

Singapore Management University

Institutional Knowledge at Singapore Management University

Research Collection School Of Computing and Information Systems

School of Computing and Information Systems

12-2023

Efficient meta neural heuristic for multi-objective combinatorial optimization

Jinbiao CHEN

Zizhen ZHANG

Te YE

Zhiguang CAO

Singapore Management University, zgcao@smu.edu.sg

Siyuan CHEN

See next page for additional authors

Follow this and additional works at: https://ink.library.smu.edu.sg/sis_research



Part of the [Databases and Information Systems Commons](#)

Citation

CHEN, Jinbiao; ZHANG, Zizhen; YE, Te; CAO, Zhiguang; CHEN, Siyuan; and WANG, Jiahai. Efficient meta neural heuristic for multi-objective combinatorial optimization. (2023). *Proceedings of the 37th Conference on Neural Information Processing, New Orleans, United States, December 12-14*. 1-25. Available at: https://ink.library.smu.edu.sg/sis_research/8403

This Conference Proceeding Article is brought to you for free and open access by the School of Computing and Information Systems at Institutional Knowledge at Singapore Management University. It has been accepted for inclusion in Research Collection School Of Computing and Information Systems by an authorized administrator of Institutional Knowledge at Singapore Management University. For more information, please email cherylds@smu.edu.sg.

Author

Jinbiao CHEN, Zizhen ZHANG, Te YE, Zhiguang CAO, Siyuan CHEN, and Jiahai WANG

Efficient Meta Neural Heuristic for Multi-Objective Combinatorial Optimization

Jinbiao Chen¹, Jiahai Wang^{1,2,3,*}, Zizhen Zhang^{1,*}, Zhiguang Cao⁴, Te Ye¹, Siyuan Chen¹

¹School of Computer Science and Engineering, Sun Yat-sen University, P.R. China

²Key Laboratory of Machine Intelligence and Advanced Computing, Ministry of Education, Sun Yat-sen University, P.R. China

³Guangdong Key Laboratory of Big Data Analysis and Processing, Guangzhou, P.R. China

⁴School of Computing and Information Systems, Singapore Management University, Singapore

chenjb69@mail2.sysu.edu.cn, {wangjiahai, zhangzizhen}@mail.sysu.edu.cn

zgcao@smu.edu.sg, {yete, chensiyuan}@mail2.sysu.edu.cn

Abstract

Recently, neural heuristics based on deep reinforcement learning have exhibited promise in solving multi-objective combinatorial optimization problems (MOCOPs). However, they are still struggling to achieve high learning efficiency and solution quality. To tackle this issue, we propose an efficient meta neural heuristic (EMNH), in which a meta-model is first trained and then fine-tuned with a few steps to solve corresponding single-objective subproblems. Specifically, for the training process, a (partial) architecture-shared multi-task model is leveraged to achieve parallel learning for the meta-model, so as to speed up the training; meanwhile, a scaled symmetric sampling method with respect to the weight vectors is designed to stabilize the training. For the fine-tuning process, an efficient hierarchical method is proposed to systematically tackle all the subproblems. Experimental results on the multi-objective traveling salesman problem (MOTSP), multi-objective capacitated vehicle routing problem (MOCVRP), and multi-objective knapsack problem (MOKP) show that, EMNH is able to outperform the state-of-the-art neural heuristics in terms of solution quality and learning efficiency, and yield competitive solutions to the strong traditional heuristics while consuming much shorter time.

1 Introduction

Multi-objective combinatorial optimization problems (MOCOPs) [1] are widely studied and applied in many real-world sectors, such as telecommunication, logistics, manufacturing, and inventory. Typically, an MOCOP requires the simultaneous optimization of multiple conflicting objectives, where the amelioration of an objective may lead to the deterioration of others. Therefore, a set of trade-off solutions, known as *Pareto-optimal* solutions, are usually sought for MOCOPs.

Generally, it is difficult to exactly find all the Pareto-optimal solutions of an MOCOP [2], especially given that the decomposed single-objective subproblem might already be NP-hard. Hence, heuristic methods [3] are usually preferred to solve MOCOPs in reality, as they can attain approximate Pareto-optimal solutions in (relatively) reasonable time. Nevertheless, the traditional heuristics are still lacking due to their reliance on handcrafted rules and massive iterative steps, and superfluous computation even for instances of the same (or similar) class.

Recently, inspired by the success of deep reinforcement learning (DRL) in learning neural heuristics for solving the single-objective combinatorial optimization problems (COPs) [4–7], a number of

*Jiahai Wang and Zizhen Zhang are the corresponding authors.

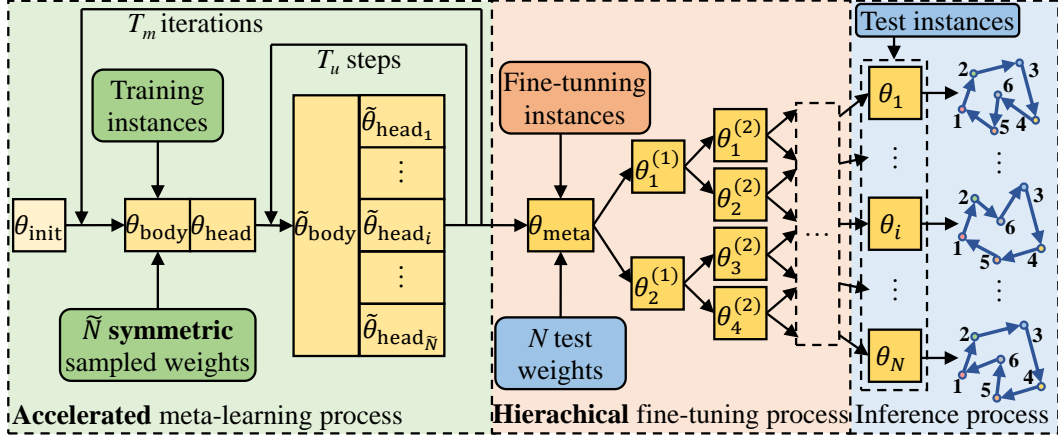


Figure 1: The overall framework of EMNH.

