DCAST: DECOUPLED CONSTRAINT ACTIVATION SIGNATURE TOKENIZATION FOR MULTI-TASK NEURAL MULTI-OBJECTIVE VEHICLE ROUTING PROBLEMS

Anonymous authors
Paper under double-blind review

ABSTRACT

Neural multi-objective combinatorial optimization (NMOCO) has attracted increasing attentions in solving multi-objective vehicle routing problems (MOVRPs) by their fast Pareto-front approximation capability. However, existing NMOCO solvers are typically trained for a single MOVRP variant, which severely restricts their applications in real-world industry problems with diverse characteristics. Therefore, this paper concentrates on developing multi-task NMOCO solvers for MOVRPs. First, we reformulate 16 single-objective VRP variants to bi-objective VRPs. Subsequently, we adapt two state-of-the-art single-task NMOCO solvers, WE-CA and POCCO, to the multi-task setting by extending their node representations and incorporating dynamic routing attributes into their context embeddings, yielding W-T and P-T, respectively. We also adapt MV-MoE, a representative multi-task single-objective NCO solver, to the multi-task NMOCO settings by embedding the preference vector as a pseudo-node and using it to modulate node representations via FiLM, yielding MOM, and further strengthening of its transferred model with a decoder-side nonlinear expert layer yields the MOM-M. Furthermore, building on MOM-M, we propose the Decoupled Constraint Activation Signature Tokenization (DCAST), which factorizes a five-dimensional binary constraint activation signature into five learnable constraint identity embeddings and two shared learnable activation-state embeddings to construct state-aware constraint tokens that are jointly encoded with node and preference embeddings. Lastly, we extend the above solvers with a size-agnostic training paradigm to further boost their cross-size generalization. Extensive experiments on the reformulated MOVRPs demonstrate the superiority of the multi-task paradigms over their original single-task counterparts in terms of both cross-problem and cross-size versatility, as well as the superiority of the proposed DCAST among multi-task solvers. The source code is available at: `https://github.com/DFSDarren/DCAST`.

1 INTRODUCTION

Vehicle routing problem (VRP) plays a critical role in modern logistics Jozefowiez et al. (2008), transportation Ahmadi et al. (2016); Zhang et al. (2023), manufacturing Ghoseiri et al. (2004), and intelligent scheduling systems Xiao et al. (2024); Ostovari et al. (2025). It is one of the most fundamental and ubiquitous combinatorial optimization problems. In real-world scenarios, routing decisions are rarely governed by a single metric. Instead, decision-makers must simultaneously balance multiple conflicting objectives, such as travel cost, delivery urgency, customer satisfaction, and workload equity. This complexity has led to the emergence of multi-objective vehicle routing problems (MOVRPs). Furthermore, practical industrial applications often involve a diverse array of operational constraints, such as time windows and pickup-and-delivery requirements Braekers et al. (2016); Vidal et al. (2020), resulting in numerous specialized MOVRP variants. However, developing dedicated algorithms for each specific variant is not only computationally expensive but also practically unmaintainable, highlighting the urgent need for a unified framework capable of solving diverse MOVRPs efficiently.

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

Unlike traditional single-objective VRPs that seek one optimal solution Berto et al. (2025a); Kwon et al. (2020b); Li et al. (2023); Wu et al. (2021); Kim et al. (2021); Zhou et al. (2023); Chen & Tian (2019); Hottung et al. (2022); Nazari et al. (2018); Hottung et al. (2025), MOVRPs aim to identify a Pareto-optimal set that captures the intricate trade-offs between conflicting goals. The NP-hard nature of these problems makes exact methods computationally intractable for large-scale instances, as the search space expands exponentially with problem size Florios & Mavrotas (2014). Traditional heuristics and metaheuristics, while effective for specific variants or instances, rely heavily on handcrafted rules and iterative local search, leading to high computational latency and limited adaptability when applied to different MOVRP settings.

Recently, multi-task neural combinatorial optimization (NCO) methods, including POMO-MTL Liu et al. (2024), RouteFinder Berto et al. (2025b), MVMoE Zhou et al. (2024), CoEKS Yu et al. (2026), and CCL Gui et al. (2026), have demonstrated strong cross-problem generalization across single-objective VRPs. Inspired by this, we conduct three lines of research on multi-task neural multi-objective optimization in this work, including extension of both the existing single-task neural multi-objective optimization solvers and multi-task neural single-objective optimization solver into multi-task neural multi-objective optimization paradigms, design of size-agnostic counterparts of these extensions, and a proposal of a stronger multi-task NMOCO framework with a novel tokenization scheme termed Decoupled Constraint Activation Signature Tokenization (DCAST).

At first, we reformulate 16 single-objective VRP variants from the NCO literature as bi-objective VRPs. Subsequently, two leading single-task NMOCO solvers, WE-CA and POCCO, are adapted to multi-task paradigms by equipping them with unified static node features and dynamic context embedding and training each shared policy on randomized task batches, yielding W-T and P-T, respectively. Moreover, MVMoE, a representative multi-task NCO model for single-objective optimization, is adapted to the NMOCO setting by embedding the preference vector as a pseudo-node and using it to modulate node representations through feature-wise linear modulation (FiLM) in the encoder. We further enhance the transferred MVMoE with a decoder-side nonlinear expert layer, resulting in MOM-M. Additionally, their size-agnostic counterparts, W-TS, P-TS, and MOM-MS, are developed to generalize jointly across problem variants and instance sizes.

Based on these extensions, we propose a stronger multi-task NMOCO framework built on MOM-M with a novel Decoupled Constraint Activation Signature Tokenization (DCAST), which explicitly models the constraint composition of each problem variant. DCAST describes a variant using a five-dimensional binary constraint activation signature and decouples its representation into five learnable constraint identity embeddings and two shared learnable activation-state embeddings. These representations are combined to construct state-aware constraint tokens, which are jointly encoded with node and preference embeddings. By representing an unseen variant as a new composition of reusable constraint identities and activation states, DCAST is designed to reduce dependence on the complete problem signatures encountered during training. Extensive experiments demonstrate that 1)the resulting multi-task model paradigm outperforms the task-specific solvers overall, while its size-agnostic variants further improve generalization across problem variants and instance sizes. 2) DCAST achieves the strongest aggregate performance across the evaluated problem variants, while its size-agnostic counterpart, DCAST-S, retains this advantage when transferred to larger, previously unseen instances. Our contributions are summarized as follows:

- We reformulate 16 single-objective VRP variants from the NCO literature as bi-objective MOVRPs, forming a unified benchmark for studying multi-task NMOCO.
- We adapt two leading single-task NMOCO solvers, WE-CA Chen et al. (2025a) and POCCO Fan et al. (2025b), into the multi-task models W-T and P-T, respectively, by augmenting their node representations with task-dependent static attributes and incorporating four dynamic routing attributes into their decoder context embeddings. Moreover, a representative multi-task NCO model MVMoE Zhou et al. (2024) is also adapted to the multi-objective setting as MOM by introducing a preference embedding as a pseudo-node into the encoder and using it to modulate node representations via FiLM. We further integrate a decoder-side nonlinear expert layer into MOM, yielding the stronger multi-task NMOCO framework MOM-M. Additionally, by adopting the size-agnostic training pipelines, above multi-task models are further extended to multi-task cross-size paradigms, denoted as W-TS, P-TS, and MOM-MS.

- Building on the MOM-M framework, we propose Decoupled Constraint Activation Signature Tokenization (DCAST). DCAST represents each problem variant using a five-dimensional binary constraint activation signature and factorizes its representation into five learnable constraint identity embeddings and two shared learnable activation-state embeddings. These representations are combined into state-aware constraint tokens and jointly encoded with node and preference embeddings. By expressing unseen variants as new compositions of reusable constraint identities and activation states, DCAST improves generalization across both problem variants and sizes.

2 PRELIMINARY

2.1 MOCVRP

A Capacitated Vehicle Routing Problem (CVRP) is defined on a complete graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, where node 0 denotes the depot and $\mathcal{V}\setminus\{0\}$ contains n customers with demands d_i. A feasible solution $\pi \in \Omega$ consists of routes that visit every customer exactly once while satisfying the vehicle capacity Q. Its multi-objective extension, MOCVRP, simultaneously minimizes $\mathbf{F}(\pi) = \big(f_1(\pi), \dots, f_M(\pi)\big)^\top$. In this work, we consider a bi-objective setting, where f_1 denotes the total tour length and f_2 denotes the maximum route length Lin et al. (2022).

Definition 1 (**Pareto Dominance**). *For two feasible solutions $\pi, \pi' \in \Omega$, π dominates π', denoted by $\pi \prec \pi'$, if $f_m(\pi) \le f_m(\pi')$ for every $m \in \{1, \dots, M\}$ and $f_j(\pi) < f_j(\pi')$ for at least one $j \in \{1, \dots, M\}$.*

Definition 2 (**Pareto Optimality**). *A feasible solution $\pi^\star \in \Omega$ is Pareto optimal if no $\pi' \in \Omega$ satisfies $\pi' \prec \pi^\star$. All Pareto-optimal solutions constitute the Pareto Set $PS = \{\pi^\star \in \Omega \mid \nexists\, \pi' \in \Omega : \pi' \prec \pi^\star\}$, whose image in the objective space is the Pareto Front $PF = \{\mathbf{F}(\pi^\star) \mid \pi^\star \in PS\}$.*

2.2 VARIANTS OF MOVRPS

We consider 16 bi-objective MOVRP variants formed by the capacity constraint (C) and four optional operational constraints: open route (O), backhaul (B), route-length limit (L), and time window (TW). These constraints can occur individually or jointly. Our goal is to learn a unified policy that approximates the PF across different constraint compositions, problem sizes, and instance distributions without task-specific retraining. Detailed definitions and illustrations are provided in Appendix B.

2.3 DECOMPOSITION-BASED NEURAL MOCO

Decomposition-based neural MOCO transforms an MOCVRP into preference-conditioned scalar subproblems Lin et al. (2022); Chen et al. (2025a). Given a preference vector $\boldsymbol{\lambda} \in \Lambda$ satisfying $\lambda_m \ge 0$ and $\sum_{m=1}^{M} \lambda_m = 1$, the Weighted Sum (WS) scalarization is $g(\pi \mid \boldsymbol{\lambda}) = \sum_{m=1}^{M} \lambda_m f_m(\pi)$. The policy constructs $\pi = (\pi_1, \dots, \pi_T)$ autoregressively according to $P_\theta(\pi \mid G, \boldsymbol{\lambda}) = \prod_{t=1}^{T} P_\theta(\pi_t \mid \pi_{1:t-1}, G, \boldsymbol{\lambda})$, where feasibility masks enforce the active constraints of the current MOVRP variant. Traversing Λ produces different objective trade-offs and thereby approximates the PF. The complete reinforcement-learning formulation is provided in Appendix B.

3 METHODOLOGY

3.1 OVERVIEW

Given a bi-objective routing instance G, a preference vector λ, and its constraint composition, we aim to learn a unified policy for multiple MOVRPs without task-specific retraining. As shown in Figure 1, DCAST builds on MOM-M, which provides preference-conditioned encoding and a decoder-side nonlinear expert layer. Rather than representing each problem variant as an indivisible task label, DCAST decomposes its constraint composition into constraint identities and activation states and injects the resulting state-aware constraint tokens into the encoder. The encoder jointly models the static node features, preference, and constraint states, while the decoder attends to the

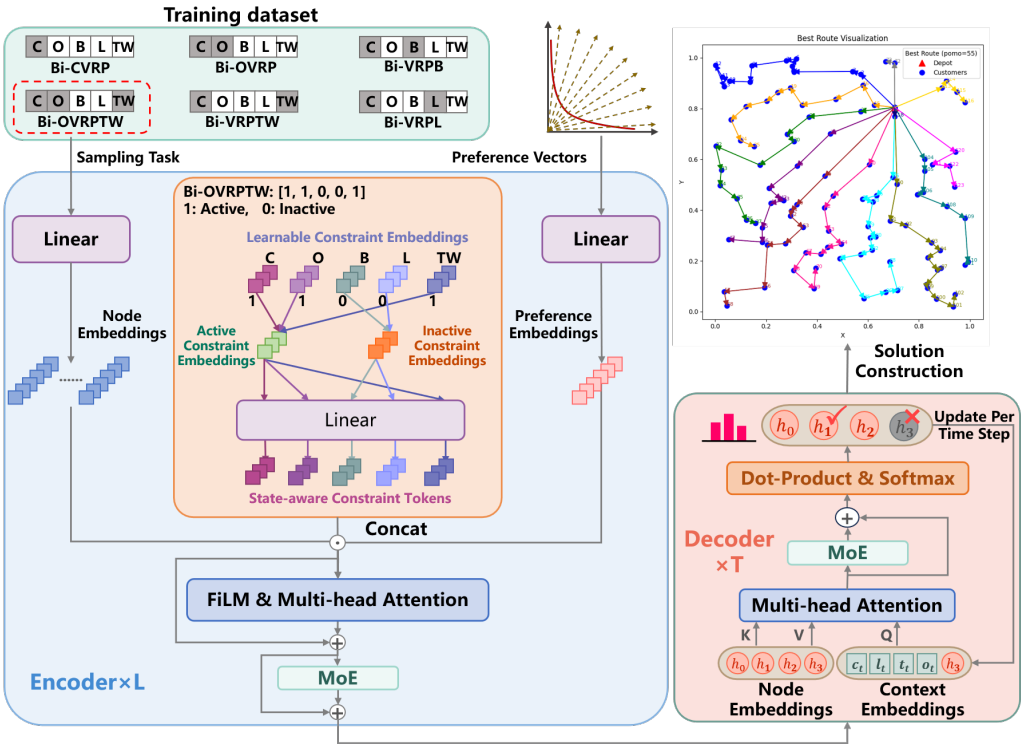


Figure 1: Overview of DCAST. Decoupled constraint-identity and activation-state tokens are jointly encoded with node and preference embeddings, enabling a shared policy to construct solutions for different MOVRP variants.

resulting representations to construct a solution autoregressively. The remainder of this section presents the formulation of DCAST.

3.2 Constraint Activation Signature

We describe the constraint composition of a problem variant τ by a five-dimensional binary *constraint activation signature*

$$\mathbf{a}^{\tau} = [a_C^{\tau}, a_O^{\tau}, a_B^{\tau}, a_L^{\tau}, a_{TW}^{\tau}] \in \{0,1\}^5, \tag{1}$$

where the entries correspond to the capacity (C), open route (O), backhaul (B), route-length (L), and time-window (TW) constraints, respectively. An entry of 0 indicates that the corresponding constraint is inactive, whereas an entry of 1 indicates that it is active. For example, Bi-CVRP has the signature $[1,0,0,0,0]$, while Bi-OVRP has $[1,1,0,0,0]$.

A constraint activation signature is compact, as each bit identifies a constraint and its activation state. but using it directly as a continuous encoder feature can tie the learned representation to the bit patterns observed during training. Such a representation is poorly aligned with our generalization setting, in which only six constraint combinations are observed during training and the remaining combinations are evaluated without retraining. DCAST therefore only uses $\mathbf{a}^{\tau}$ to select the activation-state embeddings. By this, DCAST breaks problems into individual constraints, allowing training and unseen problems to share underlying constraint structure and boost cross-variant generalization.

3.3 Decoupled Constraint Activation Signature Tokenization

DCAST factorizes the representation of equation 1 into constraint identity and activation state. Let

$$\mathbf{E}^{\mathrm{id}} = [\mathbf{e}_C, \mathbf{e}_O, \mathbf{e}_B, \mathbf{e}_L, \mathbf{e}_{TW}] \in \mathbb{R}^{5\times d} \tag{2}$$

denote five learnable constraint identity embeddings, where each vector identifies one operational constraint. We additionally introduce two shared learnable activation-state embeddings,

$$\mathbf{E}^{\text{state}} = [\mathbf{s}_0, \mathbf{s}_1] \in \mathbb{R}^{2\times d}, \tag{3}$$

representing the inactive and active states. The state embeddings are shared across all five constraints, so the meaning of activation is learned independently of any particular constraint identity.

