Appendix F.1). For TSPTW, the relaxed CMDP in Eq.(2) offers partial relief, but Bi et al. [12] show it fails in complex cases and propose a learnable mask to guide the policy toward near-feasible regions (Figure 3(d)). Still, the performance of construction methods alone remains limited due to the rigidity of inherent stepwise feasibility. For improvement methods with advanced operators (e.g., k -opt), masking is inapplicable and thus unavailable in both cases, often resulting in inefficient search trajectories through both feasible and infeasible regions. As shown in Figure 3(b) and Figure 3(e), while they can refine solutions, they are computationally inefficient, often requiring substantial time to find good initial solutions, and performance degrades significantly with multiple constraints and limited search steps (Table 1).

To address these limitations, we propose CaR, which unifies the strengths of both paradigms by using construction to generate diverse, high-quality, near-feasible initial solutions, thereby reducing the required refinement iterations. As shown in Figures 3(c) and (f), CaR applies lightweight refinement with only a few steps to help construction escape local optima and reach higher-quality local optima, enabling more effective and efficient constraint handling for complex VRPs.

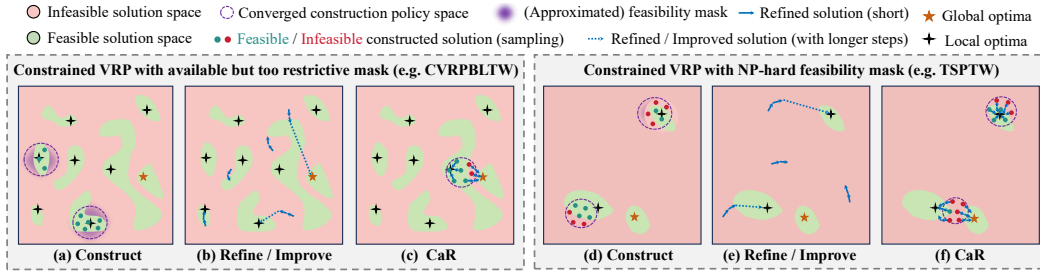


Figure 3: Illustration of the solution space and paradigm behaviors for two hard-constrained VRPs.

D Network architecture

The network architectures of mainstream neural construction and improvement solvers are typically based on a Transformer encoder with an attention-based decoder, enabling unification across both paradigms. In this paper, we adopt a unified encoder shared by the construction and refinement decoders. Specifically, we use a 6-layer Transformer encoder, following POMO [9], while retaining the original decoders from NeuOpt [11] and N2S [17], corresponding to two representative local search operators: flexible k -opt [11] and remove-and-reinsertion (R&R) [17], respectively. For the unexplored variants TSPTW and CVRPBLTW, we design new constraint-related features analogous to the contextual features used in the original decoders. The concrete forward processes are introduced below.

D.1 Encoder

As shown in Figure 4, the construction module first takes node features f_i^n —including coordinates (x_i, y_i) and constraint-related features (e.g., time windows $[l_i, u_i]$ and demand q_i)—as input. For each node v_i , these features are projected into a d -dimensional embedding $\mathbf{h}_i^{(0)} \in \mathbb{R}^d$ ($d = 128$) via a linear layer. The initial embedding is then passed through a 6-layer Transformer network [9]. At each layer j ($j = 1, \dots, 6$), the embedding $\mathbf{h}_i^{(j-1)}$ is projected into query, key, and value vectors:

$$\mathbf{q}_i^{(j)} = W_q^{(j)} \mathbf{h}_i^{(j-1)}, \quad \mathbf{k}_i^{(j)} = W_k^{(j)} \mathbf{h}_i^{(j-1)}, \quad \mathbf{v}_i^{(j)} = W_v^{(j)} \mathbf{h}_i^{(j-1)}, \quad (6)$$

where $W_q^{(j)}, W_k^{(j)}, W_v^{(j)} \in \mathbb{R}^{d \times d}$. These are fed into a multi-head attention (MHA) layer, whose output is:

$$\tilde{\mathbf{h}}_i^{(j)} = \text{Softmax} \left(\frac{(\mathbf{q}_i^{(j)})^\top \mathbf{k}_i^{(j)}}{\sqrt{d}} \right) \mathbf{v}_i^{(j)}. \quad (7)$$