DRL-based neural heuristics [8–13] have also been investigated for MOCOPs. While bypassing the handcrafted rules, these neural heuristics adopt an end-to-end paradigm to construct solutions without iterative search. Benefiting from a large amount of data (i.e., problem instances), a well-trained deep model allows the neural heuristic to automatically extract informative and expressive features for decision making and generalize to unseen instances as well.

Although demonstrating promise for solving MOCOPs, those neural heuristics still struggle to achieve high learning efficiency and solution quality. Particularly, in this line of works, the early attempts [8–11] always train multiple deep models, i.e., one for each preference (weight) combination, making them less practical. While the well-known preference-conditioned multi-objective combinatorial optimization (PMOCO) [12] realized a unified deep model by introducing a huge hypernetwork, it leaves considerable gaps in terms of solution quality. The recent Meta-DRL (MDRL) [13] has demonstrated the capability to enhance solution quality over existing state-of-the-art algorithms. However, it still faces challenges related to inefficient and unstable training procedures, as well as undesirable fine-tuning processes. This paper thereby proposes an efficient meta neural heuristic (EMNH) to further strengthen the learning efficiency and solution quality, the framework of which is illustrated in Figure 1. Following the meta-learning paradigm [13, 14], EMNH first trains a meta-model and then quickly fine-tunes it according to the weight (preference) vector to tackle the corresponding single-objective subproblems.

Our contributions are summarized as follows. (1) We propose an efficient meta-learning scheme to accelerate the training. Inspired by the feature reuse of meta-learning [15], we introduce a multi-task model composed of a parameter-shared *body* (all other layers of neural network except the last one) and respective task-related *heads* (the last layer of the network), rather than respective submodels, to train the meta-model, which is able to handle different tasks (subproblems) in parallel. (2) We design a scaled symmetric sampling method regarding the weight vectors to stabilize the training. In each meta-iteration, once a weight vector is sampled, its scaled symmetric weight vectors will also be engendered, which help avoid fluctuations of the parameter update for the meta-model, especially on problems with objectives of imbalanced domains. (3) We present a hierarchical fine-tuning method to efficiently cope with all the subproblems. Since it is redundant to derive a specific submodel for each subproblem in the early tuning process, our meta-model can be gradually fine-tuned to act as specialized submodels from low level to high level, which takes much fewer fine-tuning steps in total. Moreover, experimental results on three classic MOCOPs confirmed the effectiveness of our designs.

2 Related works

Exact and heuristic methods for MOCOPs. Over the past decades, MOCOPs have been attracting increasing attention from the computational intelligence community, and tremendous works have been proposed. Generally, exact methods [2] are able to gain the accurate Pareto-optimal solutions, but their computational time may exponentially grow. Heuristic methods such as multi-objective

evolutionary algorithms (MOEAs) are popular in practice, where the representative paradigms for the general multi-objective optimization include NSGA-II [16] and MOEA/D [17]. Among them, a number of MOEAs outfitted with local search [18–20] are specifically designed for MOCOPs, which iteratively perform a search in the solution space to find approximate Pareto-optimal solutions.

Neural heuristics for COPs. As the early attempts, neural *construction* heuristics based on deep models [21–23] were proposed to learn to directly construct solutions for (single-objective) COPs. Among them, the classic *Attention Model* (AM) [24] was proposed inspired by *Transformer* architecture [25], and known as a milestone for solving vehicle routing problems. A number of subsequent works [26–30] were developed based on AM, including the popular policy optimization with multiple optima (POMO) [26], which leveraged the solution symmetries and significantly enhanced the performance over AM. Besides, graph neural networks were also employed to learn graphic embeddings for solving COPs [31–34]. Different from the above neural construction heuristics, neural *improvement* heuristics [35–39] were proposed to assist iteration-based methods to refine an initial solution.

Neural heuristics for MOCOPs. MOCOPs seek for a set of trade-off solutions, which could be obtained by solving a series of single-objective subproblems. According to the number of trained models, neural heuristics could be roughly classified into multi-model and single-model ones. The former adopts multiple networks to tackle respective subproblems, where the set of networks could also be collaboratively trained with a neighborhood-based parameter-transfer strategy [8, 9], or evolved with MOEA [10, 11]. Differently, the latter usually exploits a unified network to tackle all the subproblems, which is more flexible and practical. For example, taking the weight (preference) vector of a subproblem as the input, PMOCO introduces a hypernetwork to learn the decoder parameters [12]. However, the hypernetwork may cause extra complexity to the original deep reinforcement learning model, rendering it less effective to learn more accurate mapping from the preference to the optimal solution of the corresponding subproblem. By contrast, MDRL [13] leverages the meta-learning to train a deep reinforcement learning model that could be fine-tuned for various subproblems. However, its learning efficiency is far from optimal, especially given the slow and unstable training due to sequential learning of randomly sampled tasks, and inefficient fine-tuning due to redundant updates.

3 Preliminary

3.1 MOCOP

In general, an MOCOP could be defined as

$$\min_{x \in \mathcal{X}} \mathbf{f}(x) = (f_1(x), f_2(x), \dots, f_M(x)), \quad (1)$$

where $\mathbf{f}(x)$ is an M -objective vector, \mathcal{X} is a discrete decision space, and the objectives might be conflicted. The Pareto-optimal solutions concerned by decision makers are defined as follows.