For constraint $k \in \{C, O, B, L, TW\}$, its signature component a_k^τ selects $\mathbf{s}_{a_k^\tau}$. DCAST concatenates this state embedding with the corresponding identity embedding and projects the result back to the model dimension:

$$\mathbf{z}_k^\tau = \mathbf{W}_z[\mathbf{e}_k \parallel \mathbf{s}_{a_k^\tau}] + \mathbf{b}_z, \qquad \mathbf{z}_k^\tau \in \mathbb{R}^d, \tag{4}$$

where $\parallel$ denotes concatenation. We refer to $\mathbf{z}_k^\tau$ as a *state-aware constraint token*. Each token identifies what the constraint is and whether it applies to the current variant. Accordingly, an unseen variant is expressed as a new composition of learned constraint identities and shared activation states rather than as an unseen task-level signature. All embeddings and the projection in equation 4 are optimized jointly with the routing policy.

3.4 Encoder and Decoder

Let node 0 denote the depot and nodes $1, \ldots, n$ denote customers. The depot feature is its coordinate vector $\mathbf{x}_0 = (x_0, y_0)$, while customer i is represented by $\mathbf{x}_i = (x_i, y_i, \delta_i, e_i, l_i)$ for $i \in \{1, \ldots, n\}$. Here, δ_i denotes demand, and e_i and l_i denote the start and end of the service time window. For variants without time windows, e_i and l_i are set to zero. Separate linear projections produce the node embeddings $\mathbf{h}_0^0 = \mathbf{W}_0\mathbf{x}_0 + \mathbf{b}_0$ and $\mathbf{h}_i^0 = \mathbf{W}_n\mathbf{x}_i + \mathbf{b}_n$. The bi-objective preference $\lambda \in \mathbb{R}^2$ is mapped to the preference embedding $\mathbf{h}_\lambda^0 = \mathbf{W}_\lambda\lambda + \mathbf{b}_\lambda$.

DCAST appends the preference embedding and five state-aware constraint tokens after the node embeddings:

$$\mathbf{H}^0 = [\mathbf{h}_0^0, \mathbf{h}_1^0, \ldots, \mathbf{h}_n^0, \mathbf{h}_\lambda^0, \mathbf{z}_C^\tau, \mathbf{z}_O^\tau, \mathbf{z}_B^\tau, \mathbf{z}_L^\tau, \mathbf{z}_{TW}^\tau] \in \mathbb{R}^{(n+7)\times d}. \tag{5}$$

The first $n + 1$ positions contain the node embeddings, followed by the preference embedding and five constraint tokens.

Within each encoder layer, the preference embedding modulates the node embeddings through FiLM. Multi-head self-attention is then applied jointly to the FiLM-conditioned node embeddings, the preference embedding, and all five state-aware constraint tokens. This is where DCAST performs constraint fusion: each state-aware constraint token interacts directly with the node embeddings at every encoder layer.

After L encoder layers, the Autoregressive Decoder receives the complete encoder output $\mathbf{H}^L$. At decoding step t, the dynamic routing attributes are $\mathbf{d}_t = [c_t, t_t, \ell_t, o_t]$, where c_t, t_t, ℓ_t, and o_t denote the remaining capacity, current time, current partial-route length, and open route indicator, respectively. The context embeddings combines these attributes with the embedding of the previously selected node as $\mathbf{q}_t = \mathbf{W}_q[\mathbf{h}_{\pi_{t-1}}^L \parallel \mathbf{d}_t]$.

The context embeddings attend to the complete sequence $\mathbf{H}^L$. Following multi-head attention and the nonlinear expert layer, the decoder obtains the embedding $\widetilde{\mathbf{u}}_t$. The preference embedding and constraint tokens participate in attention but are not candidate actions. The selection logits over the depot and customer nodes are

$$\ell_{t,i} = \begin{cases} -\infty, & \text{if node } i \text{ is infeasible,} \\ C\tanh\left(\dfrac{\widetilde{\mathbf{u}}_t^\top \mathbf{h}_i^L}{\sqrt{d}}\right), & \text{otherwise,} \end{cases} \qquad i \in \{0, \ldots, n\}, \tag{6}$$

where $C = 10$ is the clipping constant. Finally, the node-selection probability is computed as

$$P_\theta(\pi_t \mid \pi_{1:t-1}, G, \lambda, \mathbf{a}^\tau) = \text{Softmax}(\boldsymbol{\ell}_t). \tag{7}$$

Detailed framework design is provided in Appendix C.

4 EXPERIMENTS

4.1 EXPERIMENTAL SETUP

Benchmark and splits. We evaluated all models on the 16 bi-objective VRP variants described in Section 2; their constraint activation signatures and training–test split are listed in Table 5. Six variants, namely Bi-CVRP, Bi-OVRP, Bi-VRPB, Bi-VRPL, Bi-VRPTW, and Bi-OVRPTW, were used for training. The remaining ten variants containing constraint combinations that were not observed during training were used to evaluate cross-problem generalization without task-specific retraining. We considered fixed-size instances with $n = 50$ and $n = 100$, together with larger $n = 200$ instances for cross-size evaluation. When appended to a variant name, the trailing number indicates the problem size; for example, Bi-CVRP-50 denotes the Bi-CVRP variant at $n = 50$. For the training variants, each single-task model was trained independently on the corresponding MOVRP. For the unseen variants, WE-CA and POCCO used their Bi-CVRP checkpoints, whereas every multi-task model used one policy shared across all variants. The training and inference configurations are described in Appendix E.

Compared models. The comparison contains three model groups. The first group comprises the single-task NMOCO solvers WE-CA and POCCO. The second group contains the multi-task NMOCO models developed by us, i.e., W-T, P-T, MOM, MOM-M, and DCAST. W-T and P-T are multi-task adaptations of WE-CA and POCCO, MOM is the multi-objective variant of MVMoE, MOM-M augments MOM with the decoder-side nonlinear expert layer, and DCAST further equips the MOM-M with state-aware constraint tokens. The third group contains their size-agnostic counterparts W-TS, P-TS, MOM-S, MOM-MS, and DCAST-S.

We report hypervolume (HV↑), Gap (↓), and runtime indicators. A larger HV indicates a Pareto-front approximation with better convergence and coverage. Gap is reported as the HV difference relative to the reference model specified for each comparison. A negative Gap therefore denotes a larger HV than the corresponding reference. Best results are identified by the bold entries in the result tables.

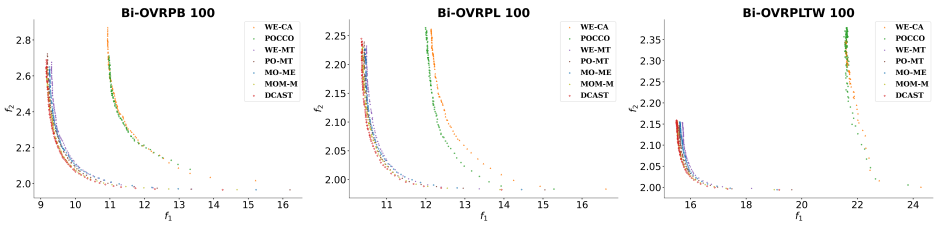
4.2 AGGREGATE PERFORMANCE ACROSS MODEL PARADIGMS

As summarized in Table 1, our multi-task model Paradigm outperformed the task-specific solvers overall across both trained and unseen MOVRP variants. Among its members, DCAST achieved the best aggregate performance. On the six training variants, DCAST obtained average Gaps of -0.17% and -0.24% for $n = 50$ and $n = 100$, respectively. On the ten unseen variants, the corresponding average Gaps were -0.37% and -0.33%. Although task-specific solvers remained competitive on individual training variants, DCAST achieved stronger average performance with one policy shared across all 16 variants.

Table 1: Average HV Gap (%) of all methods on training variants, zero-shot unseen variants, and large-scale instances (lower is better).

Method	Average Gap				
	Trained Variants		Unseen Variants		Lagre Scale Test
	n=50	n=100	n=50	n=100	n=200
WE-CA	3.31%	-0.17%	6.14%	7.70%	9.13%
POCCO	2.68%	-0.25%	5.05%	8.37%	10.51%
W-T	0.13%	0.27%	0.12%	3.18%	0.64%
P-T	-0.04%	-0.23%	-0.20%	-0.32%	-0.33%
MOM	0.00%	0.00%	0.00%	0.00%	0.82%
MOM-M	-0.10%	-0.22%	-0.22%	**-0.34%**	0.00%
DCAST	**-0.17%**	**-0.24%**	**-0.37%**	-0.33%	-0.01%
W-TS	None				-0.01%
P-TS	None				-0.27%
MOM-S	None				-0.14%
MOM-MS	None				-0.32%
DCAST-S	None				**-0.40%**

The broader comparison further distinguished cross-problem generalization from cross-size generalization. The adapted and transferred multi-task models reduced the degradation observed for single-task solvers on unseen constraint combinations. Their size-agnostic counterparts provided an



(a) Bi-OVRPB (b) Bi-OVRPL (c) Bi-OVRPLTW

Figure 2: Pareto fronts on benchmark instances

Table 2: HV, Gap, and inference time of single-task and multi-task methods on the six MOVRP variants used for training at $n = 50$ and $n = 100$.

	Method	n = 50			n = 100				Method	n = 50			n = 100		
		HV	Gap	Time	HV	Gap	Time			HV	Gap	Time	HV	Gap	Time
Bi-CVRP	WE-CA	0.3992	-0.08%	10s	0.4143	-0.15%	25s	Bi-VRPTW	WE-CA	0.3438	-0.20%	10s	0.3498	-0.14%	29s
	POCCO	**0.4007**	**-0.45%**	14s	**0.4153**	**-0.39%**	46s		POCCO	**0.3451**	**-0.58%**	21s	0.3501	-0.23%	52s
	W-T	0.3982	0.18%	8s	0.4125	0.29%	25s		W-T	0.3427	0.12%	9s	0.3478	0.43%	33s
	P-T	0.3984	0.13%	20s	0.4144	-0.17%	46S		P-T	0.3436	-0.15%	24s	**0.3507**	**-0.40%**	1m
	MOM	0.3989	0.00%	17s	0.4137	0.00%	41s		MOM	0.3431	0.00%	21s	0.3493	0.00%	53s
	MOM-M	0.3990	-0.03%	22s	0.4144	-0.17%	49s		MOM-M	0.3440	-0.26%	25s	0.3505	-0.34%	59s
	DCAST	0.3993	-0.10%	22s	0.4144	-0.17%	49s		DCAST	0.3442	-0.32%	25s	0.3505	-0.34%	1m
Bi-OVRP	WE-CA	0.4408	-0.18%	8s	0.4409	-0.20%	23s	Bi-VRPL	WE-CA	0.3974	-0.03%	11s	**0.4016**	**-0.17%**	28s
	POCCO	**0.4413**	**-0.30%**	16s	**0.4416**	**-0.36%**	43s		POCCO	0.3971	0.05%	15s	**0.4016**	**-0.17%**	42s
	W-T	0.4397	0.07%	8s	0.4396	0.09%	26s		W-T	0.3968	0.13%	8s	0.4000	0.22%	26s
	P-T	0.4408	-0.18%	22s	0.4407	-0.16%	48s		P-T	0.3970	0.08%	22s	**0.4016**	**-0.17%**	55s
	MOM	0.4400	0.00%	23s	0.4400	0.00%	43s		MOM	0.3973	0.00%	18s	0.4009	0.00%	44s
	MOM-M	0.4406	-0.14%	21s	0.4408	-0.18%	50s		MOM-M	**0.3975**	**-0.05%**	22s	**0.4016**	**-0.17%**	50s
	DCAST	0.4409	-0.20%	24s	0.4410	-0.23%	53s		DCAST	**0.3975**	**-0.05%**	24s	0.4015	-0.15%	51s
Bi-VRPB	WE-CA	0.4059	0.15%	8s	0.4268	0.05%	21s	Bi-OVRPTW	WE-CA	0.4150	-0.34%	10s	**0.4203**	**-0.38%**	28s
	POCCO	0.4061	0.10%	18s	0.4271	-0.02%	42s		POCCO	**0.4157**	**-0.51%**	21s	0.4201	-0.33%	55s
	W-T	0.4057	0.20%	8s	0.4254	0.37%	24s		W-T	0.4133	0.07%	10s	0.4177	0.24%	31s
	P-T	0.4060	0.12%	21s	0.4276	-0.14%	42s		P-T	0.4145	-0.22%	24s	0.4200	-0.31%	1.1m
	MOM	0.4065	0.00%	17s	0.4270	0.00%	41s		MOM	0.4136	0.00%	21s	0.4187	0.00%	53s
	MOM-M	0.4066	-0.02%	21s	0.4275	-0.12%	46s		MOM-M	0.4140	-0.10%	25s	0.4200	-0.31%	57s
	DCAST	**0.4070**	**-0.12%**	27s	**0.4278**	**-0.19%**	46s		DCAST	0.4146	-0.24%	32s	**0.4203**	**-0.38%**	58s

additional advantage when the test size exceeded the training range. The following analyses examine these two forms of generalization in turn.

4.3 Empirical Results

Performance on Trained MOVRPs. In this part, we test all methods on the 6 trained MOVRPs. As reported in Table 2, single-task solvers WE-CA and POCCO achieve better performance on several cases. Specifically, POCCO attained best results on Bi-CVRP-50, Bi-CVRP-100, Bi-OVRP-50, Bi-OVRP-100, Bi-VRPTW-50, and Bi-VRPL-100. after dedicated training for each of them. As a contrast, multi-task solvers achieve best results on 6 of the 12 variant–size settings, i.e., Bi-VRPB-50, Bi-VRPB-100, Bi-VRPB-50, Bi-VRPB-100, Bi-VRPL-50, Bi-VRPL-100, and Bi-OVRPTW-100, in which 4 of them are attributed to DCAST. This is unexpected since only two models are trained (one for each size) for every multi-task solver compared with 12 of that for each single-task solver. These results demonstrate that the multi-task paradigm is very competitive with conventional single-task solvers and highly promising in handling MOVRPs. Additionally, the designed DCAST exhibits the strongest advantage among all multi-task solvers, indicating that the designed decoupled constraint activation signature tokenization is an effective improvement for multi-task paradigm.

Generalization on Unseen MOVRPs. We test all methods on ten unseen MOVRPs. As reported in Table 3, the advantage of multi-task models became more pronounced on unseen MOVRPs. They outperform single-task solvers on all test cases. Particularly, the DCAST obtains the best or tied-best result on 16 of the 20 cases. On the ten unseen MOVRP variants, DCAST achieves average Gaps of -0.37% and -0.33% at $n = 50$ and $n = 100$, respectively, substantially outperforming the Bi-CVRP-trained single-task checkpoints of WE-CA (6.14% and 7.70%) and POCCO (5.05% and 8.37%). This advantage is particularly pronounced when multiple constraints interact: on Bi-VRPBTW-50, DCAST attains a Gap of -0.73%, compared with 10.88% for WE-CA and 8.79% for POCCO. Among the multi-task solvers, DCAST achieves the lowest average Gap at $n = 50$ and

Table 3: Zero-shot cross-problem generalization on the ten unseen MOVRP variants at $n = 50$ and $n = 100$, measured by HV, Gap, and inference time.