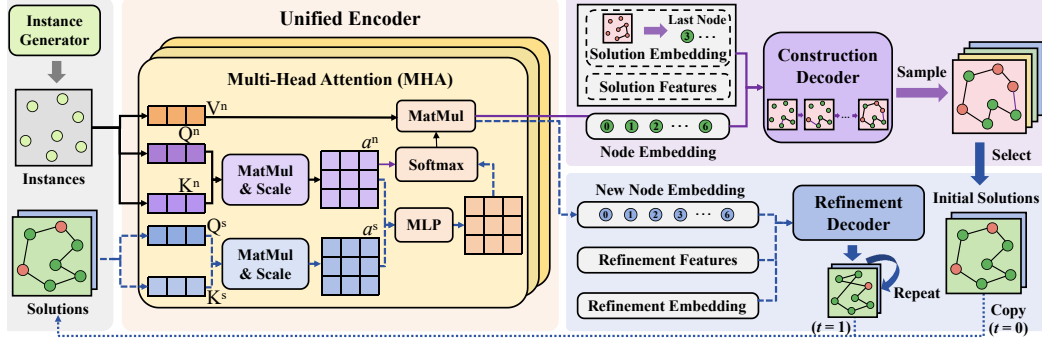


Figure 4: Overview of the unified network architecture in CaR. Blue dashed arrows indicate information flow specific to refinement, while purple dashed arrows indicate flow exclusive to construction.

587 The MHA output is then linearly transformed:

$$\hat{h}_i^{(j)} = W_{\text{MHA}}^{(j)} \tilde{h}_i^{(j)}, \quad (8)$$

588 where $W_{\text{MHA}}^{(j)} \in \mathbb{R}^{d \times d}$. This passes through instance normalization with residual connection:

$$h_i'^{(j)} = \text{IN} \left(\hat{h}_i^{(j)} + h_i^{(j-1)} \right). \quad (9)$$

589 Next, a feed-forward (FF) layer refines the output:

$$h_i''^{(j)} = W_2^{(j)} \cdot \text{ReLU} \left(W_1^{(j)} h_i'^{(j)} \right), \quad (10)$$

590 where $W_1^{(j)} \in \mathbb{R}^{d \times d'}$, $W_2^{(j)} \in \mathbb{R}^{d' \times d}$, and $d' = 512$ as in [9]. The final embedding at layer j is:

$$h_i^{(j)} = \text{IN} \left(h_i''^{(j)} + h_i'^{(j)} \right). \quad (11)$$

591 Thereafter, this process is repeated for 6 layers, and the final encoder output is $h_i = h_i^{(6)}$.

592 When it comes to the refinement module, the input changes from a node feature set to a linked list (i.e.,
 593 the solution), which consists of the original node features plus positional information and a pointer
 594 to the next node in the solution. Notably, for CVRP variants, the depot node may appear multiple
 595 times, with each occurrence treated as a distinct node due to its different position. The refinement
 596 module shares the same operations as the construction module, except for Eq. (7). Specifically, we
 597 first incorporate the cyclic positional encoding (CPE) from [57] to obtain the positional embedding
 598 $p_i \in \mathbb{R}^d$. We then apply a self-attention mechanism:

$$\tilde{p}_i = (W_q^p p_i)^\top (W_k^p p_i), \quad (12)$$

599 where $W_q^p, W_k^p \in \mathbb{R}^{d \times d}$, to compute the positional attention matrix a^s . We replace Eq. (7) from the
 600 construction module with:

$$\tilde{h}_i^{(j)} = \text{Softmax} \left(\frac{\text{MLP} \left(\left[(q_i^{(j)})^\top k_i^{(j)}, \tilde{p}_i \right] \right)}{\sqrt{d}} \right) v_i^{(j)}, \quad (13)$$

601 where the MLP reduces the dimension of the concatenated attention scores, i.e., $[a^n, a^s]$ in Figure 4,
 602 from 2 to 1, following the Syn-Att mechanism proposed in [17].