Definition 1 (Pareto dominance). Let $u, v \in \mathcal{X}$, u is said to dominate v ($u \prec v$) if and only if $f_i(u) \leq f_i(v), \forall i \in \{1, \dots, M\}$ and $\exists j \in \{1, \dots, M\}, f_j(u) < f_j(v)$.

Definition 2 (Pareto optimality). A solution $x^* \in \mathcal{X}$ is Pareto-optimal if x^* is not dominated by any other solution in \mathcal{X} , i.e., $\nexists x' \in \mathcal{X}$ such that $x' \prec x^*$. The set of all Pareto optimal solutions $\mathcal{P} = \{x^* \in \mathcal{X} \mid \nexists x' \in \mathcal{X} : x' \prec x^*\}$ is called the Pareto set. The *image* of Pareto set in the objective space, i.e., $\mathcal{PF} = \{\mathbf{f}(x) \mid x \in \mathcal{P}\}$ is called the Pareto front.

3.2 Decomposition

The decomposition strategy is always applied for solving MOCOPs [17], since it is straightforward yet effective to guide the search towards prescribed directions. Specifically, MOCOPs can be decomposed into N scalarized single-objective subproblem $g(x|\boldsymbol{\lambda})$, where the weight vector $\boldsymbol{\lambda} \in \mathcal{R}^M$ satisfies $\lambda_m \geq 0$ and $\sum_{m=1}^M \lambda_m = 1$. Then, \mathcal{PF} is approximated by solving the subproblems systematically.

Regarding the decomposition strategy, the weighted sum (WS) and Tchebycheff [40] are commonly used. As a simple representative, WS considers the linear combination of M objectives as follows,

$$\min_{x \in \mathcal{X}} g_{\text{ws}}(x|\boldsymbol{\lambda}) = \sum_{m=1}^M \lambda_m f_m(x). \quad (2)$$

Given the weight vector λ , the corresponding single-objective subproblem could be cast as a sequential decision problem and solved through DRL. In particular, a solution is represented as a sequence $\pi = \{\pi_1, \dots, \pi_T\}$ with length T , and a stochastic policy for yielding solution π from instance s is calculated as $P(\pi|s) = \prod_{t=1}^T P_{\theta}(\pi_t|\pi_{1:t-1}, s)$, where $P_{\theta}(\pi_t|\pi_{1:t-1}, s)$ is the node selection probability parameterized by θ (the subscript of θ_i, π_i, P_i for subproblem i is omitted for readability).

3.3 Meta-learning

Meta-learning [14, 41] aims to train a model that can learn to tackle various tasks (i.e., subproblems) via fine-tuning. Naturally, a meta-model can be trained to quickly adapt to N corresponding subproblems given N weight vectors. Generally, meta-learning comprises three processes for MOCOPs. In the meta-learning process, a model θ is trained by repeatedly sampling tasks from the whole task space. In the fine-tuning process, N submodels $\theta_1, \dots, \theta_N$ based on the given weight vectors are fine-tuned from θ using fine-tuning instances (also called query set). In the inference process, N subproblems of an instance are solved using the submodels to approximate the Pareto set.

4 Methodology

Our *efficient meta neural heuristic* (EMNH) includes an accelerated and stabilized meta-learning process, a hierarchical fine-tuning process, and an inference process, as depicted in Figure 1. EMNH is generic, where we adopt the first-order gradient-based *Reptile* [14] as the backbone of meta-learning, and employ the popular neural solver POMO [26] as the base model. In our meta-learning process, the model θ is trained with T_m meta-iterations, where a multi-task model is exploited to speed up the training and a scaled symmetric sampling method is designed to stabilize the training. In our fine-tuning process, the meta-model θ is hierarchically and quickly fine-tuned with a few gradient updates to solve the corresponding subproblems. The details for each design are presented below.

4.1 Accelerated training

In the accelerated training process, as shown in Algorithm 1, the meta-model θ is trained with T_m meta-iterations. In each iteration, the model needs to be optimized by learning \tilde{N} sampled tasks (subproblems) with T_u -step gradient updates. Inspired by the feature reuse of meta-learning [15], it is unnecessary and expensive to serially update the meta-model through the \tilde{N} respective submodels of the same architecture. On the other hand, it is reasonable to assume that only the *head* θ_{head} is specified for a task, while the *body* θ_{body} can be reused for all tasks. To realize such a lightweight meta-learning scheme, we further introduce a multi-task model $\tilde{\theta}$ which is composed of a shared body and \tilde{N} respective heads to learn the \tilde{N} tasks in parallel.

Specifically, since our EMNH adopts the encoder-decoder structured POMO as the base model, as depicted in Figure 2(a), 1) θ_{head} could be defined as the decoder head, i.e., $W^K \in \mathcal{R}^{d \times d}$ in the last single-head attention layer, where d is empirically set to 128 [26]; 2) θ_{body} could be composed of the whole encoder θ_{en} and the decoder body $\theta_{\text{de-body}}$. For a problem instance with n nodes, the node embeddings $\mathbf{h}_1, \dots, \mathbf{h}_n \in \mathcal{R}^d$ are computed by θ_{en} at the encoding step. At each decoding step, the *query* $\mathbf{q}_c \in \mathcal{R}^d$ is first computed by $\theta_{\text{de-body}}$ using the node embeddings and problem-specific *context* embedding \mathbf{h}_c . Then, the last single-head attention layer computes the probability of node selection $P_{\theta}(\pi|s)$ using \mathbf{q}_c and the *key* $\mathbf{k}_1, \dots, \mathbf{k}_n \in \mathcal{R}^d$, where $\mathbf{k}_{i'}$ for node i' is computed by θ_{head} , i.e., $\mathbf{k}_{i'} = W^K \mathbf{h}_{i'}$. More details are presented in Appendix A.