	Method	n = 50			n = 100				Method	n = 50			n = 100		
		HV	Gap	Time	HV	Gap	Time			HV	Gap	Time	HV	Gap	Time
Bi-OVRPB	WE-CA	0.4264	3.31%	10S	0.4327	3.72%	24s	Bi-OVRPL	WE-CA	0.4300	2.98%	10s	0.4286	2.44%	27s
	POCCO	0.4293	2.65%	14s	0.4215	6.21%	42s		POCCO	0.4324	2.44%	14s	0.4304	2.03%	47s
	W-T	0.4402	0.18%	8s	0.4485	0.20%	25s		W-T	0.4427	0.11%	8s	0.4389	0.09%	26s
	P-T	0.4414	-0.09%	23s	0.4503	-0.20%	50s		P-T	0.4436	-0.09%	22s	0.4400	-0.16%	55s
	MOM	0.4410	0.00%	19s	0.4494	0.00%	43s		MOM	0.4432	0.00%	18s	0.4393	0.00%	45s
	MOM-M	0.4414	-0.09%	22s	0.4504	-0.22%	47s		MOM-M	0.4436	-0.09%	23s	0.4401	-0.18%	52s
	DCAST	**0.4416**	**-0.14%**	23s	**0.4506**	**-0.27%**	47s		DCAST	**0.4440**	**-0.18%**	27s	**0.4404**	**-0.25%**	52s
Bi-VRPBL	WE-CA	0.4201	0.97%	11s	0.4214	2.02%	25s	Bi-VRPBTW	WE-CA	0.2942	10.88%	15s	0.2959	13.17%	34s
	POCCO	0.4238	0.09%	15s	0.4086	5.00%	43s		POCCO	0.3011	8.79%	19s	0.2980	12.56%	55s
	W-T	0.4232	0.24%	8s	0.4284	0.40%	25s		W-T	0.3297	0.12%	10s	0.3392	0.47%	32s
	P-T	0.4237	0.12%	21s	**0.4309**	**-0.19%**	49s		P-T	0.3315	-0.42%	25s	**0.3425**	**-0.50%**	1m
	MOM	0.4242	0.00%	18s	0.4301	0.00%	43s		MOM	0.3301	0.00%	23s	0.3408	0.00%	49s
	MOM-M	0.4243	-0.02%	21s	0.4309	-0.19%	47s		MOM-M	0.3317	-0.48%	27s	0.3425	-0.50%	56s
	DCAST	**0.4247**	**-0.12%**	24s	**0.4309**	**-0.19%**	46s		DCAST	**0.3325**	**-0.73%**	28s	0.3422	-0.41%	56s
Bi-VRPLTW	WE-CA	0.3132	9.32%	14s	0.3050	11.47%	36s	Bi-OVRPBL	WE-CA	0.4419	3.13%	11s	0.4297	3.39%	25s
	POCCO	0.3150	8.80%	18s	0.3071	10.86%	1.1m		POCCO	0.4450	2.46%	14s	0.4182	5.98%	43s
	W-T	0.3445	0.26%	10s	0.3435	0.29%	34s		W-T	0.4557	0.11%	9s	0.4435	0.29%	26s
	P-T	0.3456	-0.06%	28s	0.3460	-0.44%	1.1m		P-T	**0.4572**	**-0.22%**	23s	0.4456	-0.18%	52s
	MOM	0.3454	0.00%	22s	0.3445	0.00%	54s		MOM	0.4562	0.00%	19s	0.4448	0.00%	43s
	MOM-M	0.3462	-0.23%	26s	**0.3461**	**-0.46%**	62s		MOM-M	0.4567	-0.11%	23s	0.4456	-0.18%	47s
	DCAST	**0.3467**	**-0.38%**	28s	**0.3461**	**-0.46%**	1m		DCAST	0.4571	-0.20%	24s	**0.4458**	**-0.22%**	49s
Bi-OVRPBTW	WE-CA	0.3797	6.73%	14s	0.3664	9.84%	32s	Bi-OVRPLTW	WE-CA	0.3838	6.73%	13s	0.3716	8.38%	35s
	POCCO	0.3861	5.16%	17s	0.3647	10.26%	54s		POCCO	0.3885	5.59%	18s	0.3738	7.84%	56s
	W-T	0.4071	0.00%	11s	0.4053	0.27%	31s		W-T	0.4113	0.05%	10s	0.4046	0.25%	32s
	P-T	0.4086	-0.37%	27s	0.4077	-0.32%	1m		P-T	0.4122	-0.17%	25s	0.4067	-0.27%	58s
	MOM	0.4071	0.00%	23s	0.4064	0.00%	55s		MOM	0.4115	0.00%	23s	0.4056	0.00%	52s
	MOM-M	0.4085	-0.34%	28s	**0.4078**	**-0.34%**	57s		MOM-M	0.4120	-0.12%	27s	0.4071	-0.37%	1m
	DCAST	**0.4092**	**-0.52%**	34s	**0.4078**	**-0.34%**	57s		DCAST	**0.4126**	**-0.27%**	30s	**0.4072**	**-0.39%**	1m
Bi-VRPBLTW	WE-CA	0.3012	10.62%	16s	0.2955	13.32%	36s	Bi-OVRPBLTW	WE-CA	0.3822	6.73%	15s	0.3613	9.88%	33s
	POCCO	0.3080	8.61%	20s	0.2975	12.73%	57s		POCCO	0.3885	5.20%	19s	0.3601	10.18%	54s
	W-T	0.3366	0.12%	11s	0.3391	0.53%	35s		W-T	0.4097	0.02%	11s	0.3998	0.27%	35s
	P-T	0.3381	-0.33%	24s	**0.3429**	**-0.59%**	1m		P-T	0.4114	-0.39%	26s	0.4023	-0.35%	59s
	MOM	0.3370	0.00%	23s	0.3409	0.00%	53s		MOM	0.4098	0.00%	24s	0.4009	0.00%	54s
	MOM-M	0.3384	-0.42%	29s	0.3425	-0.47%	1m		MOM-M	0.4112	-0.34%	28s	**0.4027**	**-0.45%**	1m
	DCAST	**0.3395**	**-0.74%**	33s	0.3424	-0.44%	58s		DCAST	**0.4117**	**-0.46%**	28s	0.4023	-0.35%	58s

remains highly competitive at $n = 100$, suggesting that decoupling constraint identities from shared activation states helps the policy generalize to unseen constraint compositions. The representative Pareto fronts in Figure 2 further illustrate this behavior. More broadly, the multi-task solvers consistently achieve lower average Gaps than the single-task checkpoints at both sizes, demonstrating the value of multi-task learning for zero-shot cross-variant generalization.

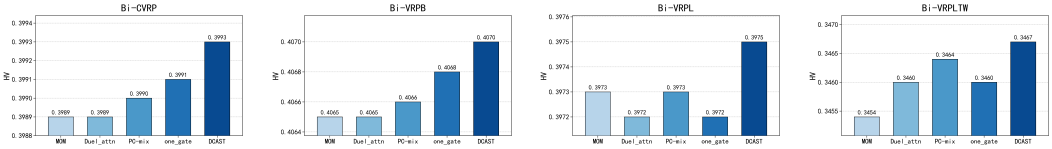
Cross-Size Generalization. We train the fixed-size single-task and multi-task solvers on instances with $n = 100$, whereas their size-agnostic counterparts, denoted by the suffix "S", are trained on instances with n sampled from $\{50, \ldots, 100\}$ according to the distribution defined in Equation 23. All solvers are subsequently evaluated zero-shot on the 16 MOVRP variants at the unseen problem size $n = 200$ to assess their cross-size generalization. As reported in Table 4, the multi-task solvers substantially outperform the single-task solvers overall, demonstrating the advantage of the multi-task paradigm for generalization to larger problem sizes. Moreover, size-agnostic training generally further improves the cross-size performance of the multi-task solvers over their fixed-size counterparts. Among the size-agnostic solvers, DCAST-S achieves the best or tied-best performance on 12 of the 16 variants and obtains the lowest overall average Gap of -0.40%, compared with -0.01%, -0.27%, -0.14%, and -0.32% for W-TS, P-TS, MOM-S, and MOM-MS, respectively. These results indicate that combining the compositional constraint-state modeling of DCAST with size-agnostic training provides the strongest overall zero-shot generalization across problem sizes.

4.4 Ablation Study

We compared DCAST with four variants that incorporate constraint and preference information differently. MOM is the transferred multi-task NMOCO baseline. Duel_attn uses two encoder branches, separately appending the constraint embeddings and preference token to the physical-node embeddings; the branch outputs are merged before the MoE layer. PC-mix instead fuses the constraint and preference embeddings into one composite embedding before appending it to the physical-node embeddings. one_gate is a model that keeps only one gating vector. As shown in

Table 4: Zero-shot cross-size generalization on all 16 MOVRP variants at $n = 200$. Fixed-size models are trained at $n = 100$, whereas size-agnostic models are trained with $n \in \{50, \ldots, 100\}$.

	Method	n = 200 HV	Gap	Time		Method	n = 200 HV	Gap	Time		Method	n = 200 HV	Gap	Time		Method	n = 200 HV	Gap	Time
Bi-CVRP	WE-CA	0.5361	-0.04%	2.6m	Bi-VRPL	WE-CA	0.5344	-0.07%	2.7m	Bi-VRPLTW	WE-CA	0.3313	18.92%	3.8m	Bi-VRPBTW	WE-CA	0.3248	20.29%	4.0m
	POCCO	0.5347	0.22%	3.9m		POCCO	0.5330	0.19%	4.1m		POCCO	0.3258	20.26%	5.3m		POCCO	0.3209	21.25%	4.8m
	W-T	0.5343	0.30%	2.1m		W-T	0.5327	0.24%	2.4m		W-T	0.4033	1.30%	3.3m		W-T	0.4020	1.35%	2.5m
	P-T	0.5377	-0.34%	5.5m		P-T	0.5362	-0.41%	5.9m		P-T	0.4082	0.10%	6.3m		P-T	0.4075	0.00%	6.4m
	MOM	0.5319	0.75%	4.5m		MOM	0.5299	0.77%	4.5m		MOM	0.4041	1.10%	6.5m		MOM	0.4029	1.13%	6.0m
	MOM-M	0.5359	0.00%	4.4m		MOM-M	0.5340	0.00%	4.6m		MOM-M	0.4086	0.00%	5.9m		MOM-M	0.4075	0.00%	5.4m
	DCAST	0.5363	-0.07%	5.1m		DCAST	0.5346	-0.11%	5.0m		DCAST	0.4074	0.29%	6.7m		DCAST	0.4072	0.07%	6.5m
	W-TS	0.5367	-0.15%	3.2m		W-TS	0.5350	-0.19%	3.2m		W-TS	0.4082	0.10%	4.2m		W-TS	0.4073	0.05%	4.0m
	P-TS	0.5374	-0.28%	4.2m		P-TS	0.5359	-0.36%	4.2m		P-TS	0.4090	-0.10%	5.0m		P-TS	0.4087	-0.29%	4.6m
	MOM-S	0.5374	-0.28%	3.6m		MOM-S	0.5360	-0.37%	4.0m		MOM-S	0.4085	0.02%	4.8m		MOM-S	0.4072	0.07%	4.5m
	MOM-MS	0.5375	-0.30%	4.4m		MOM-MS	0.5359	-0.36%	4.6m		MOM-MS	**0.4101**	**-0.37%**	5.5m		MOM-MS	**0.4094**	**-0.47%**	5.1m
	DCAST-S	**0.5382**	**-0.43%**	5.1m		DCAST-S	**0.5366**	**-0.49%**	5.0m		DCAST-S	0.4098	-0.29%	6.7m		DCAST-S	0.4091	-0.39%	6.5m
Bi-OVRP	WE-CA	0.5691	2.23%	2.6m	Bi-OVRPTW	WE-CA	0.4586	12.56%	3.4m	Bi-OVRPBTW	WE-CA	0.4389	13.82%	3.3m	Bi-OVRPBL	WE-CA	0.5618	3.57%	2.8m
	POCCO	0.5685	2.34%	3.9m		POCCO	0.4540	13.44%	4.8m		POCCO	0.4320	15.18%	4.6m		POCCO	0.5440	6.63%	3.7m
	W-T	0.5817	0.07%	2.1m		W-T	0.5205	0.76%	2.6m		W-T	0.5058	0.69%	2.3m		W-T	0.5811	0.26%	1.9m
	P-T	0.5835	-0.24%	5.1m		P-T	0.5246	-0.02%	5.7m		P-T	0.5094	-0.02%	5.8m		P-T	0.5855	-0.50%	5.3m
	MOM	0.5797	0.41%	4.6m		MOM	0.5214	0.59%	5.2m		MOM	0.5066	0.53%	6.0m		MOM	0.5783	0.74%	4.8m
	MOM-M	0.5821	0.00%	4.5m		MOM-M	0.5245	0.00%	5.3m		MOM-M	0.5093	0.00%	5.2m		MOM-M	0.5826	0.00%	4.5m
	DCAST	0.5829	-0.14%	5.0m		DCAST	0.5238	0.13%	6.0m		DCAST	0.5082	0.22%	6.3m		DCAST	0.5852	-0.45%	5.2m
	W-TS	0.5820	0.02%	3.1m		W-TS	0.5230	0.29%	3.8m		W-TS	0.5084	0.18%	3.9m		W-TS	0.5829	-0.05%	3.1m
	P-TS	0.5833	-0.21%	4.0m		P-TS	0.5247	-0.04%	4.8m		P-TS	0.5098	-0.10%	4.5m		P-TS	0.5857	-0.53%	3.8m
	MOM-S	0.5827	-0.10%	3.7m		MOM-S	0.5230	0.29%	4.5m		MOM-S	0.5086	0.14%	4.4m		MOM-S	0.5855	-0.50%	3.7m
	MOM-MS	0.5831	-0.17%	4.7m		MOM-MS	0.5248	-0.06%	5.1m		MOM-MS	0.5097	-0.08%	5.0m		MOM-MS	0.5856	-0.51%	4.3m
	DCAST-S	**0.5841**	**-0.34%**	5.0m		DCAST-S	**0.5260**	**-0.29%**	6.0m		DCAST-S	**0.5102**	**-0.18%**	6.3m		DCAST-S	**0.5861**	**-0.60%**	5.2m
Bi-VRPB	WE-CA	0.5496	1.42%	2.4m	Bi-OVRPB	WE-CA	0.5651	3.43%	2.5m	Bi-VRPLTW	WE-CA	0.3153	21.16%	3.9m	Bi-OVRPLTW	WE-CA	0.4504	12.51%	3.5m
	POCCO	0.5323	4.52%	3.6m		POCCO	0.5453	6.82%	3.6m		POCCO	0.3127	21.81%	5.0m		POCCO	0.4465	13.27%	5.0m
	W-T	0.5560	0.27%	1.8m		W-T	0.5840	0.21%	1.9m		W-T	0.3929	1.75%	2.8m		W-T	0.5108	0.78%	2.8m
	P-T	0.5614	-0.70%	3.4m		P-T	0.5884	-0.55%	4.6m		P-T	0.3997	0.05%	6.5m		P-T	0.5145	0.06%	6.0m
	MOM	0.5512	1.13%	4.4m		MOM	0.5814	0.65%	4.4m		MOM	0.3951	1.20%	6.2m		MOM	0.5116	0.62%	6.0m
	MOM-M	0.5575	0.00%	4.3m		MOM-M	0.5852	0.00%	4.2m		MOM-M	0.3999	0.00%	5.8m		MOM-M	0.5148	0.00%	5.4m
	DCAST	0.5591	-0.29%	4.8m		DCAST	0.5879	-0.46%	5.1m		DCAST	0.3989	0.25%	6.7m		DCAST	0.5138	0.19%	6.3m
	W-TS	0.5594	-0.34%	2.9m		W-TS	0.5859	-0.12%	3.0m		W-TS	0.3995	0.10%	4.2m		W-TS	0.5131	0.33%	4.0m
	P-TS	0.5612	-0.66%	3.6m		P-TS	0.5885	-0.56%	3.7m		P-TS	0.4007	-0.20%	4.7m		P-TS	0.5149	-0.02%	4.8m
	MOM-S	0.5611	-0.65%	3.4m		MOM-S	0.5884	-0.55%	3.8m		MOM-S	0.4001	-0.05%	4.9m		MOM-S	0.5134	0.27%	4.8m
	MOM-MS	0.5609	-0.61%	4.3m		MOM-MS	0.5885	-0.56%	4.1m		MOM-MS	**0.4017**	**-0.45%**	5.3m		MOM-MS	0.5149	-0.02%	5.3m
	DCAST-S	**0.5620**	**-0.81%**	4.8m		DCAST-S	**0.5892**	**-0.68%**	5.1m		DCAST-S	0.4011	-0.30%	6.7m		DCAST-S	**0.5159**	**-0.21%**	6.3m
Bi-VRPTW	WE-CA	0.3330	18.72%	3.6m	Bi-VRPBL	WE-CA	0.5565	1.49%	2.7m	Bi-OVRPL	WE-CA	0.5597	2.27%	2.8m	Bi-OVRPBLTW	WE-CA	0.4455	13.93%	2.8m
	POCCO	0.3283	19.87%	5.1m		POCCO	0.5368	4.97%	3.8m		POCCO	0.5598	2.25%	4.0m		POCCO	0.4389	15.20%	3.7m
	W-T	0.4052	1.10%	2.9m		W-T	0.5628	0.37%	2.0m		W-T	0.5724	0.05%	2.3m		W-T	0.5135	0.79%	2.5m
	P-T	0.4094	0.07%	6.8m		P-T	0.5682	-0.58%	5.2m		P-T	**0.5855**	**-2.24%**	5.1m		P-T	0.5174	0.04%	6.5m
	MOM	0.4052	1.10%	5.4m		MOM	0.5574	1.33%	4.6m		MOM	0.5704	0.40%	5.0m		MOM	0.5145	0.60%	5.9m
	MOM-M	0.4097	0.00%	5.5m		MOM-M	0.5649	0.00%	4.4m		MOM-M	0.5727	0.00%	4.6m		MOM-M	0.5176	0.00%	5.4m
	DCAST	0.4086	0.27%	6.3m		DCAST	0.5660	-0.19%	5.2m		DCAST	0.5737	-0.17%	5.3m		DCAST	0.5159	0.33%	6.3m
	W-TS	0.4095	0.05%	4.0m		W-TS	0.5658	-0.16%	3.1m		W-TS	0.5728	-0.02%	3.2m		W-TS	0.5165	0.21%	4.0m
	P-TS	0.4105	-0.20%	4.9m		P-TS	0.5679	-0.53%	3.8m		P-TS	0.5738	-0.19%	4.2m		P-TS	0.5178	-0.04%	4.6m
	MOM-S	0.4100	-0.07%	4.7m		MOM-S	0.5679	-0.53%	3.7m		MOM-S	0.5733	-0.10%	4.0m		MOM-S	0.5169	0.14%	6.0m
	MOM-MS	**0.4109**	**-0.29%**	5.6m		MOM-MS	0.5679	-0.53%	4.2m		MOM-MS	0.5737	-0.17%	4.5m		MOM-MS	0.5181	-0.10%	5.1m
	DCAST-S	**0.4109**	**-0.29%**	6.3m		DCAST-S	**0.5687**	**-0.67%**	5.2m		DCAST-S	0.5747	-0.35%	5.3m		DCAST-S	**0.5183**	**-0.14%**	6.3m



(a) Bi-CVRP (b) Bi-VRPB (c) Bi-VRPL (d) Bi-VRPLTW

Figure 3: Ablation results on four representative MOVRP variants with $n = 50$

Figure 3, DCAST outperformed the alternative architectures, supporting the effectiveness of our design.