603 D.2 Decoder

604 The construction decoder primarily employs an MHA layer, where the key and value vectors are
 605 linearly transformed from the node embedding h_i . The query vector is computed using contextual
 606 information from the construction process, including: 1) the embedding of the partial solution h^s (i.e.,
 607 the node embedding of the last node in the partial solution), and 2) the step-wise solution features

608 f^s , which include the vehicle’s remaining load, current time, and current sub-tour length. The MHA
 609 output is calculated as:

$$a = \sum_{i=0}^n \text{Softmax} \left(\frac{q_i^\top k_i}{\sqrt{d}} + \xi_i \right) v_i, \quad (14)$$

610 where ξ_i is the feasibility mask for node v_i . We mask out visited nodes to ensure solution validity
 611 across all variants. For CVRP, as discussed in Appendix F.1, we also mask nodes that violate capacity
 612 constraints at each construction step. The output of each head is then passed through a linear layer
 613 parameterized by W_o , yielding $h_a = W_o a$. Finally, the decoder computes selection probabilities for
 614 all candidate nodes using a single-head attention layer:

$$p_i = \text{Softmax} \left(\zeta \cdot \tanh \left(\frac{h_a^\top h_i}{\sqrt{d}} \right) + \xi_i \right), \quad (15)$$

615 where ζ scales the logits to encourage exploration [9].

616 *The refinement decoder* follows the original design in the NeuOpt [11] and N2S [17], which performs
 617 two representative local search operators: the flexible k -opt¹ and remove-and-reinsertion (R&R),
 618 respectively. Figure 5 illustrates the detailed improvement process in a single time step for different
 619 local search operators, where each node selection corresponds to a full computation of the network
 620 decoder. Note that although the original NeuOpt and N2S decoders adopt different architectures, their
 621 input and output formats are similar. At each refinement step $t = 1, \dots, T_R$, the inputs include: (i)
 622 the updated *node embedding* h_i^t derived from the previous solution τ_{t-1} (with τ_0 as the top- p initial
 623 construction solutions); (ii) *refinement features* encoding feasibility transition history for k -opt or
 624 removal action history for R&R; and (iii) *refinement embedding* of the last action node, analogous to
 625 the solution embedding in construction. The output is the selection probability over candidate nodes
 626 for the next refinement operation (e.g., k -opt, removal, or insertion).

627 To adapt improvement solvers to new variants (e.g., TSPTW, CVRPBLTW), we follow their original
 628 design principles while integrating variant-specific features, such as refinement history and node-level
 629 feasibility information. Specifically, when using the NeuOpt decoder, we encode the feasibility of the
 630 most recent three refinement steps as a binary vector to represent refinement history and incorporate
 631 node-level feasibility features to enhance constraint awareness. For TSPTW, these features include
 632 the arrival time at each node, time window violation value, last node arrival time, and an indicator of
 633 whether the solution becomes infeasible after visiting the current node. For CVRPBLTW, node-level
 634 feasibility features include binary indicators for depot and backhaul nodes, violation values for each
 635 constraint, constraint-related attributes (e.g., arrival time and cumulative distance/demand), and
 636 infeasibility markers indicating whether the solution becomes infeasible before or after visiting the
 637 current node. Adapted NeuOpt variants are marked with ‡ in the main tables.

638 E Experimental settings

639 We preset refinement steps $T_R = 5$ during training. Models are trained using Adam with a learning
 640 rate of 1×10^{-4} , weight decay of 1×10^{-6} . Follow the setting of [9], we use instance normalization
 641 in the MHA and set the embedding dimension to 128. To prevent out-of-memory issues early in
 642 training, refinement is activated after a ~ 10 -epoch warmup of the construction module. To optimize
 643 GPU memory utilization, we adjust the number of refined solutions p separately for each problem
 644 size during training. We set $\alpha_1 = 0.01$ and $\alpha_2 = 1$ in the construction loss, and set $\omega = 100$ for
 645 TSPTW and CVRP, and 10 for CVRPBLTW based on the relative loss scales between modules
 646 observed during warm-up epochs. During inference, TSPTW and CVRPBLTW results are obtained
 647 without multi-starting, consistent with the training setup. For each augmented instance, we sample
 648 one solution for TSPTW and greedily generate one for CVRPBLTW ($p = 1$), which are then passed
 649 to the refinement module. For CVRP, Table 3 shows that multi-starting improves performance, so we
 650 apply it during construction but refine only the top p solutions for efficiency, setting $p = 2$ for $n = 50$
 651 and $p = 1$ for $n = 100$ due to the GPU memory constraint. We use only one GPU for all problem
 652 sizes during inference to ensure fair comparison across variants. All experiments are conducted on
 653 servers equipped with NVIDIA GeForce RTX 4090 GPUs and Intel(R) Core i9-10940X CPUs at
 654 3.30GHz.