As demonstrated in Figure 2(b), for natural accommodation, the multi-task model $\tilde{\theta}$ consists of $\tilde{\theta}_{\text{body}}$ and $\tilde{\theta}_{\text{head}_1}, \dots, \tilde{\theta}_{\text{head}_{\tilde{N}}}$, where $\tilde{\theta}_{\text{body}}$ and $\tilde{\theta}_{\text{head}_i}$ have the same architecture as θ_{body} and θ_{head} , respectively. Note that $\tilde{\theta}_{\text{head}_i}$ is individually updated for subproblem i , while $\tilde{\theta}_{\text{body}}$ is shared for \tilde{N} tasks. Concretely, since the subproblems are captured by different weight vectors but the same (or similar) node features, the shared node embeddings are computed by $\tilde{\theta}_{\text{en}}$. At each decoding step, for subproblem i , the *query* $\mathbf{q}_{c,i}$ is first computed by $\tilde{\theta}_{\text{de-body}}$ using the node embeddings and *context* embedding $\mathbf{h}_{c,i}$. Then, the *key* is computed as $\mathbf{k}_{i',i} = W_i^K \mathbf{h}_{i'}$ for node i' . Finally, the single-head attention layer computes the probability of node selection $P_{\tilde{\theta}_i}(\pi|s)$ for subproblem i using $\mathbf{q}_{c,i}$ and $\mathbf{k}_{1,i}, \dots, \mathbf{k}_{n,i}$. In practice, W_1^K, \dots, W_N^K are concatenated to tackle the \tilde{N} tasks in parallel.

Algorithm 1 Accelerated training process

Input: initialized meta-model θ , number of symmetric sampled weight vectors \tilde{N} , initialized multi-task model $\tilde{\theta}$, initial meta-learning rate ϵ_0 , number of meta-iterations T_m , number of update steps of the multi-task model T_u , batch size B , problem size n

- 1: $\epsilon \leftarrow \epsilon_0$
 - 2: **for** $t_m = 1$ to T_m **do**
 - 3: λ_i is obtained by the scaled symmetric sampling method, $\forall i \in \{1, \dots, \tilde{N}\}$
 - 4: $\tilde{\theta}_{\text{body}} \leftarrow \theta_{\text{body}}$
 - 5: $\tilde{\theta}_{\text{head}_i} \leftarrow \theta_{\text{head}}, \forall i$
 - 6: **for** $t_u = 1$ to T_u **do**
 - 7: $s_j \sim \text{SampleInstance}(\mathcal{S}), \forall j \in \{1, \dots, B\}$
 - 8: $\{\pi^k | s_j, \lambda_i\} \sim \text{SampleRollout}(P_{\tilde{\theta}_i}(\cdot | s_j)), \forall k \in \{1, \dots, n\}, \forall i, \forall j$
 - 9: $b_{ij} \leftarrow \frac{1}{n} \sum_{k=1}^n g(\pi^k | s_j, \lambda_i)$
 - 10: $\nabla \mathcal{L}(\tilde{\theta}) \leftarrow \frac{1}{\tilde{N}Bn} \sum_{i=1}^{\tilde{N}} \sum_{j=1}^B \sum_{k=1}^n [(g_{ij}^k - b_{ij}) \nabla \log P_{\tilde{\theta}_i}(\pi^k | s_j)]$
 - 11: $\tilde{\theta} \leftarrow \text{Adam}(\tilde{\theta}, \nabla \mathcal{L}(\tilde{\theta}))$
 - 12: **end for**
 - 13: $\theta_{\text{body}} \leftarrow \tilde{\theta}_{\text{body}}$
 - 14: $\theta_{\text{head}} \leftarrow \theta_{\text{head}} + \epsilon (\frac{1}{\tilde{N}} \sum_{i=1}^{\tilde{N}} \tilde{\theta}_{\text{head}_i} - \theta_{\text{head}})$
 - 15: $\epsilon \leftarrow \epsilon - \epsilon_0 / T_m$
 - 16: **end for**
- Output:** The trained meta-model θ
-




Figure 2: Model architectures. (a) Base model (POMO). (b) Multi-task model.

We extend REINFORCE [42] to train the multi-task model $\tilde{\theta}$, including $\tilde{\theta}_{\text{body}}$ for all tasks and $\tilde{\theta}_{\text{head}_i}$ for task i . Specifically, $\tilde{\theta}$ is optimized by gradient descent with the multi-task loss $\nabla \mathcal{L} = \frac{1}{\tilde{N}} \sum_{i=1}^{\tilde{N}} \nabla \mathcal{L}_i$, where $\mathcal{L}_i(\tilde{\theta} | s)$ is the i -th task loss and its gradient is estimated as follows,

$$\nabla \mathcal{L}_i(\tilde{\theta} | s) = \mathbf{E}_{P_{\tilde{\theta}_i}(\pi | s)} [(g(\pi | s, \lambda_i) - b_i(s)) \nabla \log P_{\tilde{\theta}_i}(\pi | s)], \quad (3)$$

where $b_i(s) = \frac{1}{n} \sum_{k=1}^n g(\pi^k | s, \lambda_i)$ is a baseline to reduce the gradient variance; π^1, \dots, π^n are solutions produced by POMO with n different starting nodes. In practice, $\nabla \mathcal{L}$ is approximated using Monte Carlo sampling, and computed by a batch with B instances as follows,

$$\nabla \mathcal{L}(\tilde{\theta}) \approx \frac{1}{\tilde{N}Bn} \sum_{i=1}^{\tilde{N}} \sum_{j=1}^B \sum_{k=1}^n [(g_{ij}^k - b_{ij}) \nabla \log P_{\tilde{\theta}_i}(\pi^k | s_j)], \quad (4)$$

where $g_{ij}^k = g(\pi^k | s_j, \lambda_i)$ and $b_{ij} = b_i(s_j)$. $\tilde{\theta}$ is then optimized with T_u steps using Adam [43].