5 Conclusion

This paper presented a unified multi-task NMOCO framework for MOVRPs. We reformulated 16 established VRP variants as bi-objective problems, adapted WE-CA and POCCO into W-T and P-T, and transferred MVMoE to the NMOCO setting as MOM, with MOM-M further incorporating a decoder-side nonlinear expert layer. Building on MOM-M, DCAST factorizes a five-dimensional constraint activation signature into constraint identity embeddings and shared activation-state embeddings to construct state-aware constraint tokens, which are jointly encoded with node and preference embeddings. Across the training and unseen variants, a single DCAST policy achieved the strongest aggregate performance, while DCAST-S attained the best average zero-shot performance on larger $n = 200$ instances. The model paradigm also reduced the training cost of covering multiple variants and generalized across heterogeneous CVRPLIB subsets. These results establish explicit, compositional constraint-state modeling as an effective basis for unified neural solvers within the evaluated bi-objective MOVRP settings.

AI USE STATEMENT

During the research, writing and preparation of this paper, no generative artificial intelligence tools were used.

REPRODUCIBILITY STATEMENT

We provide the information required to reproduce our methods and experiments throughout the paper and appendices. The formulations of the bi-objective VRP variants, optimization objective, and constraint definitions are described in Section 2 and Appendix B. The DCAST architecture is specified in Section 3, with complete encoder and decoder formulations given in Appendix C. The multi-task and size-agnostic training procedures are detailed in Appendix D. Experimental configurations, including model hyperparameters, optimization settings, preference sampling, instance sizes, inference protocol, hardware, and the training–test split of all 16 MOVRP variants, are reported in Appendix E. Additional results on training-task composition, checkpoint selection, training efficiency, and CVRPLIB instances are provided in the corresponding appendices. The source code and configurations needed to reproduce the reported experiments are available at `https://github.com/DFSDarren/DCAST`.

REFERENCES

Ehsan Ahmadi, Mostafa Zandieh, Mojtaba Farrokh, and Seyed Mohammad Emami. A multi objective optimization approach for flexible job shop scheduling problem under random machine breakdown by evolutionary algorithms. *Computers & Operations Research*, 73:56–66, 2016. ISSN 0305-0548.

Yoshua Bengio, Andrea Lodi, and Antoine Prouvost. Machine learning for combinatorial optimization: a methodological tour d'horizon. *European Journal of Operational Research*, 290(2): 405–421, 2021.

Federico Berto, Chuanbo Hua, Junyoung Park, Laurin Luttmann, 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. In *Proceedings of the 31st ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 2025a.

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. *Transactions on Machine Learning Research*, 2025b. ISSN 2835-8856.

Aigerim Bogyrbayeva, Meraryslan Meraliyev, Taukekhan Mustakhov, and Bissenbay Dauletbayev. Machine learning to solve vehicle routing problems: A survey. *IEEE Transactions on Intelligent Transportation Systems*, 25(6):4754–4772, 2024. doi: 10.1109/TITS.2023.3334976.

Kris Braekers, Katrien Ramaekers, and Inneke Van Nieuwenhuyse. The vehicle routing problem: State of the art classification and review. *Computers & Industrial Engineering*, 99:300–313, 2016. ISSN 0360-8352.

Jinbiao Chen, Jiahai Wang, Zizhen Zhang, Zhiguang Cao, Te Ye, and Siyuan Chen. Efficient meta neural heuristic for multi-objective combinatorial optimization. *Advances in neural information processing systems*, 36:56825–56837, 2023.

Jinbiao Chen, Zhiguang Cao, Jiahai Wang, Yaoxin Wu, Hanzhang Qin, Zizhen Zhang, and Yue-Jiao Gong. Rethinking neural multi-objective combinatorial optimization via neat weight embedding. In *The Thirteenth International Conference on Learning Representations*, 2025a.

Jinbiao Chen, Jiahai Wang, Zhiguang Cao, and Yaoxin Wu. Neural multi-objective combinatorial optimization via graph-image multimodal fusion. In *The Thirteenth International Conference on Learning Representations*, 2025b.

Jinbiao Chen, Jiahai Wang, Zhiguang Cao, and Yaoxin Wu. Neural multi-objective combinatorial optimization via graph-image multimodal fusion. In *The Thirteenth International Conference on Learning Representations*, 2025c.

Xinyun Chen and Yuandong Tian. Learning to perform local rewriting for combinatorial optimization. *Advances in neural information processing systems*, 32, 2019.

Darko Drakulić, Sofia Michel, and Jean-Marc Andreoli. Goal: A generalist combinatorial optimization agent learner. In *International Conference on Learning Representations*, volume 2025, pp. 52465–52488, 2025.

Mingfeng Fan, Yaoxin Wu, Zhiguang Cao, Wen Song, Guillaume Sartoretti, Huan Liu, and Guohua Wu. Conditional neural heuristic for multiobjective vehicle routing problems. *IEEE Transactions on Neural Networks and Learning Systems*, 36(3):4677–4689, 2024.

Mingfeng Fan, Jianan Zhou, Yifeng Zhang, Yaoxin Wu, Jinbiao Chen, and Guillaume Sartoretti. Preference-driven multi-objective combinatorial optimization with conditional computation. In D. Belgrave, C. Zhang, H. Lin, R. Pascanu, P. Koniusz, M. Ghassemi, and N. Chen (eds.), *Advances in Neural Information Processing Systems*, volume 38, pp. 11594–11627. Curran Associates, Inc., 2025a.

Mingfeng Fan, Jianan Zhou, Yifeng Zhang, Yaoxin Wu, Jinbiao Chen, and Guillaume Adrien Sartoretti. Preference-driven multi-objective combinatorial optimization with conditional computation. *arXiv preprint arXiv:2506.08898*, 2025b.

Zhanhong Fang, Debing Wang, Jinbiao Chen, Jiahai Wang, and Zizhen Zhang. Ucpo: A universal constrained combinatorial optimization method via preference optimization. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 40, pp. 36900–36908, 2026.

Kostas Florios and George Mavrotas. Generation of the exact pareto set in multi-objective traveling salesman and set covering problems. *Applied Mathematics and Computation*, 237:1–19, 2014.

Xiaoyu Fu, Shenshen Gu, and Chee-Meng Chew. Optimizing the multi-objective traveling salesman problem with a deep reinforcement learning algorithm using cross fusion attention networks. *Neural Networks*, pp. 107904, 2025.

Keivan Ghoseiri, Ferenc Szidarovszky, and Mohammad Jawad Asgharpour. A multi-objective train scheduling model and solution. *Transportation research part B: Methodological*, 38(10):927–952, 2004.

Shuangchun Gui, Suyu Liu, Xuehe Wang, and Zhiguang Cao. Chain-of-context learning: Dynamic constraint understanding for multi-task vrps. In C. Vondrick, B. Hariharan, C. Raffel, L. Pinto, D. Yang, and A. Faust (eds.), *International Conference on Learning Representations*, volume 2026, pp. 15776–15798, 2026.

André Hottung, Yeong-Dae Kwon, and Kevin Tierney. Efficient active search for combinatorial optimization problems. *International Conference on Learning Representations*, 2022.

André Hottung, Mridul Mahajan, and Kevin Tierney. Polynet: Learning diverse solution strategies for neural combinatorial optimization. In Y. Yue, A. Garg, N. Peng, F. Sha, and R. Yu (eds.), *International Conference on Learning Representations*, volume 2025, pp. 80417–80435, 2025.

Ziwei Huang, Jianan Zhou, Zhiguang Cao, and Yixin Xu. Rethinking light decoder-based solvers for vehicle routing problems. *CoRR*, abs/2503.00753, March 2025.

Xia Jiang, Yaoxin Wu, Yuan Wang, and Yingqian Zhang. Unco: Towards unifying neural combinatorial optimization through large language model. 2024.

Nicolas Jozefowiez, Frédéric Semet, and El-Ghazali Talbi. Multi-objective vehicle routing problems. *European Journal of Operational Research*, 189(2):293–309, 2008. ISSN 0377-2217. doi: https://doi.org/10.1016/j.ejor.2007.05.055.

Minsu Kim, Jinkyoo Park, et al. Learning collaborative policies to solve np-hard routing problems. *Advances in Neural Information Processing Systems*, 34:10418–10430, 2021.

Wouter Kool, Herke van Hoof, and Max Welling. Attention, learn to solve routing problems! In *International Conference on Learning Representations*, 2019.

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

Yeong-Dae Kwon, Jinho Choo, Byoungjip Kim, Iljoo Yoon, Youngjune Gwon, and Seungjai Min. Pomo: Policy optimization with multiple optima for reinforcement learning. *Advances in neural information processing systems*, 33:21188–21198, 2020b.

Jingwen Li, Yining Ma, Zhiguang Cao, Yaoxin Wu, Wen Song, Jie Zhang, and Yeow Meng Chee. Learning feature embedding refiner for solving vehicle routing problems. *IEEE Transactions on Neural Networks and Learning Systems*, 35(11):15279–15291, 2023.

Kaiwen Li, Tao Zhang, and Rui Wang. Deep reinforcement learning for multiobjective optimization. *IEEE transactions on cybernetics*, 51(6):3103–3114, 2020.

Xi Lin, Zhiyuan Yang, and Qingfu Zhang. Pareto set learning for neural multi-objective combinatorial optimization. *International Conference on Learning Representations*, 2022.

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*, pp. 1898–1908, 2024.

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

Fu Luo, Xi Lin, Yaoxin Wu, Zhenkun Wang, Tong Xialiang, Mingxuan Yuan, and Qingfu Zhang. Boosting neural combinatorial optimization for large-scale vehicle routing problems. In *The Thirteenth International Conference on Learning Representations*, 2025.

Mohammadreza Nazari, Afshin Oroojlooy, Lawrence Snyder, and Martin Takác. Reinforcement learning for solving the vehicle routing problem. *Advances in neural information processing systems*, 31, 2018.

Amir Ostovari, Lyes Benyoucef, and Hacène Haddou-Benderbal. Multi-objective production scheduling and workforce planning in sustainable reconfigurable manufacturing system. *International Journal on Interactive Design and Manufacturing*, 19:3803–3823, 2025.

Thibaut Vidal, Gilbert Laporte, and Piotr Matl. A concise guide to existing and emerging vehicle routing problem variants. *European Journal of Operational Research*, 286(2):401–416, 2020. ISSN 0377-2217.

Oriol Vinyals, Meire Fortunato, and Navdeep Jaitly. Pointer networks. *Advances in neural information processing systems*, 28, 2015.

Zhenkun Wang, Shunyu Yao, Genghui Li, and Qingfu Zhang. Multiobjective combinatorial optimization using a single deep reinforcement learning model. *IEEE transactions on cybernetics*, 54 (3):1984–1996, 2023.

Hong Wu, Jiahai Wang, and Zizhen Zhang. Modrl/d-am: Multiobjective deep reinforcement learning algorithm using decomposition and attention model for multiobjective optimization. In *International Symposium on Intelligence Computation and Applications*, pp. 575–589. Springer, 2019.

Yaoxin 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.

Ming Xiao, Lihua Chen, Haoxiong Feng, Zhigao Peng, and Qiong Long. Smart city public transportation route planning based on multi-objective optimization: A review. *Archives of Computational Methods in Engineering*, 31:3351–3375, 2024.

Zikang Yu, Jinbiao Chen, and Jiahai Wang. Combination-of-experts with knowledge sharing for cross-task vehicle routing problems. In *The Fourteenth International Conference on Learning Representations*, 2026.

Cong Zhang, Yaoxin Wu, Yining Ma, Wen Song, Zhang Le, Zhiguang Cao, and Jie Zhang. A review on learning to solve combinatorial optimisation problems in manufacturing. *IET Collaborative Intelligent Manufacturing*, 5(1):e12072, 2023. doi: 10.1049/cim2.12072.

Zizhen Zhang, Zhiyuan Wu, Hang Zhang, and Jiahai Wang. Meta-learning-based deep reinforcement learning for multiobjective optimization problems. *IEEE Transactions on Neural Networks and Learning Systems*, 34(10):7978–7991, 2022.

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*, pp. 42769–42789. PMLR, 2023.

Jianan Zhou, Zhiguang Cao, Yaoxin Wu, Wen Song, Yining Ma, Jie Zhang, and Chi Xu. Mvmoe: Multi-task vehicle routing solver with mixture-of-experts. In *International Conference on Machine Learning*, 2024.

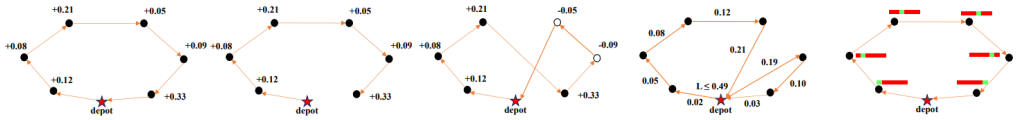
A Related Works

A.1 Neural Approaches for Single-Objective VRPs

In the field of neural combinatorial optimization (NCO), end-to-end solvers have emerged as powerful alternatives to traditional heuristics Huang et al. (2025); Bengio et al. (2021); Bogyrbayeva et al. (2024). These models typically follow an autoregressive construction paradigm, primarily bifurcated into heavy-decoder architectures (e.g., LEHD Luo et al. (2023; 2025)), which prioritize generalization through dynamic updates, and light-decoder architectures (e.g., AM Kool et al. (2019), POMO Kwon et al. (2020a)), which favor inference efficiency via static embeddings Huang et al. (2025); Vinyals et al. (2015). However, these canonical frameworks often encounter performance bottlenecks when dealing with zero-shot generalization or out-of-distribution (OOD) instances. To address these limitations, recent research has successfully enhanced model robustness across unseen tasks and complex scenarios by either strengthening architectural representational capacities Huang et al. (2025); Jiang et al. (2024); Berto et al. (2025b); Zhou et al. (2024); Liu et al. (2024); Drakulić et al. (2025); Gui et al. (2026) or implementing multi-task knowledge-sharing mechanisms Yu et al. (2026). Nevertheless, they are designed primarily for single-objective optimization and therefore do not explicitly model the preference-dependent objective trade-offs required in NMOCO.