¹Following [11], we set $k \leq 4$.

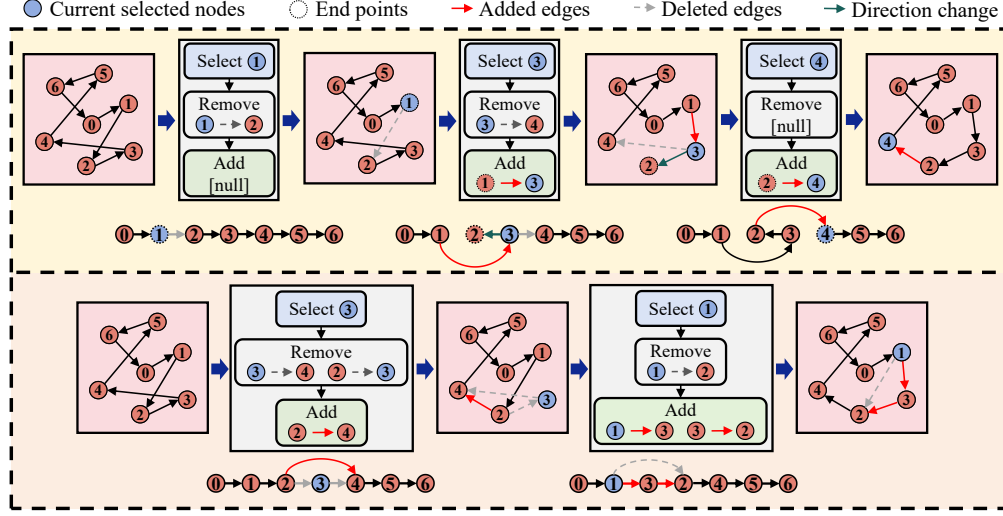


Figure 5: Illustration of actions and operations in k -opt (top) and remove-and-reinsertion (bottom).

F Additional experiments and discussion

F.1 Discussion of the existing feasibility masking

We now present our insights into the limitations of existing feasibility masking mechanisms. As shown in Table 3, removing feasibility masking for CVRP hampers its performance, in contrast to more complex problems like CVRPBLTW, where its removal leads to improved results (see results of POMO* vs. POMO in Table 1). This suggests that *while existing feasibility masking is sufficient for simple problems, it is far from a silver bullet for handling complex constraints*. In fact, for highly constrained problems, feasibility masking is often intractable (e.g., TSPTW) or ineffective (e.g., CVRPBLTW), resulting in infeasibility and sub-optimality.

Table 4: CaR module impact on CVRPBLTW feasibility and quality (**Bold**: reported in Table 1).

Method	Module	n=50			n=100		
		Obj.↓	Gap↓	Infsb%↓	Obj.↓	Gap↓	Infsb%↓
CaR ($T_R = 20$)	Construction (w/o mask)	14.798	2.116%	3.50%	24.404	-2.050%	2.80%
	Refinement (k -opt)	14.759	1.682%	2.50%	24.360	-2.285%	2.60%
	Re-construction (w. mask)	14.844	2.114%	0.00%	24.585	-1.724%	0.00%
CaR (R&R) ($T_R = 20$)	Construction (w/o mask)	14.843	2.235%	2.60%	24.460	-1.715%	3.10%
	Refinement (R&R)	14.567	0.221%	1.10%	24.215	-2.886%	1.90%
	Re-construction (w. mask)	14.601	0.463%	0.00%	24.400	-2.448%	0.00%

To address these challenges, we propose the CaR framework. For TSPTW, CaR leverages refinement to correct infeasibility and improve solution quality. For CVRPBLTW, we found that removing this

Table 3: Effects of the multi-starting strategy in [9] and feasibility masking during construction on CVRP-50.

w. multi-start	w. mask	Obj.↓	Gap↓	Infsb%↓
×	×	10.550	2.093%	0.00%
×	✓	10.520	1.793%	0.00%
✓	×	10.431	0.921%	0.00%
✓	✓	10.424	0.857%	0.000%




Figure 6: Masked node count in CVRPBLTW construction.