Figure 3: (a) Random sampling. (b) Symmetric sampling. (c) Scaled symmetric sampling.

4.2 Stabilized training

In each meta-iteration, the deviation of \tilde{N} (usually a small number) randomly sampled weight vectors may cause fluctuations to the parameter update of the meta-model, thereby leading to unstable training performance. For example, as shown in Figure 3(a), the biased samples of weight vectors, i.e., λ_1 and λ_2 , make the meta-model tend to favor the second objective more in this iteration.

To address this issue, we first design a symmetric sampling method for the case of balanced objective domain, as illustrated in Figure 3(b). Once a weight vector $\tilde{\lambda}_1$ is sampled, its $M-1$ rotational symmetric weight vectors $\tilde{\lambda}_2, \dots, \tilde{\lambda}_M$ are calculated and used to reduce the bias as follows,

$$\tilde{\lambda}_{i,m} = \begin{cases} \tilde{\lambda}_{i-1,M}, & m = 1 \\ \tilde{\lambda}_{i-1,m-1}, & \text{otherwise} \end{cases} \quad (5)$$

For the case of imbalanced objective domains, we further design a scaled symmetric sampling method. In particular, the m -th weight λ_m is scaled as λ_m/f_m^* , where f_m^* is the ideal objective value but difficult to be calculated beforehand. Thus, f_m^* is replaced with f'_m , which is dynamically estimated using the meta-model on a validation dataset with $\lambda = (1/M, \dots, 1/M)$ during training.

Concretely, for a sampled $\tilde{\lambda}_1$, it is first multiplied by f' to scale objectives into the balanced domains. Then, its $M-1$ rotational symmetric weight vectors are obtained by Eq. (5). Finally, they are divided by f' to scale back to the original domain to attain the genuine rotational symmetric weight vectors. An example for $\tilde{\lambda}_1 = (0.2, 0.8)$ and its scaled symmetric $\tilde{\lambda}_2 = (0.43, 0.57)$ are demonstrated in Figure 3(c). In summary, the scaled symmetric $\tilde{\lambda}_2, \dots, \tilde{\lambda}_M$ for a given $\tilde{\lambda}_1$ are calculated as follows,

$$\lambda_{i,m}^s = \begin{cases} \tilde{\lambda}_{i-1,M} \times f'_M/f'_m, & m = 1 \\ \tilde{\lambda}_{i-1,m-1} \times f'_{m-1}/f'_m, & \text{otherwise} \end{cases} \quad (6)$$

Then, the vectors are normalized as $\tilde{\lambda}_{i,m} = \frac{\lambda_{i,m}^s}{\sum_{m=1}^M \lambda_{i,m}^s}$. The pseudo code of the scaled symmetric sampling method for the weight vectors is presented in Appendix B.

4.3 Hierarchical fine-tuning

Once the meta-model is trained, it can be fine-tuned with a number of gradient updates to derive customized submodels for given weight vectors. However, it still might be expensive for MOCOPs since numerous submodels need to be individually fine-tuned from the meta-model. From a systematic perspective, the early coarse-tuning processes in which the parameters of neighboring submodels might be close, could be merged. In this sense, we propose a hierarchical method to efficiently fine-tune the numerous subproblems from the lowest to the highest level, as illustrated in Figure 1.

We suggest an L -level and a -section hierarchy to gradually produce more specific submodels. Particularly, all $N^{(l)}$ weight vectors in level l are defined as the centers of their corresponding subspaces that uniformly divide the whole weight vector space based on the Das and Dennis [44] method. $N^{(l)}$ submodels are then fine-tuned to derive $N^{(l+1)} = aN^{(l)}$ submodels with $K^{(l+1)}$ fine-tuning steps in level $l+1$. Finally, $N^{(L-1)}$ submodels are fine-tuned to derive $N^{(L)} = N$ submodels, which are used to solve the N given subproblems. More details of this fine-tuning are presented in Appendix C.

For the L -level and a -section hierarchy with $a^L = N$ and $K^{(1)} = \dots = K^{(L)} = K$, the total fine-tuning step is $T_f^h = \sum_{l=1}^L K a^l = K a(a^L - 1)/(a - 1)$. For the vanilla process, each submodel needs KL fine-tuning steps to achieve comparable performance, thus $T_f^v = KLN = KLa^L$ steps in total. Hence, the hierarchical fine-tuning only needs $T_f^h/T_f^v \approx 1/L$ of the vanilla fine-tuning steps.

5 Experimental results

5.1 Experimental settings

We conduct computational experiments to evaluate the proposed method on the multi-objective traveling salesman problem (MOTSP), multi-objective capacitated vehicle routing problem (MOCVRP), and multi-objective knapsack problem (MOKP). Following the convention in [26, 12], we consider the instances of different sizes $n=20/50/100$ for MOTSP/MOCVRP and $n=50/100/200$ for MOKP. All experiments are run on a PC with an Intel Xeon 4216 CPU and an RTX 3090 GPU.