A.2 Neural Approaches for MOVRPs

In recent years, the rise of deep reinforcement learning (DRL) has motivated increasing interest in applying neural methods to MOCO. A common strategy is decomposition, where a preference vector is used to scalarize multi-objective problems into a series of single-objective subproblems Li et al. (2020); Vinyals et al. (2015); Wu et al. (2019). Early studies followed a one-to-one paradigm, in which an independent neural model was trained or fine-tuned for each scalarized subproblem. Pointer networks Vinyals et al. (2015) were initially employed as the backbone, and subsequent works introduced attention mechanisms Kwon et al. (2020a); Kool et al. (2019) and meta-learning techniques Zhang et al. (2022) to improve generalization across subproblems. Nevertheless, this paradigm incurred prohibitively high training costs and has gradually fallen out of favor. More recent research has shifted to a one-to-many paradigm, where a single neural network is designed to handle all subproblems simultaneously Lin et al. (2022); Fu et al. (2025); Chen et al. (2025b; 2023); Fan et al. (2025a); Chen et al. (2025a); Fan et al. (2024); Wang et al. (2023). Representative methods include hypernetwork-based approaches Lin et al. (2022), fixed-preference encoders Wang et al. (2023), multimodal architectures Chen et al. (2025c), conditional neural heuristics Fan



(a) C (b) O (c) B (d) L (e) TW

Figure 4: Illustration of the five routing constraints considered in our benchmark: (a) capacity (C), (b) open route (O), (c) backhaul (B), (d) route-length limit (L), and (e) time window (TW).

et al. (2024); Chen et al. (2025a); Fang et al. (2026), and preference-optimization frameworks Fan et al. (2025a). Although these one-to-many methods solve multiple preference-induced scalarized subproblems with a single network, they remain confined to a single MOVRP variant. To the best of our knowledge, no prior NMOCO model learns a unified policy across different MOVRP variants. Consequently, a separate model must be trained whenever the constraint composition changes, resulting in substantial computational overhead and severely limited practical applicability.

B Problem Formulation and Optimization Details

CVRP formulation. The CVRP is defined on a complete graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, where $\mathcal{V} = \{0, 1, \ldots, n\}$. Node 0 is the central depot, and nodes $\mathcal{V} \setminus \{0\}$ are customers. Each customer $i \in \{1, \ldots, n\}$ has demand $d_i > 0$, and a fleet of vehicles with uniform capacity Q is stationed at the depot. A feasible solution is a set of routes such that (i) every route starts and ends at the depot unless the open route constraint is active, (ii) every customer is visited exactly once, and (iii) the total demand served by each vehicle does not exceed Q.

MOVRP variants. Capacity (C) limits the total demand served on each route. Under open route (O), a vehicle need not return to the depot after serving its final customer. Backhaul (B) introduces pickup customers in addition to delivery customers. Route-length limit (L) bounds the total length of each route. Under time window (TW), each customer must be served within a prescribed interval; an early vehicle waits, whereas a late arrival is infeasible. The four optional constraints are not mutually exclusive. Combining them with the capacity constraint produces the 16 bi-objective MOVRP variants considered in this work. Figure 4 illustrates the five constraint types.

Policy optimization. The construction process is formulated as a Markov Decision Process. At step t, the policy produces a distribution over feasible actions conditioned on the instance G, the partial solution $\pi_{1:t-1}$, and preference $\boldsymbol{\lambda}$. Using the WS cost defined in Section 2, the training objective is

$$J(\theta) = \mathbb{E}_{G\sim p(G),\, \boldsymbol{\lambda}\sim p(\boldsymbol{\lambda})} \left[\mathbb{E}_{\pi\sim P_\theta(\cdot|G,\boldsymbol{\lambda})} [g(\pi \mid \boldsymbol{\lambda})] \right]. \tag{8}$$

Equivalently, reinforcement learning maximizes the reward $R(\pi) = -g(\pi \mid \boldsymbol{\lambda})$. With a baseline $b(G, \boldsymbol{\lambda})$, the REINFORCE estimator is

$$\nabla_\theta J(\theta) \approx \mathbb{E}_{G,\boldsymbol{\lambda},\pi} \left[\big(g(\pi \mid \boldsymbol{\lambda}) - b(G, \boldsymbol{\lambda}) \big) \nabla_\theta \log P_\theta(\pi \mid G, \boldsymbol{\lambda}) \right]. \tag{9}$$

Sampling or traversing the preference space Λ trains the policy to produce solutions covering different regions of the PF.

C Encoder and Decoder Formulation

This section presents the formulation of the encoder and autoregressive decoder summarized in Section 3.4. The constraint activation signature and the construction of state-aware constraint tokens remain as defined in Sections 3.2 and 3.3.

C.1 Encoder

Let node 0 denote the depot and nodes $1, \ldots, n$ denote customers. The depot is represented by its coordinate vector $\mathbf{x}_0 = (x_0, y_0)$, while customer i is represented by

$$\mathbf{x}_i = (x_i, y_i, \delta_i, e_i, l_i), \qquad i \in \{1, \ldots, n\}, \tag{10}$$

where δ_i denotes demand, and e_i and l_i denote the start and end of the service time window. For variants without time windows, e_i and l_i are set to zero. This unified feature representation allows the same input interface to be used across all MOVRP variants considered in this work.

Separate linear projections map the depot and customer features to the model dimension:

$$\mathbf{h}_0^0 = \mathbf{W}_0\mathbf{x}_0 + \mathbf{b}_0, \qquad \mathbf{h}_i^0 = \mathbf{W}_n\mathbf{x}_i + \mathbf{b}_n. \tag{11}$$

The bi-objective preference $\lambda \in \mathbb{R}^2$ is mapped to the preference embedding

$$\mathbf{h}_\lambda^0 = \mathbf{W}_\lambda\lambda + \mathbf{b}_\lambda. \tag{12}$$

For variant τ, the preference embedding and five state-aware constraint tokens $\mathbf{z}_C^\tau$, $\mathbf{z}_O^\tau$, $\mathbf{z}_B^\tau$, $\mathbf{z}_L^\tau$, and $\mathbf{z}_{TW}^\tau$ are appended after the node embeddings. This produces the initial encoder sequence $\mathbf{H}^0$ defined in Eq. equation 5. Its first $n+1$ positions contain the node embeddings, followed by the preference embedding and five state-aware constraint tokens. This ordering is retained throughout all encoder layers.

At encoder layer ℓ, the preference embedding from the preceding layer, $\mathbf{h}_\lambda^{\ell-1}$, generates feature-wise scaling and shifting parameters that modulate only the node embeddings:

$$[\boldsymbol{\gamma}^\ell, \boldsymbol{\beta}^\ell] = \mathbf{W}_F^\ell \mathbf{h}_\lambda^{\ell-1}, \qquad \widehat{\mathbf{h}}_i^{\ell-1} = \boldsymbol{\gamma}^\ell \odot \mathbf{h}_i^{\ell-1} + \boldsymbol{\beta}^\ell, \quad i \in \{0, \dots, n\}, \tag{13}$$

where $\odot$ denotes element-wise multiplication. The resulting node sequence is

$$\widehat{\mathbf{H}}_{\text{node}}^{\ell-1} = [\widehat{\mathbf{h}}_0^{\ell-1}, \widehat{\mathbf{h}}_1^{\ell-1}, \dots, \widehat{\mathbf{h}}_n^{\ell-1}]. \tag{14}$$

The preference embedding and state-aware constraint tokens are not FiLM-modulated. Instead, they remain explicit elements of the sequence supplied to multi-head self-attention:

$$\widehat{\mathbf{H}}^{\ell-1} = [\widehat{\mathbf{h}}_0^{\ell-1}, \dots, \widehat{\mathbf{h}}_n^{\ell-1}, \mathbf{h}_\lambda^{\ell-1}, \mathbf{z}_C^{\tau,\ell-1}, \mathbf{z}_O^{\tau,\ell-1}, \mathbf{z}_B^{\tau,\ell-1}, \mathbf{z}_L^{\tau,\ell-1}, \mathbf{z}_{TW}^{\tau,\ell-1}]. \tag{15}$$

Multi-head self-attention is applied jointly over this complete sequence. Consequently, the node embeddings can exchange information with the preference embedding and all five state-aware constraint tokens. The constraint tokens can also be contextualized by the routing instance and the other constraints. The encoder preserves the same sequence structure across layers and produces

$$\mathbf{H}^L = [\mathbf{h}_0^L, \mathbf{h}_1^L, \dots, \mathbf{h}_n^L, \mathbf{h}_\lambda^L, \mathbf{z}_C^{\tau,L}, \mathbf{z}_O^{\tau,L}, \mathbf{z}_B^{\tau,L}, \mathbf{z}_L^{\tau,L}, \mathbf{z}_{TW}^{\tau,L}]. \tag{16}$$

The complete sequence $\mathbf{H}^L$ is passed to the decoder. Thus, the state-aware constraint tokens are not reduced to a single task embedding before route construction and remain individually accessible through decoder attention.

C.2 Autoregressive Decoder

At decoding step t, the decoder receives the complete encoder output $\mathbf{H}^L$. Its dynamic routing attributes are

$$\mathbf{d}_t = [c_t, t_t, \ell_t, o_t], \tag{17}$$

where c_t, t_t, ℓ_t, and o_t denote the remaining capacity, current time, current partial-route length, and open route indicator, respectively. These attributes are combined with the embedding of the previously selected node $\mathbf{h}_{\pi_{t-1}}^L$ to form the decoder query:

$$\mathbf{q}_t = \mathbf{W}_q[\mathbf{h}_{\pi_{t-1}}^L \parallel \mathbf{d}_t]. \tag{18}$$

Keys and values are computed from the complete sequence $\mathbf{H}^L$. The query can therefore attend to the encoded depot and customers, the preference embedding, and all five state-aware constraint tokens. The preference embedding and constraint tokens participate in attention but are not candidate actions.

Let $\mathbf{u}_t$ denote the output of multi-head attention. The nonlinear expert layer combines m MLP experts using a sparse top-k gate:

$$\text{MoE}(\mathbf{u}_t) = \sum_{j=1}^{m} g_j(\mathbf{u}_t) E_j(\mathbf{u}_t), \qquad \sum_{j=1}^{m} g_j(\mathbf{u}_t) = 1, \tag{19}$$

where only the selected top-k gating weights are nonzero. A residual connection and normalization produce the decoder embedding:

$$\widetilde{\mathbf{u}}_t = \mathrm{IN}\left(\mathbf{u}_t + \mathrm{MoE}(\mathbf{u}_t)\right). \tag{20}$$

Node-selection logits are computed only over the depot and customer nodes. Feasibility masks assign $-\infty$ to infeasible nodes, whereas the remaining logits are obtained from the compatibility between the decoder embedding and the corresponding node embedding:

$$\ell_{t,i} = \begin{cases} -\infty, & \text{if node } i \text{ is infeasible,} \\ C\tanh\left(\dfrac{\widetilde{\mathbf{u}}_t^\top \mathbf{h}_i^L}{\sqrt{d}}\right), & \text{otherwise,} \end{cases} \qquad i \in \{0,\dots,n\}, \tag{21}$$

where $C = 10$ is the clipping constant. Applying a softmax over the depot and customer nodes gives the node-selection probability

$$P_\theta\left(\pi_t \mid \pi_{1:t-1}, G, \lambda, \mathbf{a}^\tau\right) = \mathrm{Softmax}(\boldsymbol{\ell}_t). \tag{22}$$

After selecting π_t, the dynamic routing attributes and feasibility masks are updated according to the active constraints of variant τ. This autoregressive process continues until a complete feasible solution has been constructed.

D Multi-Task and Size-Agnostic Training

DCAST is trained on six bi-objective variants: Bi-CVRP, Bi-OVRP, Bi-VRPB, Bi-VRPL, Bi-VRPTW, and Bi-OVRPTW. Each mini-batch contains instances from one uniformly sampled training variant, and each instance is paired with a preference sampled from $\mathrm{Dirichlet}(1,1)$. The signature $\mathbf{a}^\tau$ is determined by the sampled variant and shared by all instances in that mini-batch. We train the complete model end to end by optimizing the weighted-sum objective with the REINFORCE estimator in Eq. equation 9.

DCAST uses a fixed problem size during each training run. Its size-agnostic counterpart, DCAST-S, additionally samples the number of customers from $\{50, 51, \dots, 100\}$. To place greater emphasis on more demanding large instances, the sampling probability increases linearly with size:

$$P(n) = \frac{n-49}{\sum_{r=50}^{100}(r-49)}, \qquad n \in \{50,\dots,100\}. \tag{23}$$

All other components of DCAST and DCAST-S are identical. This controlled difference allows the experiments to distinguish the generalization enabled by the compositional constraint representation from the additional effect of size-agnostic training.

E Detailed Experimental Setup

All neural solvers were optimized with Adam for 200 epochs using a learning rate of 3×10^{-4}, a weight decay of 10^{-6}, 100,000 randomly generated instances per epoch, and a batch size of 64. Training preferences were sampled from $\mathrm{Dirichlet}(1,1)$. At inference, we evaluated 101 evenly distributed preferences between the two objectives to approximate the PF. Unless stated otherwise, fixed-size models were trained separately using instances with $n = 50$ or $n = 100$, whereas size-agnostic models were trained with the number of customers sampled from $\{50,\dots,100\}$ according to Eq. equation 23. All experiments were conducted on machines equipped with a 32-GB virtual GPU and an Intel Xeon Platinum 8352V processor. Runtime is reported together with HV and Gap using the same inference protocol.

F Training-Task Composition

As shown in Table 6, we compared two training-task compositions to determine the six variants used throughout the main experiments. The orthogonal constraint suite produced an average Gap of 0.00%, compared with 0.29% for the cumulative suite, and was therefore used for all multi-task models.

Table 5: Constraint activation signatures and data splits for the 16 MOVRP variants. Signature entries follow the order $[C, O, B, L, TW]$.

Variant	Signature	Split	Variant	Signature	Split
Bi-CVRP	[1,0,0,0,0]	Training	Bi-OVRPB	[1,1,1,0,0]	Unseen
Bi-OVRP	[1,1,0,0,0]	Training	Bi-OVRPL	[1,1,0,1,0]	Unseen
Bi-VRPB	[1,0,1,0,0]	Training	Bi-VRPBL	[1,0,1,1,0]	Unseen
Bi-VRPL	[1,0,0,1,0]	Training	Bi-VRPBTW	[1,0,1,0,1]	Unseen
Bi-VRPTW	[1,0,0,0,1]	Training	Bi-VRPLTW	[1,0,0,1,1]	Unseen
Bi-OVRPTW	[1,1,0,0,1]	Training	Bi-OVRPBL	[1,1,1,1,0]	Unseen
Bi-OVRPBTW	[1,1,1,0,1]	Unseen	Bi-OVRPLTW	[1,1,0,1,1]	Unseen
Bi-VRPBLTW	[1,0,1,1,1]	Unseen	Bi-OVRPBLTW	[1,1,1,1,1]	Unseen

Table 6: Effect of training-task composition: HV on all 16 MOVRP variants with $n = 50$ under the orthogonal and cumulative constraint suites.

Problems/size50	constraint orthogonal HV	constraint increment HV
Bi-CVRP	**0.3989**	0.3986
Bi-OVRP	0.4400	**0.4405**
Bi-VRPB	**0.4065**	0.4024
Bi-OVRPB	**0.4410**	0.3309
Bi-VRPBL	**0.4242**	0.3969
Bi-VRPLTW	**0.3454**	0.4132
Bi-OVRPBTW	0.4071	**0.4424**
Bi-VRPBLTW	0.3370	**0.4200**
Bi-VRPTW	**0.3431**	0.3323
Bi-VRPL	0.3973	**0.4086**
Bi-OVRPTW	**0.4136**	0.3249
Bi-OVRPL	0.4432	**0.4435**
Bi-VRPBTW	**0.3301**	0.3168
Bi-OVRPBL	0.4562	**0.4578**
Bi-OVRPLTW	**0.4115**	0.4112
Bi-OVRPBLTW	0.4098	**0.4109**
Avg.Gap	**0.00%**	0.29%

G Single-Task Checkpoint Selection

Table 7: Cross-variant performance of WE-CA checkpoints trained on each candidate training variant, used to select the checkpoint for zero-shot evaluation.