Table 5: Effects of unified encoder on TSPTW.

Method	TSPTW-50 Medium		TSPTW-50 Hard		TSPTW-100 Medium	
	Infsb%↓	Gap↓	Infsb%↓	Gap↓	Infsb%↓	Gap↓
Construction-only	3.77%	5.230%	37.270%	1.635%	0.120%	10.931%
Separate	0.010%	2.047%	0.675%	0.199%	0.000%	7.589%
Unified (Ours)	0.000%	2.173%	0.010%	0.014%	0.000%	5.815%

Table 6: Model parameters and performance of different unifications on TSPTW-50 Hard.

Encoder	Decoder	#Params	Infsb%↓	Gap↓
Separate	Separate	2.8M	0.675%	0.199%
Unified	Unified	1.4M	0.010%	0.041%
Unified	Separate	1.6M	0.010%	0.014%




Figure 7: Effects of separate decoder.

restrictive mask significantly improves the performance (Gap: 9.17% vs. 2.31%), but it leads to unexpected infeasibility (2.6%). CaR addresses this by applying refinement, which further reduces infeasibility and improves solution quality. However, a small fraction of infeasible solutions can still remain. Since feasibility can ultimately be guaranteed through masking, we apply a final reconstruction step with feasibility masking during inference to ensure all reported solutions are valid (see Table 4). This strategy allows CaR to relax constraints during construction and refinement to enhance quality while maintaining overall feasibility through targeted post-processing. Note that in this paper, “without masks” does not imply removing all masks. All VRP variants must satisfy the fundamental constraint that each node (except depots) is visited exactly once. To ensure the validity of the generated solutions, we retain the masking of visited nodes during the construction process.

F.2 Discussion of the multi-start strategy and the diversification

Following the insights from [30] and [12], the multi-start strategy proposed in POMO [9] negatively impacts performance when training on VRP variants with TW constraints. Therefore, in this paper, we sample S start nodes instead of enumerating all nodes for constrained variants such as TSPTW and CVRPBLTW. Motivated by the advantage of diversification by multi-start, we add the diversity loss in Eq. (4), which is empirically proven to be effective (see Figure 10).

F.3 Discussion of the unified model architecture

Recall that CaR employs a unified encoder shared across the construction and refinement decoders. As demonstrated in Table 2, this architecture significantly enhances performance. To further assess the impact of the unified encoder under varying constraint complexities, we present comprehensive results in Table 5 for *Medium* TSPTW-50, *Hard* TSPTW-50, and *Medium* TSPTW-100. Results indicate that as problem complexity increases, evidenced by higher infeasibility rates (e.g., 37.270% in *Hard* TSPTW-50) and greater optimality gaps (e.g., 10.931% in TSPTW-100), the performance improvements attributed to the unified encoder become more pronounced. This suggests that shared representations facilitate better generalization and constraint handling in more challenging scenarios.

We further evaluate this design on CVRPBLTW-100, where using separate encoders reduces the optimality gap slightly from -2.448% to -2.544% but increases model parameters by 1.7× compared to the unified encoder version of CaR (2.9M vs. 1.7M). To balance performance and computational efficiency, we advocate for the unified encoder design.

Moreover, we explore the unification of the encoder and decoder components. As shown in Table 6 and Figure 7, results on TSPTW-50 indicate that employing a separate decoder increases the number of model parameters by 1.1×, yet it further reduces the optimality gap. Therefore, this paper adopts an architecture with a unified encoder and separate decoders for each module.