Problems. Four classes of MOTSP [45] are considered, including bi-objective TSP type 1&2 (Bi-TSP-1, Bi-TSP-2) and tri-objective TSP type 1&2 (Tri-TSP-1, Tri-TSP-2). For the M -objective TSP type 1, node i has M 2D coordinates $\{\mathbf{x}_i^1, \dots, \mathbf{x}_i^M\}$, where the m -th cost between node i and j is defined as the Euclidean distance $c_{ij}^m = \|\mathbf{x}_i^m - \mathbf{x}_j^m\|_2$. For the M -objective TSP type 2, a node contains $M-1$ 2D coordinates together with the altitude, so the M -th cost is defined as the altitude variance. For the bi-objective CVRP (Bi-CVRP) [46], two conflicting objectives with imbalanced domains, i.e., the total traveling distance and the traveling distance of the longest route (also called makespan), are considered. For the bi-objective KP (Bi-KP) [47], item i has a single weight and two distinct values. More details of these MOCOPs are presented in Appendix D.

Hyper-parameters. The meta-learning rate ϵ is linearly annealed to 0 from $\epsilon_0 = 1$ initially. A constant learning rate of the Adam optimizer is set to 10^{-4} . We set $B = 64$, $T_m = 3000$, $T_u = 100$, and $\tilde{N} = M$. The N weight vectors in \mathcal{PF} construction are produced according to [44], where $N = C_{H+M-1}^{M-1}$. H is set to 100 ($N = 101$) and 13 ($N = 105$) for $M = 2$ and $M = 3$, respectively. WS scalarization is adopted. In the hierarchy of fine-tuning, the whole weight space is uniformly divided by the Das and Dennis [44] method with $H^{(l)} = 2^l$, i.e., level l has 2^l or 4^l subspaces and $L = 7$ or $L = 4$ for $M = 2$ or $M = 3$. The number of fine-tuning steps at each level is $K^{(l)} = K = 20$ or $K^{(l)} = K = 25$ for $M = 2$ or $M = 3$.

Baselines. Two kinds of strong baselines are introduced, and all of them adopt WS (weighted sum) scalarization for fair comparisons. (1) The state-of-the-art neural heuristics, including **MDRL** [13], **PMOCO** [12], and DRL-based multiobjective optimization algorithm (**DRL-MOA**) [8]. Same as EMNH, all these neural heuristics adopt POMO as the base model for single-objective subproblems. (2) The state-of-the-art traditional heuristics, including **PPLS/D-C** [20], **WS-LKH**, and **WS-DP**. PPLS/D-C is a local-search-based MOEA proposed for MOCOPs, where a 2-opt heuristic is used for MOTSP and MOCVRP, and a greedy transformation heuristic [48] is used for MOKP, running with 200 iterations for all the three problems. WS-LKH and WS-DP are based on the state-of-the-art heuristics for decomposed single-objective subproblems, i.e., LKH [49, 50] for MOTSP and dynamic programming (DP) for MOKP, respectively. Our code is publicly available².

Metrics. Our EMNH is evaluated in terms of solution quality and learning efficiency. Solution quality is mainly measured by hypervolume (HV) [51], where a higher HV means a better solution set (the definition is presented in Appendix E). Learning efficiency mainly refers to the training and fine-tuning efficiency, which are measured by training time for the same amount of training instances and solution quality with the same total fine-tuning steps (see Appendix F), respectively.

²<https://github.com/bill-cjb/EMNH>

Table 1: Results on 200 random instances for MOCOPs with balanced objective domains.

Method	Bi-TSP-1 ($n=20$)			Bi-TSP-1 ($n=50$)			Bi-TSP-1 ($n=100$)		
	HV \uparrow	Gap \downarrow	Time	HV \uparrow	Gap \downarrow	Time	HV \uparrow	Gap \downarrow	Time
WS-LKH	<u>0.6270</u>	<u>0.02%</u>	10m	0.6415	-0.11%	1.8h	0.7090	-0.95%	6.0h
PPLS/D-C	<u>0.6256</u>	0.24%	26m	0.6282	1.97%	2.8h	0.6844	2.55%	11h
DRL-MOA	0.6257	0.22%	6s	0.6360	0.75%	9s	0.6970	0.75%	21s
PMOCO	0.6259	0.19%	6s	0.6351	0.89%	10s	0.6957	0.94%	19s
MDRL	0.6271	0.00%	5s	0.6364	0.69%	9s	0.6969	0.77%	17s
EMNH	0.6271	0.00%	5s	0.6364	0.69%	9s	0.6969	0.77%	16s
PMOCO-Aug	<u>0.6270</u>	<u>0.02%</u>	45s	0.6395	0.20%	2.3m	0.7016	0.10%	15m
MDRL-Aug	0.6271	0.00%	33s	<u>0.6408</u>	<u>0.00%</u>	1.7m	<u>0.7022</u>	<u>0.01%</u>	14m
EMNH-Aug	0.6271	0.00%	33s	<u>0.6408</u>	<u>0.00%</u>	1.7m	<u>0.7023</u>	<u>0.00%</u>	14m
Method	Bi-KP ($n=50$)			Bi-KP ($n=100$)			Bi-KP ($n=200$)		
	HV \uparrow	Gap \downarrow	Time	HV \uparrow	Gap \downarrow	Time	HV \uparrow	Gap \downarrow	Time
WS-DP	0.3561	0.00%	22m	0.4532	<u>0.07%</u>	2.0h	<u>0.3601</u>	0.06%	5.8h
PPLS/D-C	0.3528	0.93%	18m	0.4480	1.21%	47m	0.3541	1.72%	1.5h
DRL-MOA	<u>0.3559</u>	<u>0.06%</u>	9s	0.4531	0.09%	18s	<u>0.3601</u>	<u>0.06%</u>	1.0m
PMOCO	0.3552	0.25%	6s	0.4523	0.26%	22s	0.3595	0.22%	1.3m
MDRL	0.3530	0.87%	6s	0.4532	<u>0.07%</u>	21s	<u>0.3601</u>	0.06%	1.2m
EMNH	0.3561	0.00%	6s	0.4535	0.00%	21s	0.3603	0.00%	1.2m
Method	Tri-TSP-1 ($n=20$)			Tri-TSP-1 ($n=50$)			Tri-TSP-1 ($n=100$)		
	HV \uparrow	Gap \downarrow	Time	HV \uparrow	Gap \downarrow	Time	HV \uparrow	Gap \downarrow	Time
WS-LKH	0.4712	0.00%	12m	0.4440	-0.50%	1.9h	0.5076	-2.07%	6.6h
PPLS/D-C	0.4698	0.30%	1.4h	0.4174	5.52%	3.9h	0.4376	12.00%	14h
DRL-MOA	<u>0.4699</u>	<u>0.28%</u>	6s	0.4303	2.60%	9s	0.4806	3.36%	19s
PMOCO	0.4693	0.40%	5s	0.4315	2.33%	8s	0.4858	2.31%	18s
MDRL	<u>0.4699</u>	<u>0.28%</u>	5s	0.4317	2.29%	9s	0.4852	2.43%	16s
EMNH	<u>0.4699</u>	<u>0.28%</u>	5s	0.4324	2.13%	9s	0.4866	2.15%	16s
PMOCO-Aug	0.4712	0.00%	3.2m	0.4409	0.20%	28m	0.4956	0.34%	1.7h
MDRL-Aug	0.4712	0.00%	2.6m	0.4408	0.23%	25m	0.4958	0.30%	1.7h
EMNH-Aug	0.4712	0.00%	2.6m	<u>0.4418</u>	<u>0.00%</u>	25m	<u>0.4973</u>	<u>0.00%</u>	1.7h