Method	Bi-CVRP	OVRP	VRPB	VRPL	VRPTW	OVRPTW	Average
WE-CA-Bi-CVRP	**-0.08%**	2.98%	0.86%	**-0.03%**	9.44%	6.67%	3.3067%
WE-CA-Bi-OVRP	2.93%	**-0.18%**	4.72%	3.05%	13.49%	3.31%	4.5533%
WE-CA-Bi-VRPB	0.43%	3.00%	**0.15%**	0.45%	10.23%	6.75%	3.5017%
WE-CA-Bi-VRPL	-0.05%	3.00%	0.66%	**-0.03%**	9.30%	6.87%	**3.2917%**
WE-CA-Bi-VRPTW	43.52%	42.02%	91.54%	38.99%	**-0.20%**	3.19%	36.3433%
WE-CA-Bi-OVRPTW	33.84%	33.43%	78.03%	30.61%	5.71%	**-0.34%**	30.2133%

As reported in Table 7, for zero-shot evaluation on unseen variants, we used the Bi-CVRP checkpoint for the single-task models. A preliminary comparison of WE-CA checkpoints trained on each training variant showed that the Bi-CVRP checkpoint was within 0.015 percentage points of the best average Gap and attained the best result on two of the six validation variants.

H Training Efficiency

As reported in Table 8, unified multi-task training also reduced the cost of covering multiple problem variants. Training separate WE-CA models for the six training variants required 56.16 and 96.41 hours at $n = 50$ and $n = 100$, respectively, while the corresponding POCCO runs required 164.4 and 382.56 hours. In contrast, a single DCAST policy required 43.61 hours at $n = 50$ and 79.35 hours at $n = 100$. Furthermore, training a DCAST-S policy with stronger generalization capability required only 61 hours.

Table 8: Training time of task-specific, multi-task, and size-agnostic models. WE-CA and POCCO report the total time required to train six task-specific models, whereas each multi-task method uses one shared policy.

Method	n=50	n=100
WE-CA	56.16h	96.41h
POCCO	164.4h	382.56h
W-T	10.59h	30.07h
P-T	40.31h	61.22h
MOM	39.51h	74.18h
MOM-M	51.21h	86.48h
DCAST	43.61h	79.35h
W-TS	23.96h	
P-TS	64.81h	
MOM-S	65.11h	
MOM-M-S	100.19h	
DCAST-S	61.24h	

Table 9: Average Gap (%) of all methods across the six CVRPLIB benchmark subsets.

Dataset	WE-CA	POCCO	W-T	P-T	MOM	MOM-M	DCAST	W-TS	P-TS	MOM-S	MOM-MS	DCAST-S
Set-X	0.86%	-0.22%	0.40%	-0.08%	0.61%	0.00%	-0.48%	-0.08%	-0.13%	0.14%	0.20%	**-0.72%**
Set-Solomon	6.55%	5.02%	0.52%	-1.38%	-0.44%	0.00%	3.19%	**-1.78%**	1.83%	-1.17%	0.77%	0.45%
Set-A	1.00%	0.78%	**-1.43%**	0.58%	-1.42%	0.00%	1.81%	1.58%	0.28%	1.84%	1.16%	0.30%
Set-B	2.61%	1.27%	**-0.62%**	2.48%	-0.06%	0.00%	3.12%	3.93%	1.52%	2.19%	1.75%	0.46%
Set-E	0.09%	-1.66%	-2.75%	0.99%	**-3.10%**	0.00%	1.21%	-2.23%	1.33%	0.65%	0.92%	-1.55%
Set-P	0.55%	-0.32%	-1.20%	0.44%	**-1.64%**	0.00%	1.07%	-0.66%	0.56%	-0.16%	0.72%	-0.29%

I COMPLETE CVRPLIB RESULTS

The CVRPLIB evaluation assessed out-of-distribution generalization across Set-A, Set-B, Set-E, Set-P, Set-X, and Set-Solomon. Table 9 summarizes the average Gap results for all methods across these six subsets. Tables 10, 11, 12, 13, 14, and 15 report the corresponding per-instance HV and Gap results.

Across all six CVRPLIB subsets, a member of the multi-task model paradigm achieved the lowest average Gap. DCAST-S led on Set-X, W-TS on Set-Solomon, W-T on Set-A and Set-B, and MOM on Set-E and Set-P. The leading method varied across subsets, indicating complementary performance within the multi-task model paradigm.

Table 10: Per-instance HV and Gap (%) of all compared methods on the CVRPLIB Set-X instances.

Set-X	WE-CA		POCCO		W-T		P-T		MOM		MOM-M		DCAST		W-TS		P-TS		MOM-S		MOM-MS		DCAST-S	
	HV	Gap	HV	Gap	HV	Gap	HV	Gap	HV	Gap	HV	Gap	HV	Gap	HV	Gap	HV	Gap	HV	Gap	HV	Gap	HV	Gap
X-n101-k25	0.7448	0.12%	0.7458	-0.01%	0.7459	-0.03%	0.7460	-0.04%	0.7454	0.04%	0.7457	0.00%	0.7461	-0.05%	0.7449	0.11%	0.7457	0.00%	**0.7464**	**-0.09%**	0.7461	-0.05%	0.7459	-0.03%
X-n106-k14	0.6534	1.91%	**0.6679**	**-0.27%**	0.6603	0.87%	0.6635	0.39%	0.6648	0.20%	0.6661	0.00%	0.6656	0.08%	0.6630	0.47%	0.6642	0.29%	0.6628	0.50%	0.6619	0.63%	0.6678	-0.26%
X-n110-k13	0.8209	-0.02%	0.8216	-0.11%	0.8197	0.12%	**0.8217**	**-0.12%**	0.8203	0.02%	0.8207	0.00%	0.8215	-0.10%	0.8205	0.02%	0.8215	-0.10%	0.8213	-0.07%	0.8208	-0.01%	0.8207	0.00%
X-n115-k10	0.8233	0.34%	0.8130	1.59%	0.8223	0.46%	0.8242	0.23%	0.8250	0.13%	**0.8261**	**0.00%**	0.8251	0.12%	0.8232	0.35%	0.8244	0.21%	0.8247	0.17%	0.8252	0.11%	0.8256	0.06%
X-n120-k6	0.6954	0.61%	0.7013	-0.23%	0.6986	0.16%	0.7013	-0.23%	0.6988	0.13%	0.6997	0.00%	**0.7018**	**-0.30%**	0.7002	-0.07%	0.7009	-0.17%	0.7005	-0.11%	0.6986	0.16%	0.7003	-0.09%
X-n125-k30	0.6188	0.10%	0.6174	0.32%	0.6176	0.29%	0.6177	0.27%	0.6186	0.13%	0.6194	0.00%	0.6188	0.10%	0.6162	0.52%	**0.6204**	**-0.16%**	0.6177	0.27%	0.6198	-0.06%	0.6193	0.02%
X-n129-k18	0.6645	0.06%	0.6647	0.03%	0.6632	0.26%	0.6639	0.15%	0.6634	0.23%	**0.6649**	**0.00%**	0.6642	0.11%	0.6639	0.15%	0.6638	0.17%	0.6637	0.18%	0.6643	0.09%	0.6636	0.20%
X-n134-k13	0.7783	0.40%	0.7959	-1.86%	0.7827	-0.17%	0.7927	-1.45%	0.7809	0.06%	0.7814	0.00%	**0.7961**	**-1.88%**	0.7881	-0.86%	0.7918	-1.33%	0.7935	-1.55%	0.7811	0.04%	0.7943	-1.65%
X-n139-k10	0.8266	-0.28%	0.8281	-0.46%	0.8184	0.72%	0.8281	-0.46%	0.8225	0.22%	0.8243	0.00%	0.8280	-0.45%	0.8256	-0.16%	0.8254	-0.13%	0.8272	-0.35%	0.8263	-0.24%	**0.8284**	**-0.50%**
X-n143-k7	0.6841	-0.35%	**0.6846**	**-0.43%**	0.6833	-0.23%	0.6839	-0.32%	0.6819	-0.03%	0.6817	0.00%	0.6841	-0.35%	0.6832	-0.22%	0.6845	-0.41%	0.6838	-0.31%	0.6807	0.15%	0.6843	-0.38%
X-n148-k46	0.6978	0.11%	0.6989	-0.05%	0.6978	0.11%	0.6997	-0.16%	0.6962	0.34%	0.6986	0.00%	0.7001	-0.21%	0.7001	-0.21%	0.6999	-0.19%	0.6995	-0.13%	0.7000	-0.20%	**0.7006**	**-0.29%**
X-n153-k22	0.7908	0.43%	0.7960	-0.23%	0.7869	0.92%	0.7923	0.24%	0.7785	1.98%	0.7942	0.00%	**0.7982**	**-0.50%**	0.7907	0.44%	0.7960	-0.23%	0.7866	0.96%	0.7796	1.84%	0.7968	-0.33%
X-n157-k13	0.7303	3.97%	0.7731	-1.66%	0.7653	-0.63%	0.7431	2.29%	0.7819	-2.81%	0.7605	0.00%	**0.7882**	**-3.64%**	0.7639	-0.45%	0.7660	-0.72%	0.7721	-1.53%	0.7741	-1.79%	0.7825	-2.89%
X-n162-k11	0.8216	-0.09%	0.8239	-0.37%	0.8177	0.39%	0.8229	-0.24%	0.8196	0.16%	0.8209	0.00%	0.8243	-0.41%	0.8231	-0.27%	0.8222	-0.16%	0.8215	-0.07%	0.8197	0.15%	**0.8248**	**-0.48%**
X-n167-k10	0.6733	0.18%	0.6747	-0.03%	0.6739	0.09%	0.6750	-0.07%	0.6735	0.15%	0.6745	0.00%	0.6749	-0.06%	0.6739	0.09%	0.6748	-0.04%	0.6743	0.03%	0.6738	0.10%	**0.6754**	**-0.13%**
X-n172-k51	0.7258	0.10%	0.7277	-0.17%	0.7249	0.22%	0.7291	-0.36%	0.7232	0.45%	0.7265	0.00%	**0.7295**	**-0.41%**	0.7269	-0.06%	0.7279	-0.19%	0.7249	0.22%	0.7240	0.34%	0.7262	0.04%
X-n176-k26	0.6011	0.38%	0.6030	0.07%	0.5991	0.71%	0.6024	0.17%	0.6049	-0.25%	0.6034	0.00%	0.6046	-0.20%	0.6055	-0.35%	0.6033	0.02%	0.6044	-0.17%	0.6001	0.55%	**0.6063**	**-0.48%**
X-n181-k23	0.7752	-0.58%	0.7740	-0.43%	0.7669	0.49%	0.7726	-0.25%	0.7681	0.34%	0.7707	0.00%	**0.7762**	**-0.71%**	0.7755	-0.62%	0.7707	0.00%	0.7753	-0.60%	0.7732	-0.32%	0.7752	-0.58%
X-n186-k15	0.6996	0.13%	0.7007	-0.03%	0.6994	0.16%	0.7008	-0.04%	0.6962	0.61%	0.7005	0.00%	0.7010	-0.07%	0.6997	0.11%	0.7010	-0.07%	0.7016	-0.16%	0.7011	-0.09%	**0.7017**	**-0.17%**
X-n190-k8	0.6543	2.98%	**0.6984**	**-3.56%**	0.6729	0.22%	0.6596	2.19%	0.6628	1.72%	0.6744	0.00%	0.6750	-0.09%	0.6922	-2.64%	0.6683	0.90%	0.6630	1.69%	0.6603	2.09%	0.6904	-2.37%
X-n195-k51	0.7236	0.15%	0.7251	-0.05%	0.7160	1.20%	0.7261	-0.19%	0.7223	0.33%	0.7247	0.00%	0.7249	-0.03%	0.7231	0.22%	0.7264	-0.23%	0.7248	-0.01%	0.7240	0.10%	**0.7275**	**-0.39%**
X-n200-k36	0.5941	4.13%	0.6223	-0.42%	0.6078	1.92%	0.6225	-0.45%	0.5702	7.99%	0.6197	0.00%	0.6034	2.63%	0.6162	0.56%	0.6184	0.21%	0.6016	2.92%	0.6175	0.36%	**0.6233**	**-0.58%**
X-n209-k16	0.6436	-0.44%	**0.6448**	**-0.62%**	0.6420	-0.19%	**0.6448**	**-0.62%**	0.6434	-0.41%	0.6408	0.00%	0.6445	-0.58%	0.6438	-0.47%	0.6434	-0.41%	0.6438	-0.47%	0.6416	-0.12%	**0.6448**	**-0.62%**
X-n219-k73	0.3756	10.89%	0.4019	4.65%	0.4141	1.76%	0.4269	-1.28%	0.3963	5.98%	0.4215	0.00%	0.4298	-1.97%	0.4231	-0.38%	0.4304	-2.11%	0.4112	2.44%	0.4218	-0.07%	**0.4322**	**-2.54%**
X-n228-k23	0.7776	-0.12%	0.7748	0.24%	0.7676	1.17%	0.7737	0.39%	0.7717	0.64%	0.7767	0.00%	0.7774	-0.09%	0.7743	0.31%	0.7761	0.08%	0.7782	-0.19%	0.7737	0.39%	**0.7795**	**-0.36%**
X-n237-k14	0.6441	2.50%	0.6603	0.05%	0.6572	0.51%	0.6637	-0.47%	0.6621	-0.23%	0.6606	0.00%	0.6642	-0.54%	0.6627	-0.32%	0.6474	2.00%	0.6656	-0.76%	0.6515	1.38%	**0.6658**	**-0.79%**
X-n247-k50	0.7489	-3.71%	0.7341	-1.66%	0.7329	-1.50%	0.7341	-1.66%	0.7356	-1.87%	0.7221	0.00%	0.7499	-3.85%	0.7099	1.69%	0.7290	-0.96%	0.7120	1.40%	0.7226	-0.07%	**0.7526**	**-4.22%**
X-n251-k28	0.6856	0.17%	**0.6888**	**-0.29%**	0.6789	1.15%	0.6869	-0.01%	0.6808	0.87%	0.6868	0.00%	0.6870	-0.03%	0.6874	-0.09%	0.6869	-0.01%	0.6886	-0.26%	0.6861	0.10%	0.6884	-0.23%
Per-Model Gaps	0.86%		-0.22%		0.40%		-0.07%		0.61%		0.00%		-0.48%		-0.08%		-0.13%		0.14%		0.20%		**-0.72%**	
Avg.Gaps	0.320%								0.092%										**-0.118%**					

Table 11: Per-instance HV and Gap (%) of all compared methods on the CVRPLIB Set-Solomon instances.