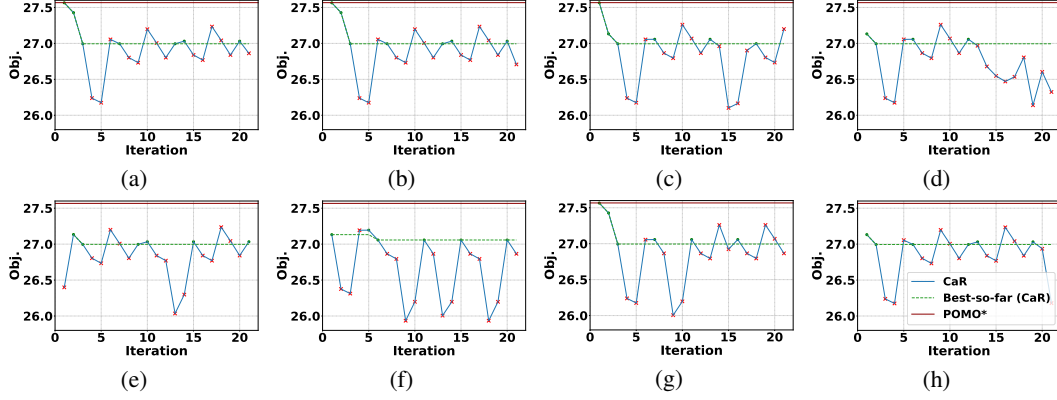


Figure 8: CaR’s refinement on one TSPTW instance. Green dots and red crosses denote the feasible and infeasible solutions, respectively. Green dashed lines mark the best-so-far feasible objective; the red line indicates the best objective with data augmentation. CaR’s refinement process is shown over time, with $t = 0$ representing the initially constructed solution.

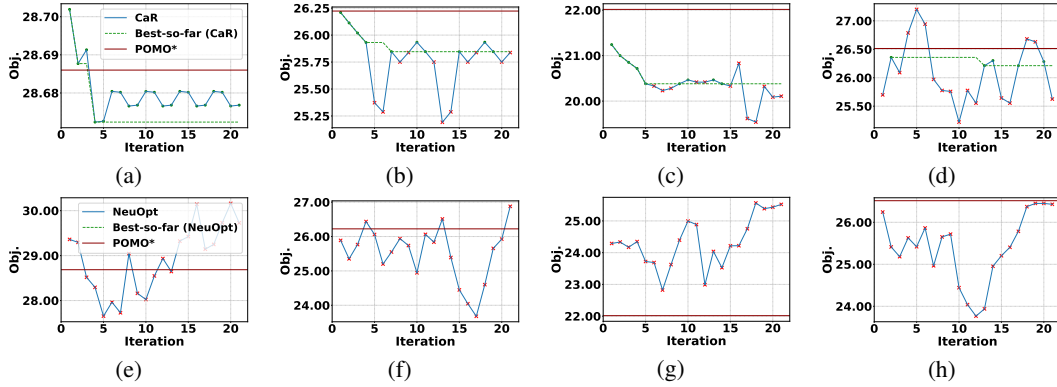


Figure 9: Refinement efficiency of CaR (*top*) and NeuOpt (*bottom*) on typical TSPTW-50 instances.

699 F.4 Additional results for the pattern of CaR’s refinement

700 During inference, we follow [9] to augment each instance $8\times$, generating eight initial solutions
 701 for refinement. Figure 8 shows four representative refinement trajectories on a TSPTW instance to
 702 illustrate the diversity in initial solutions and their subsequent refinement.

703 Moreover, to better understand how CaR refines solutions for feasibility and optimality, we also
 704 visualize the objective trajectories on other randomly generated TSPTW-50 instances in Figure 9.
 705 CaR dynamically navigates both feasible and infeasible regions, adjusting its focus based on the
 706 current solution state. For example, in the instance of Figure 9 (d), CaR begins with a near-feasible
 707 solution ($t = 0$), achieves feasibility in one step ($t = 1$), explores infeasible regions ($t = 2-12$), and
 708 ultimately improves optimality ($t = 13$). In contrast, in the instance of Figure 9 (c), the construction
 709 already yields a feasible, high-quality solution, and refinement continues to enhance optimality
 710 efficiently. These patterns illustrate CaR’s ability to balance feasibility and optimality by exploring
 711 both feasible and infeasible spaces. By comparison, NeuOpt fails to escape infeasible regions within
 712 limited steps, highlighting CaR’s advantage in constraint-aware refinement.