5.2 Solution quality

The results for MOTSP, MOCVRP, and MOKP are recorded in Tables 1 and 2, including the average HV, gap, and total running time for solving 200 random test instances. To further show the significant differences between the results, a Wilcoxon rank-sum test at 1% significance level is adopted. The best result and the one without statistical significance to it are highlighted as **bold**. The second-best result and the one without statistical significance to it are highlighted as underline. The method with “-Aug” represents the inference results using instance augmentation [12] (see Appendix G).

Optimality gap. The gaps of HV with respect to EMNH-Aug are reported for all methods. According to the results in Tables 1 and 2, EMNH outperforms other neural heuristics without instance augmentation on all problems. When equipped with instance augmentation, EMNH-Aug further improves the solution, and performs superior to most of the baselines, while slightly inferior to WS-LKH on MOTSP with $n=50$ and $n=100$. However, as iteration-based methods, WS-LKH and MOEAs take quite long running time on MOCOPs. For the neural heuristics, DRL-MOA is less agile since it trains multiple fixed models for a priori weight vectors. For new weight vectors concerned by decision makers, EMNH, MDRL, and PMOCO are able to efficiently produce high-quality trade-off solutions, where EMNH achieves the smallest gap.

Imbalanced objective domains. As demonstrated in Table 2, the gaps between EMNH and PMOCO or MDRL are further increased on the problems with imbalanced objective domains, i.e., Bi-CVRP, Bi-TSP-2, and Tri-TSP-2. EMNH exhibits more competitive performance on these problems, revealing

Table 2: Results on 200 random instances for MOCOPs with imbalanced objective domains.

Method	Bi-CVRP ($n=20$)			Bi-CVRP ($n=50$)			Bi-CVRP ($n=100$)		
	HV \uparrow	Gap \downarrow	Time	HV \uparrow	Gap \downarrow	Time	HV \uparrow	Gap \downarrow	Time
PPLS/D-C	0.4287	0.35%	1.6h	0.4007	2.41%	9.7h	0.3946	3.26%	38h
DRL-MOA	0.4287	0.35%	10s	0.4076	0.73%	12s	0.4055	0.59%	33s
PMOCO	0.4267	0.81%	7s	0.4036	1.70%	12s	0.3913	4.07%	32s
MDRL	0.4291	0.26%	8s	0.4082	0.58%	13s	0.4056	0.56%	32s
EMNH	<u>0.4299</u>	<u>0.07%</u>	8s	<u>0.4098</u>	<u>0.19%</u>	12s	<u>0.4072</u>	<u>0.17%</u>	32s
PMOCO-Aug	0.4294	0.19%	13s	0.4080	0.63%	36s	0.3969	2.70%	2.7m
MDRL-Aug	0.4294	0.19%	11s	0.4092	0.34%	36s	<u>0.4072</u>	<u>0.17%</u>	2.8m
EMNH-Aug	0.4302	0.00%	11s	0.4106	0.00%	35s	0.4079	0.00%	2.8m
Method	Bi-TSP-2 ($n=20$)			Bi-TSP-2 ($n=50$)			Bi-TSP-2 ($n=100$)		
	HV \uparrow	Gap \downarrow	Time	HV \uparrow	Gap \downarrow	Time	HV \uparrow	Gap \downarrow	Time
WS-LKH	0.6660	0.13%	11m	0.7390	-0.07%	1.8h	0.8055	-0.74%	6.1h
PPLS/D-C	<u>0.6662</u>	<u>0.10%</u>	27m	0.7300	1.15%	3.3h	0.7859	1.71%	10h
DRL-MOA	0.6657	0.18%	6s	0.7359	0.35%	8s	0.7965	0.39%	18s
PMOCO	0.6590	1.18%	5s	0.7347	0.51%	8s	0.7944	0.65%	17s
MDRL	0.6669	0.00%	5s	0.7361	0.32%	9s	0.7965	0.39%	15s
EMNH	0.6669	0.00%	5s	0.7361	0.32%	9s	0.7965	0.39%	15s
PMOCO-Aug	0.6653	0.24%	13s	0.7375	0.14%	33s	0.7988	0.10%	3.5m
MDRL-Aug	0.6669	0.00%	8s	<u>0.7385</u>	<u>0.00%</u>	22s	<u>0.7996</u>	<u>0.00%</u>	2.8m
EMNH-Aug	0.6669	0.00%	8s	<u>0.7385</u>	<u>0.00%</u>	22s	<u>0.7996</u>	<u>0.00%</u>	2.8m
Method	Tri-TSP-2 ($n=20$)			Tri-TSP-2 ($n=50$)			Tri-TSP-2 ($n=100$)		
	HV \uparrow	Gap \downarrow	Time	HV \uparrow	Gap \downarrow	Time	HV \uparrow	Gap \downarrow	Time
WS-LKH	0.5035	0.00%	13m	0.5305	-0.49%	2.0h	0.5996	-1.70%	6.6h
PPLS/D-C	0.5035	0.00%	1.4h	0.5045	4.42%	4.1h	0.5306	10.01%	15h
DRL-MOA	0.5019	0.32%	6s	0.5101	3.37%	8s	0.5488	6.92%	19s
PMOCO	0.5020	0.30%	5s	0.5176	1.95%	8s	0.5777	2.02%	18s
MDRL	<u>0.5024</u>	<u>0.22%</u>	5s	0.5183	1.82%	9s	0.5806	1.53%	16s
EMNH	<u>0.5024</u>	<u>0.22%</u>	5s	0.5205	1.40%	9s	0.5813	1.41%	16s
PMOCO-Aug	0.5035	0.00%	55s	0.5258	0.40%	6.1m	0.5862	0.58%	32m
MDRL-Aug	0.5035	0.00%	37s	0.5267	0.23%	4.2m	0.5886	0.17%	30m
EMNH-Aug	0.5035	0.00%	37s	<u>0.5279</u>	<u>0.00%</u>	4.2m	<u>0.5896</u>	<u>0.00%</u>	30m