Set-Solomon	WE-CA		POCCO		W-T		P-T		MOM		MOM-M		DCAST		W-TS		P-TS		MOM-S		MOM-MS		DCAST-S	
	HV	Gap	HV	Gap	HV	Gap	HV	Gap	HV	Gap	HV	Gap	HV	Gap	HV	Gap	HV	Gap	HV	Gap	HV	Gap	HV	Gap
R101	0.4073	14.54%	0.4046	15.11%	0.4568	4.15%	0.4777	-0.23%	0.4741	0.52%	0.4766	0.00%	0.4798	-0.67%	0.4708	1.22%	0.4668	2.06%	0.4799	-0.69%	0.4812	-0.97%	**0.4851**	**-1.78%**
R102	0.4370	12.98%	0.4399	12.41%	0.4794	4.54%	0.4952	1.39%	0.4979	0.86%	**0.5022**	**0.00%**	0.5010	0.24%	0.5008	0.28%	0.4861	3.21%	0.5009	0.26%	0.4989	0.65%	0.5012	0.20%
R103	0.4928	6.44%	0.4983	5.39%	0.4987	5.32%	0.5284	-0.32%	0.5322	-1.04%	0.5267	0.00%	**0.5287**	**-0.38%**	0.5252	0.28%	0.5203	1.22%	0.5265	0.04%	0.5263	8.00%	0.5244	0.44%
R104	0.5443	1.57%	0.5457	1.32%	0.5441	1.61%	**0.5540**	**-0.18%**	0.5514	0.29%	0.5530	0.00%	0.5512	0.33%	0.5531	-0.02%	0.5304	4.09%	0.5489	0.74%	0.5521	0.16%	0.5535	-0.09%
R105	0.4587	10.60%	0.4625	9.86%	0.5060	1.38%	0.5057	1.44%	0.5068	1.23%	0.5131	0.00%	0.5131	0.00%	0.5057	1.44%	0.5110	0.41%	0.5060	1.38%	0.5122	0.18%	**0.5158**	**-0.53%**
R106	0.4857	7.33%	0.4755	9.27%	0.5178	1.20%	**0.5269**	**-0.53%**	0.5207	0.65%	0.5241	0.00%	0.5267	-0.50%	0.5211	0.57%	0.5266	-0.48%	0.5153	1.68%	0.5265	-0.46%	0.5261	-0.38%
R107	0.5149	3.94%	0.5187	3.23%	0.5308	0.97%	0.5438	-1.46%	0.5405	-0.84%	0.5360	0.00%	0.5434	-1.38%	0.5403	-0.80%	0.5351	1.10%	0.5374	-0.26%	0.5409	-0.91%	**0.5455**	**-1.77%**
R108	0.5540	-0.86%	0.5538	-0.82%	0.5519	-0.47%	0.5594	-1.84%	0.5574	-1.47%	0.5493	0.00%	0.5605	-2.04%	0.5567	-1.35%	0.5513	-0.36%	0.5496	-0.05%	0.5497	-0.07%	**0.5607**	**-2.08%**
R109	0.4999	6.51%	0.5093	4.75%	0.5351	-0.07%	0.5357	-0.19%	0.5324	0.43%	0.5347	0.00%	0.5352	-0.09%	0.5353	-0.11%	0.5356	-0.17%	0.5362	-0.28%	0.5354	-0.13%	**0.5371**	**-0.45%**
R110	0.5178	5.04%	0.5291	2.97%	0.5426	0.50%	0.5451	0.04%	0.5418	0.64%	0.5453	0.00%	**0.5470**	**-0.31%**	0.5444	0.17%	0.5451	0.04%	0.5444	0.17%	0.5428	0.46%	0.5459	-0.11%
R111	0.5187	4.55%	0.5278	2.87%	0.5422	0.22%	0.5472	-0.70%	0.5431	0.06%	0.5434	0.00%	0.5470	-0.66%	0.5455	-0.39%	**0.5474**	**-0.74%**	0.5422	0.22%	0.5441	-0.13%	0.5431	0.06%
R112	0.5541	-0.14%	0.5536	-0.05%	0.5475	1.05%	0.5603	-1.27%	0.5572	-0.70%	0.5533	0.00%	0.5605	-1.30%	0.5602	-1.25%	0.5561	-0.51%	0.5519	0.25%	**0.5607**	**-1.34%**	0.5600	-1.21%
RC101	0.4459	15.90%	0.4543	14.32%	0.5192	2.07%	**0.5321**	**-0.36%**	0.5280	0.41%	0.5302	0.00%	0.5261	0.77%	0.5264	0.72%	0.5257	0.85%	0.5253	0.92%	0.5309	-0.13%	0.5273	0.55%
RC102	0.4740	12.32%	0.4874	9.84%	0.5289	2.16%	0.5414	-0.15%	0.5380	0.48%	0.5406	0.00%	0.5394	0.22%	0.5455	-0.91%	0.5362	0.81%	0.5367	0.72%	**0.5490**	**-1.55%**	0.5477	-1.31%
RC103	0.5245	6.12%	0.5223	6.52%	0.5310	4.96%	0.5574	0.23%	0.5581	0.11%	0.5587	0.00%	0.5513	1.32%	0.5550	0.66%	0.5446	2.52%	0.5596	-0.16%	0.5550	0.66%	**0.5637**	**-0.89%**
RC104	0.5672	0.37%	0.5642	0.90%	0.5656	0.65%	0.5755	-1.09%	0.5640	0.93%	0.5693	0.00%	0.5688	0.09%	0.5726	-0.58%	0.5699	-0.11%	0.5664	0.51%	0.5741	-0.84%	**0.5773**	**-1.41%**
RC105	0.4599	15.21%	0.4777	11.93%	0.5346	1.44%	0.5386	0.70%	0.5411	0.24%	**0.5424**	**0.00%**	0.5366	1.07%	0.5423	0.02%	0.5370	1.00%	0.5358	1.22%	0.5411	0.24%	0.5411	0.24%
RC106	0.5118	7.05%	0.5153	6.41%	0.5525	-0.35%	**0.5554**	**-0.87%**	0.5535	-0.53%	0.5506	0.00%	0.5518	-0.22%	0.5667	-2.92%	0.5525	-0.35%	0.5485	0.38%	0.5537	-0.56%	0.5540	-0.62%
RC107	0.5318	6.08%	0.5405	4.54%	0.5622	0.71%	**0.5713**	-0.90%	0.5623	0.69%	0.5662	0.00%	0.5618	0.78%	0.5542	2.12%	0.5631	0.55%	0.5640	0.39%	0.5658	0.07%	0.5664	-0.04%
RC108	0.5603	2.30%	0.5616	2.07%	0.5719	0.28%	**0.5809**	**-1.29%**	0.5746	-0.19%	0.5735	0.00%	0.5773	-0.66%	0.5750	-0.26%	0.5718	0.30%	0.5724	0.19%	0.5733	0.03%	0.5776	-0.71%
RC201	0.7205	3.78%	0.7429	0.79%	**0.7884**	**-5.29%**	0.7709	-2.95%	0.7490	-0.03%	0.7488	0.00%	0.6860	8.39%	0.7997	-6.80%	0.6837	8.69%	0.7592	-1.39%	0.7281	2.76%	0.7179	4.13%
RC202	0.6845	7.44%	0.7408	-0.18%	0.7177	2.95%	**0.7995**	**-8.11%**	0.7392	0.04%	0.7395	0.00%	0.6167	16.61%	0.7890	-6.69%	0.6920	6.42%	0.7743	-4.71%	0.7223	2.33%	0.7146	3.37%
RC203	0.7035	4.04%	0.7371	-0.55%	0.7500	-2.31%	0.7584	-3.45%	**0.7836**	**-6.89%**	0.7331	0.00%	0.6440	12.15%	0.7706	-5.12%	0.7319	0.16%	0.7683	-4.80%	0.7213	1.61%	0.7001	4.50%
RC204	0.6938	8.42%	0.7123	5.98%	0.7733	-2.07%	0.7797	-2.92%	**0.7978**	**-5.31%**	0.7576	0.00%	0.6667	12.00%	0.7728	-2.01%	0.7232	4.54%	0.7965	-5.13%	0.6446	14.92%	0.6939	8.41%
RC205	0.7254	-0.21%	0.7557	-4.39%	0.7596	-4.93%	**0.7865**	**-8.65%**	0.7137	1.41%	0.7239	0.00%	0.6398	11.62%	0.7844	-8.36%	0.7316	-1.06%	0.7726	-6.73%	0.7035	2.82%	0.7148	1.26%
RC206	0.6994	7.11%	0.7177	4.68%	0.7650	-1.61%	0.7728	-2.64%	**0.7739**	**-2.79%**	0.7529	0.00%	0.6831	9.27%	0.8271	-9.86%	0.7188	4.53%	0.7860	-4.40%	0.7532	-0.04%	0.7320	2.78%
RC207	0.7044	4.91%	0.7134	3.70%	**0.7763**	**-4.79%**	0.7703	-3.98%	0.7305	1.39%	0.7408	0.00%	0.6760	8.75%	0.7799	-5.28%	0.6838	7.69%	0.7837	-5.79%	0.7335	0.99%	0.7466	-0.78%
RC208	0.6896	9.97%	0.7062	7.81%	0.7632	0.37%	0.7541	1.55%	0.7884	-2.92%	0.7660	0.00%	0.6594	13.92%	0.8008	-4.54%	0.7288	4.86%	**0.8226**	**-7.39%**	0.7605	0.72%	0.7589	0.93%
Per-Model Gaps	6.55%		5.02%		0.52%		-1.38%		-0.44%		0.00%		3.19%		**-1.78%**		1.83%		-1.17%		1.05%		0.45%	
Avg.Gaps	5.785%								0.378%										**0.076%**					

Table 12: Per-instance HV and Gap (%) of all compared methods on the CVRPLIB Set-A instances.

Set-A	WE-CA		POCCO		W-T		P-T		MOM		MOM-M		DCAST		W-TS		P-TS		MOM-S		MOM-MS		DCAST-S	
	HV	Gap	HV	Gap	HV	Gap	HV	Gap	HV	Gap	HV	Gap	HV	Gap	HV	Gap	HV	Gap	HV	Gap	HV	Gap	HV	Gap
A-n32-k5	0.8185	1.42%	0.8295	-1.71%	0.8476	-3.68%	0.8198	1.26%	**0.8565**	**-5.16%**	0.8303	0.00%	0.8291	0.14%	0.8439	-1.67%	0.8442	-1.76%	0.8274	0.35%	0.8147	1.88%	0.8247	0.67%
A-n33-k5	0.8433	2.55%	0.8548	1.22%	**0.8789**	**-1.56%**	0.8518	1.57%	0.8680	-0.30%	0.8654	0.00%	0.8661	-0.08%	**0.8789**	**-1.56%**	0.8644	0.12%	0.8438	2.50%	0.8570	0.97%	0.8617	0.43%
A-n33-k6	0.8466	1.98%	0.8544	1.08%	0.8618	0.22%	**0.8655**	**-0.21%**	0.8660	-0.27%	0.8637	0.00%	0.8547	1.04%	0.8451	2.15%	0.8638	-0.01%	0.8532	1.22%	0.8553	0.97%	0.8625	0.14%
A-n34-k5	0.8433	-0.66%	0.8500	-1.46%	**0.8668**	**-3.46%**	0.8426	-0.57%	0.8654	-3.29%	0.8378	0.00%	0.8415	-0.44%	0.8618	-2.86%	0.8462	-1.00%	0.8260	1.41%	0.8337	0.49%	0.8495	-1.40%
A-n36-k5	0.8343	-0.99%	0.8282	-0.25%	**0.8451**	**-2.30%**	0.8428	-2.02%	0.8440	-2.17%	0.8261	0.00%	0.8301	-0.48%	0.8379	-1.43%	0.8327	-0.80%	0.8087	2.11%	0.8084	2.14%	0.8238	0.28%
A-n37-k5	0.8623	-2.29%	0.8281	1.77%	**0.8755**	**-3.86%**	0.8475	-0.53%	0.8714	-3.37%	0.8430	0.00%	0.8463	-0.39%	0.8646	-2.56%	0.8544	-1.35%	0.8200	2.73%	0.8403	0.32%	0.8343	1.03%
A-n37-k6	0.8061	2.17%	0.8165	0.91%	0.8273	-0.40%	0.8293	-0.64%	0.8336	-1.17%	0.8240	0.00%	0.8182	0.70%	0.8181	0.72%	0.8177	0.76%	0.8074	2.01%	0.8177	0.76%	**0.8340**	**-1.21%**
A-n38-k5	0.8635	-2.54%	0.8361	0.71%	**0.8682**	**-3.10%**	0.8403	0.21%	0.8600	-2.13%	0.8421	0.00%	0.8342	0.94%	0.8664	-2.89%	0.8597	-2.09%	0.8215	2.45%	0.8213	2.47%	0.8425	-0.05%
A-n39-k5	0.8234	0.99%	0.8225	1.09%	0.8387	-0.85%	0.8189	1.53%	0.8369	-0.64%	0.8316	0.00%	0.8240	0.91%	**0.8398**	**-0.99%**	0.8190	1.52%	0.8374	-0.70%	0.8260	0.67%	0.8136	2.16%
A-n39-k6	0.8027	5.19%	0.8191	3.25%	0.8423	0.51%	0.8389	0.91%	**0.8483**	**-0.20%**	0.8466	0.00%	0.8039	5.04%	0.7894	6.75%	0.8069	4.69%	0.8145	3.79%	0.8217	2.94%	0.8066	4.72%
A-n44-k6	0.8171	-0.20%	0.8023	1.62%	0.8251	-1.18%	**0.8346**	**-2.34%**	0.8206	-0.63%	0.8155	0.00%	0.7949	2.53%	0.8092	0.77%	0.8179	-0.29%	0.8133	0.27%	0.8031	1.52%	0.8139	0.20%
A-n45-k6	0.7781	3.72%	0.8007	0.93%	**0.8187**	**-1.30%**	0.7939	1.77%	0.8144	-0.77%	0.8082	0.00%	0.7876	2.55%	0.8134	-0.64%	0.8130	-0.59%	0.7830	3.12%	0.7868	2.65%	0.8088	-0.07%
A-n45-k7	0.7964	1.58%	0.8057	0.43%	**0.8238**	**-1.80%**	0.8022	0.87%	0.8199	-1.32%	0.8092	0.00%	0.7868	2.77%	0.7948	1.78%	0.8074	0.22%	0.7983	1.35%	0.7909	2.26%	0.8097	-0.06%
A-n46-k7	0.8216	-0.58%	0.8080	1.09%	0.8286	-1.43%	0.8072	1.19%	**0.8398**	**-2.80%**	0.8169	0.00%	0.8161	0.10%	0.8160	0.11%	0.8233	-0.78%	0.8077	1.13%	0.8151	0.22%	0.8229	-0.73%
A-n48-k7	0.7706	3.49%	0.7990	0.42%	0.7882	1.95%	0.8039	**-0.73%**	**0.8043**	**-0.73%**	0.7985	0.00%	0.7854	1.62%	0.7971	1.18%	0.7994	0.11%	0.7769	2.71%	0.7788	2.47%	0.8004	-0.24%
A-n53-k7	0.7969	-1.21%	0.8023	-1.89%	0.8014	-1.78%	0.7726	1.88%	**0.8055**	**-2.30%**	0.7874	0.00%	0.7840	0.43%	0.7878	-0.05%	0.7918	-0.56%	0.7785	1.13%	0.7818	0.71%	0.7761	1.44%
A-n54-k7	0.7626	2.53%	0.7676	1.89%	**0.8162**	**-4.32%**	0.7706	1.51%	0.7990	-2.12%	0.7824	0.00%	0.7643	2.31%	0.7711	1.44%	0.7751	0.93%	0.7595	2.93%	0.7780	0.56%	0.7920	-1.23%
A-n55-k9	0.8034	-0.21%	0.7906	1.38%	**0.8113**	**-1.20%**	0.7870	1.83%	0.8024	-0.09%	0.8017	0.00%	0.7876	1.76%	0.8037	-0.25%	0.7936	1.01%	0.7813	2.54%	0.7866	1.88%	0.7969	0.60%
A-n60-k9	0.7570	0.84%	0.7583	0.67%	0.7522	1.47%	0.7532	1.34%	**0.7729**	**-1.24%**	0.7634	0.00%	0.7429	2.69%	0.6891	9.73%	0.7502	1.73%	0.7421	2.79%	0.7584	0.65%	0.7526	1.41%
A-n61-k9	0.8035	1.71%	0.8057	1.44%	0.8184	-0.11%	0.7908	3.27%	**0.8201**	**-0.32%**	0.8175	0.00%	0.7981	2.37%	0.7476	8.55%	0.7907	3.28%	0.7838	4.12%	0.7903	3.33%	0.7952	2.73%
A-n62-k8	0.7259	3.57%	0.7328	2.66%	**0.7591**	**-0.84%**	0.7338	2.52%	0.7490	0.50%	0.7528	0.00%	0.7168	4.78%	0.7429	1.32%	0.7342	2.47%	0.7426	1.35%	0.7378	1.99%	0.7455	0.97%
A-n63-k9	0.7093	3.10%	0.7285	0.48%	**0.7339**	**-0.26%**	0.7237	1.13%	0.7249	0.97%	0.7320	0.00%	0.7018	4.13%	0.7058	3.58%	0.7184	1.86%	0.7135	2.53%	0.7204	1.58%	0.7270	0.68%
A-n63-k10	0.7494	1.26%	**0.7816**	**-2.98%**	0.7728	-1.82%	0.7703	-1.49%	0.7756	-2.19%	0.7590	0.00%	0.7400	2.50%	0.7356	3.08%	0.7581	0.12%	0.7655	-0.86%	0.7733	-1.88%	0.7714	-1.63%
A-n64-k9	0.7456	-1.73%	0.7310	0.26%	0.7572	-3.32%	0.7578	-3.40%	**0.7598**	**-3.67%**	0.7329	0.00%	0.7273	0.76%	0.7538	-2.85%	0.7244	1.16%	0.7184	1.98%	0.7428	-1.35%	0.7494	-2.25%
A-n65-k9	0.7741	-1.82%	0.7488	1.07%	**0.7743**	**-2.17%**	0.7414	2.49%	0.7717	-1.50%	0.7603	0.00%	0.7349	3.34%	0.7531	0.95%	0.7706	-1.35%	0.7507	1.26%	0.7592	0.14%	0.7699	-1.26%
A-n69-k9	0.7229	4.00%	0.7451	1.05%	0.7768	-3.16%	0.7556	-0.35%	**0.7881**	**-4.66%**	0.7530	0.00%	0.7200	4.38%	0.7381	1.98%	0.7714	-2.44%	0.7456	0.98%	0.7439	1.21%	0.7586	-0.74%
A-n80-k10	**0.6810**	**-0.99%**	0.6691	0.77%	0.6721	0.33%	0.6740	0.04%	0.6666	1.14%	0.6743	0.00%	0.6583	2.37%	0.5575	17.32%	0.6688	0.82%	0.6580	2.42%	0.6752	-0.13%	0.6642	1.50%
Per-Model Gaps	1.00%		0.78%		-1.43%		0.58%		**-1.42%**		0.00%		1.81%		1.58%		0.28%		1.84%		1.16%		0.30%	
Avg.Gaps	0.890%								**-0.092%**										1.032%					

Table 13: Per-instance HV and Gap (%) of all compared methods on the CVRPLIB Set-B instances.