713 F.5 Effects of the joint training framework.

714 To assess our joint training framework, we evaluate simple combinations without our joint training
 715 framework. As shown in Table 7, joint training is crucial, showing that simply combining the two
 716 paradigms is insufficient.

Table 7: Effects of joint training on TSPTW-50. Inference time is fixed.

ID	Construction	Refinement	Joint train	Infsb%↓	Gap↓
0	Random	Trainable NeuOpt*	×	100.00%	/
1	Pretrained PIP	LKH3	×	0.15%	0.003%
2	Pretrained PIP	Pretrained NeuOpt*	×	0.03%	0.034%
3	Pretrained PIP	Trainable NeuOpt*	×	2.44%	0.161%
4 (CaR)	Trainable PIP	Trainable NeuOpt*	✓	0.00%	0.005%
5	Pretrained POMO*	LKH3	×	82.81%	0.024%
6	Pretrained POMO*	Pretrained NeuOpt*	×	0.22%	0.083%
7	Pretrained POMO*	Trainable NeuOpt*	×	33.89%	1.542%
8 (CaR)	Trainable POMO*	Trainable NeuOpt*	✓	0.01%	0.014%




Figure 10: \mathcal{L}_{DIV} and \mathcal{L}_{SL} .

Table 8: Standard deviation of the results in Table 1 under five different random seeds.

Problem	Method	n=50		n=100	
		Gap ↓	Infsb% ↓	Gap ↓	Infsb% ↓
TSPTW	CaR-POMO ($T_R = 20$)	0.014% ± 0.000%	0.00% ± 0.01%	0.404% ± 0.001%	2.30% ± 0.03%
	CaR-PIP ($T_R = 20$)	0.005% ± 0.000%	0.00% ± 0.01%	0.144% ± 0.001%	0.02% ± 0.01%
CVRPBLTW	CaR (k -opt) ($T_R = 20$)	2.147% ± 0.029%	0.00% ± 0.00%	-1.730% ± 0.004%	0.00% ± 0.00%
	CaR (R&R) ($T_R = 20$)	0.465% ± 0.007%	0.00% ± 0.00%	-2.449% ± 0.010%	0.00% ± 0.00%

F.6 Effects of \mathcal{L}_{DIV} and \mathcal{L}_{SL} .

We further ablate the diversity loss and the supervised loss for the construction module in Eq. (4) and Eq. (5), respectively. Figure 10 shows that diverse initial solutions aid refinement, and supervision with improved solutions strengthens cross-paradigm collaboration. Additional CVRP results in Appendix F.2 show that \mathcal{L}_{DIV} remains effective with multi-starting.

F.7 Efficiency under different T_R

As shown in Table 1, increasing the number of refinement steps (T_R) improves CaR performance, particularly for larger problem instances. For example, the optimality gaps on TSPTW-50 and TSPTW-100 decrease from 0.034% and 1.065% at $T_R = 5$ to 0.014% and 0.406% at $T_R = 20$, respectively. This supports the intuition that longer solution sequences benefit from more refinement. Notably, CaR achieves competitive solution quality and feasibility with only 5 refinement steps, significantly reducing computational effort compared to traditional methods requiring up to 5,000 steps. Further, we illustrate the evolution of the solution feasibility and optimality on hard-constrained problems (TSPTW-50 and CVRPBLTW-50). Results indicate that more constrained instances, such as CVRPBLTW, require additional refinement steps for convergence. This suggests that dynamically adjusting T_R based on constraint complexity could further enhance CaR.

F.8 Sensitivity to random seeds and statistical significance

Sensitivity to random seeds. To assess the robustness of CaR, we evaluate its performance under five different random seeds, i.e., 2023, 2024, 2025, 2026 and 2027, during inference. As shown in Table 8, the standard deviations are small, indicating that CaR’s performance is stable with respect to random initialization.

Statistical significance. Figure 11 illustrates the distribution of optimality gaps for NeuOpt, CaR-POMO, and CaR-PIP, corresponding to the results on TSPTW-50 in Table 1. We additionally conduct significance tests, revealing that all pairwise differences are statistically significant ($p < 0.001$). This suggest that CaR delivers a significant performance improvement over the representative baselines.




Figure 11: Boxplot of diverse methods on TSPTW-50.