Set-B	WE-CA		POCCO		W-T		P-T		MOM		MOM-M		DCAST		W-TS		P-TS		MOM-S		MOM-MS		DCAST-S	
	HV	Gap	HV	Gap	HV	Gap	HV	Gap	HV	Gap	HV	Gap	HV	Gap	HV	Gap	HV	Gap	HV	Gap	HV	Gap	HV	Gap
B-n31-k5	0.8063	3.01%	**0.8358**	**-0.54%**	0.8205	1.30%	0.8268	0.54%	0.8350	-0.45%	0.8313	0.00%	0.8253	0.72%	0.8183	1.56%	0.8304	0.11%	0.8258	0.66%	0.8233	0.96%	0.8307	0.07%
B-n34-k5	0.8061	-1.37%	0.7997	-0.57%	**0.8203**	**-3.16%**	0.7760	2.41%	0.7929	0.29%	0.7952	0.00%	0.7676	3.47%	0.8014	-0.78%	0.7789	1.05%	0.7858	1.18%	0.7935	0.21%	0.7802	1.89%
B-n35-k5	0.7482	-1.20%	0.7402	-0.12%	0.7508	-1.56%	0.7279	1.54%	0.7502	-1.47%	0.7393	0.00%	0.7222	2.31%	**0.7538**	**-1.96%**	0.7438	-0.61%	0.7393	0.00%	0.7260	1.80%	0.7436	-0.58%
B-n38-k6	0.7714	2.07%

Table 14: Per-instance HV and Gap (%) of all compared methods on the CVRPLIB Set-E instances.

Set-E	WE-CA		POCCO		W-T		P-T		MOM		MOM-M		DCAST		W-TS		P-TS		MOM-S		MOM-MS		DCAST-S	
	HV	Gap	HV	Gap	HV	Gap	HV	Gap	HV	Gap	HV	Gap	HV	Gap	HV	Gap	HV	Gap	HV	Gap	HV	Gap	HV	Gap
E-n22-k4	0.8663	1.06%	0.8954	-2.26%	0.8934	-2.03%	0.8853	-1.11%	0.8949	-2.20%	0.8756	0.00%	**0.9026**	**-3.08%**	0.8953	-2.25%	0.8892	-1.55%	0.8796	-0.46%	0.8750	0.07%	0.8914	-1.80%
E-n23-k3	0.8274	-1.04%	**0.8447**	**-3.15%**	0.8421	-2.83%	0.7799	4.76%	0.8378	-2.31%	0.8189	0.00%	0.7845	4.20%	0.8437	-3.03%	0.7780	4.99%	0.8344	-1.89%	0.8252	-0.77%	0.8363	-2.12%
E-n30-k3	0.8464	-5.29%	0.8415	-4.68%	0.8541	-6.24%	0.8365	-4.06%	**0.8557**	**-6.44%**	0.8039	0.00%	0.8349	-3.86%	0.8479	-5.47%	0.8223	-2.29%	0.8384	-4.29%	0.8228	-2.35%	0.8493	-5.65%
E-n33-k4	0.7540	0.21%	0.7735	-2.37%	0.7526	0.40%	0.7653	-1.28%	0.7811	-3.37%	0.7556	0.00%	0.7815	-3.43%	0.7601	-0.60%	**0.7968**	**-5.45%**	0.7443	1.50%	0.7605	-0.65%	0.7735	-2.37%
E-n51-k5	0.8146	-0.77%	0.7964	1.48%	0.8363	-3.45%	0.8020	0.79%	**0.8463**	**-4.69%**	0.8084	0.00%	0.8019	0.80%	0.8330	-3.04%	0.7941	1.77%	0.7975	1.35%	0.7947	1.69%	0.8170	-1.06%
E-n76-k7	0.7544	-1.96%	0.7280	1.61%	0.7601	-2.73%	0.7251	2.00%	**0.7865**	**-6.30%**	0.7399	0.00%	0.7232	2.26%	0.7802	-5.45%	0.7165	3.16%	0.7211	2.54%	0.7258	1.91%	0.7344	0.74%
E-n76-k8	0.7306	1.43%	0.7510	-1.32%	**0.7718**	**-4.13%**	0.7359	0.72%	0.7531	-1.61%	0.7412	0.00%	0.7247	2.23%	0.7550	-1.86%	0.7319	1.25%	0.7193	2.95%	0.7385	0.36%	0.7338	1.00%
E-n76-k10	0.7459	0.09%	0.7576	-1.47%	**0.7755**	**-3.87%**	0.7271	2.61%	0.7444	0.29%	0.7466	0.00%	0.7482	-0.21%	0.7519	-0.71%	0.7359	1.43%	0.7577	-1.49%	0.7251	2.88%	0.7636	-2.28%
E-n76-k14	0.7350	0.77%	**0.7628**	**-2.98%**	0.7248	2.15%	0.7342	0.88%	0.7614	-2.79%	0.7407	0.00%	0.7221	2.51%	0.7574	-2.25%	0.7292	1.55%	0.7201	2.78%	0.7343	0.86%	0.7354	0.72%
E-n101-k8	0.6541	4.78%	0.6856	0.19%	**0.7220**	**-5.11%**	0.6597	3.96%	0.7185	-4.60%	0.6869	0.00%	0.6353	7.51%	0.6956	-1.27%	0.6374	7.21%	0.6551	4.63%	0.6622	3.60%	0.7062	-2.81%
E-n101-k14	0.6645	1.73%	**0.6983**	**-3.27%**	0.6925	-2.41%	0.6652	1.63%	0.6764	-0.03%	0.6762	0.00%	0.6469	4.33%	0.6666	1.42%	0.6592	2.51%	0.6797	-0.52%	0.6591	2.53%	0.6857	-1.40%
Per-Model Gaps	0.09%		-1.66%		**-2.75%**		0.99%		-3.10%		0.00%		1.21%		-2.23%		1.33%		0.65%		0.92%		-1.55%	
Avg.Gaps	**-0.785%**				-0.730%										-0.176%									

Table 15: Per-instance HV and Gap (%) of all compared methods on the CVRPLIB Set-P instances.

Set-P	WE-CA		POCCO		W-T		P-T		MOM		MOM-M		DCAST		W-TS		P-TS		MOM-S		MOM-MS		DCAST-S	
	HV	Gap	HV	Gap	HV	Gap	HV	Gap	HV	Gap	HV	Gap	HV	Gap	HV	Gap	HV	Gap	HV	Gap	HV	Gap	HV	Gap
P-n16-k8	0.9199	1.30%	0.9314	0.06%	0.9290	0.32%	**0.9340**	**-0.21%**	0.9320	0.00%	0.9320	0.00%	0.9286	0.36%	0.9319	0.01%	0.9327	-0.08%	0.9318	0.02%	0.9314	0.06%	0.9325	-0.05%
P-n19-k2	0.9402	-0.85%	0.9312	0.12%	0.9376	-0.57%	0.9399	-0.82%	0.9398	-0.80%	0.9323	0.00%	0.9286	0.40%	**0.9415**	**-0.99%**	0.9335	-0.13%	0.9401	-0.84%	0.9282	0.44%	0.9338	-0.16%
P-n20-k2	0.9385	-0.95%	0.9322	-0.27%	0.9393	-1.03%	0.9321	-0.26%	0.9371	-0.80%	0.9297	0.00%	0.9252	0.48%	0.9394	-1.04%	0.9316	-0.20%	**0.9419**	**-1.31%**	0.9207	0.97%	0.9285	0.13%
P-n21-k2	0.9313	0.60%	0.9303	0.70%	0.9355	0.15%	0.9271	1.05%	0.9375	-0.06%	0.9369	0.00%	0.9261	1.15%	0.9332	0.39%	0.9342	0.29%	**0.9394**	**-0.27%**	0.9164	2.19%	0.9282	0.93%
P-n22-k2	0.9241	0.69%	0.9285	0.21%	**0.9404**	**-1.06%**	0.9290	0.16%	0.9379	-0.80%	0.9305	0.00%	0.9217	0.95%	0.9243	0.67%	0.9226	0.85%	0.9386	-0.87%	0.9190	1.24%	0.9230	0.81%
P-n22-k8	0.8804	1.80%	0.8955	0.11%	0.9004	-0.44%	0.9042	-0.86%	0.8940	0.28%	0.8965	0.00%	**0.9053**	**-0.98%**	0.8921	0.49%	0.8963	0.02%	0.8943	0.25%	0.8986	-0.23%	0.9027	-0.69%
P-n23-k8	0.9129	0.17%	0.9115	0.33%	0.9106	0.43%	0.9125	0.22%	0.9131	0.15%	**0.9145**	**0.00%**	0.9100	0.49%	0.9015	1.42%	0.9117	0.31%	0.9080	0.71%	0.9135	0.11%	0.9137	0.09%
P-n40-k5	0.8735	-1.94%	0.8697	-1.49%	0.8798	-2.67%	0.8603	-0.40%	**0.8841**	**-3.17%**	0.8569	0.00%	0.8565	0.05%	0.8770	-2.35%	0.8585	-0.19%	0.8618	-0.57%	0.8487	0.96%	0.8535	0.40%
P-n45-k5	0.8411	-0.36%	0.8398	-0.20%	0.8636	-3.04%	0.8467	-1.03%	**0.8718**	**-4.02%**	0.8381	0.00%	0.8344	0.44%	0.8495	-1.36%	0.8474	-1.11%	0.8435	-0.64%	0.8316	0.78%	0.8440	-0.70%
P-n50-k7	0.8471	0.42%	0.8480	0.32%	0.8637	-1.53%	0.8334	2.03%	**0.8722**	**-2.53%**	0.8507	0.00%	0.8504	0.04%	0.8622	-1.35%	0.8459	0.56%	0.8451	0.66%	0.8498	0.11%	0.8408	1.16%
P-n50-k8	0.8418	0.54%	0.8550	-1.02%	0.8552	-1.04%	0.8397	0.79%	**0.8634**	**-2.01%**	0.8464	0.00%	0.8338	1.49%	0.8495	-0.37%	0.8511	-0.56%	0.8428	0.43%	0.8354	1.30%	0.8471	-0.08%
P-n50-k10	0.8453	0.41%	0.8482	0.07%	**0.8531**	**-0.51%**	0.8380	1.27%	0.8490	-0.02%	0.8488	0.00%	0.8388	1.18%	0.8435	0.62%	0.8481	0.08%	0.8481	0.08%	0.8379	1.28%	0.8399	1.05%
P-n51-k10	0.8141	2.88%	0.8316	0.79%	0.8329	0.63%	0.8370	0.14%	**0.8467**	**-1.01%**	0.8382	0.00%	0.8257	1.49%	0.8422	-0.48%	0.8372	0.12%	0.8378	0.05%	0.8349	0.39%	0.8295	1.04%
P-n55-k7	0.8295	-0.04%	0.8440	-1.78%	0.8480	-2.27%	0.8304	-0.14%	**0.8593**	**-3.63%**	0.8292	0.00%	0.8353	-0.74%	0.8469	-2.13%	0.8335	-0.52%	0.8581	-3.49%	0.8369	-0.93%	0.8349	-0.69%
P-n55-k8	0.8274	0.91%	0.8430	-0.96%	0.8461	-1.33%	0.8361	-0.13%	0.8655	-3.65%	0.8350	0.00%	0.8350	0.00%	**0.8695**	**-4.13%**	0.8313	0.44%	0.8626	-3.31%	0.8349	0.01%	0.8336	0.17%
P-n55-k10	0.8306	1.45%	0.8419	0.11%	0.8379	0.58%	0.8346	0.97%	**0.8475**	**-0.56%**	0.8428	0.00%	0.8308	1.42%	0.8405	0.27%	0.8346	0.97%	0.8271	1.86%	0.8253	2.08%	0.8430	-0.02%
P-n55-k15	0.8202	0.01%	**0.8309**	**-1.29%**	0.8184	0.23%	0.8297	-1.15%	0.8242	-0.48%	0.8203	0.00%	0.8099	1.27%	0.7892	3.79%	0.8239	-0.44%	0.8149	0.66%	0.8166	0.45%	0.8266	-0.77%
P-n60-k10	0.8071	1.03%	0.8136	0.23%	0.8214	-0.72%	0.8013	1.74%	0.8272	-1.43%	0.8155	0.00%	0.7983	2.11%	**0.8317**	**-1.99%**	0.8097	0.71%	0.8068	1.07%	0.8040	1.41%	0.8253	-1.20%
P-n60-k15	0.8074	-0.79%	0.7986	0.31%	0.8074	-0.79%	0.8077	-0.82%	0.8052	-0.51%	0.8011	0.00%	0.7940	0.89%	0.7670	4.26%	**0.8133**	**-1.52%**	0.7997	0.17%	0.7957	0.67%	0.8090	-0.99%
P-n65-k10	0.7787	1.82%	0.8075	-1.82%	**0.8123**	**-2.42%**	0.7849	1.03%	0.8122	-2.41%	0.7931	0.00%	0.7910	0.26%	0.8058	-1.60%	0.7896	0.44%	0.8061	-1.64%	0.7876	0.69%	0.7979	-0.61%
P-n70-k10	0.7741	1.43%	**0.7987**	**-1.71%**	0.7966	-1.44%	0.7829	0.31%	0.7951	-1.25%	0.7853	0.00%	0.7894	-0.52%	0.7826	0.34%	0.7707	1.86%	0.7765	1.12%	0.7914	-0.78%	0.7954	-1.29%
P-n76-k4	0.7851	-0.72%	0.7796	-0.01%	0.8054	-3.32%	0.7487	3.95%	**0.8102**	**-3.94%**	0.7795	0.00%	0.7542	3.25%	0.8084	-3.71%	0.7396	5.12%	0.7722	0.94%	0.7784	0.14%	0.8064	-3.45%
P-n76-k5	0.7744	-0.81%	0.7749	-0.87%	0.8030	-4.53%	0.7503	2.33%	0.8033	-4.57%	0.7682	0.00%	0.7749	-0.87%	**0.8066**	**-5.00%**	0.7590	1.20%	0.7695	-0.17%	0.7756	-0.96%	0.7797	-1.50%
P-n101-k4	0.7044	4.27%	0.7326	0.43%	**0.7534**	**-2.39%**	0.7333	0.34%	0.7523	-2.24%	0.7358	0.00%	0.6538	11.14%	0.7476	-1.60%	0.6978	5.16%	0.7268	1.22%	0.7000	4.87%	0.7390	-0.43%
Per-Model Gaps	0.55%		-0.32%		-1.20%		0.44%		**-1.64%**		0.00%		1.07%		-0.66%		0.56%		-0.16%		0.72%		-0.29%	
Avg.Gaps	0.115%				**-0.266%**										0.034%									