the effectiveness of the proposed scaled symmetric sampling method in tackling the imbalance of objective domains. We further equip PMOCO with the same sampling method, which actually also improved the solutions for those problems, but still performed inferior to EMNH (see Appendix H).

Generalization ability. We test the generalization ability of the model on 200 larger-scale random instances ($n=150/200$) and 3 commonly used MOTSP benchmark instances (KroAB100/150/200) in TSPLIB [52]. The zero-shot generalization performance of the model trained and fine-tuned both on the instances with $n=100$ is reported in Appendix I. The results suggest that EMNH has a superior generalization ability compared with MOEAs and other neural heuristics for larger problem sizes.

Hyper-parameter study. The results in Appendix J showed that the number of sampled weight vectors $\tilde{N} = M$ is a more desirable setting and WS is a simple yet effective scalarization method.

5.3 Learning efficiency

As verified above, EMNH is able to produce superior solutions to the state-of-the-art neural heuristics, especially demonstrating a significant advantage over PMOCO in terms of solution quality. Then we further show that EMNH also has favorable learning efficiency against MDRL.

Training efficiency. EMNH, MDRL, and PMOCO only train one model to tackle all subproblems, where EMNH and MDRL use the same amount of training instances, and PMOCO requires a few




Figure 4: Learning efficiency. (a) Training efficiency. (b) Training stability. (c) Fine-tuning efficiency.

more training instances according to its original setting. DRL-MOA needs to train multiple submodels with more training instances. More details are presented in Appendix F. Figure 4(a) displays the speed-up ratio, i.e., the ratio of the training time of MDRL to EMNH. Due to the multi-task based training for (partial) architecture reuse, training time of EMNH is only about $1/\tilde{N}$ of that of MDRL.

Training stability. Figure 4(b) shows the stability of the training process, where EMNH-S and EMNH-R refer to EMNH with the symmetric sampling and the random sampling, respectively. EMNH achieves the stablest and best training performance, compared to EMNH-S, EMNH-R, and MDRL, as the proposed sampling method in EMNH considers symmetric weight vectors and (imbalanced) objective domains. More details are presented in Appendix F.

Fine-tuning efficiency. Figure 4(c) compares the hierarchical fine-tuning in EMNH with the vanilla fine-tuning method in MDRL, where the HV of EMNH with K fine-tuning steps at each level and the HV of MDRL with approximately equal total fine-tuning steps are presented. The results demonstrate that EMNH attains higher fine-tuning efficiency than MDRL, i.e., larger gaps with smaller fine-tuning steps. Furthermore, PMOCO is also equipped with the hierarchical fine-tuning (for both versions) for a fair comparison. It is worth noting that our fine-tuning process is performed individually for each weight vector on dedicated fine-tuning instances. As can be seen, PMOCO could hardly get improved by fine-tuning for zero-shot inference on test instances, since it has already converged for the corresponding subproblems. This means that EMNH has favorable potential to derive more desirable submodels to tackle specific tasks. Notably, EMNH with a few fine-tuning steps (e.g., $K = 5$) outperforms PMOCO in most cases, as demonstrated in Appendix F.

6 Conclusion

This paper proposes an efficient meta neural heuristic (EMNH) for MOCOPs. Specifically, EMNH reduces training time, stabilizes the training process, and curtails total fine-tuning steps via (partial) architecture-shared multi-task learning, scaled symmetric sampling (of weight vectors), and hierarchical fine-tuning, respectively. The experimental results on MOTSP, MOCVRP, and MOKP verified the superiority of EMNH in learning efficiency and solution quality. A limitation is that our EMNH, as the same as other neural heuristics, can not guarantee obtaining the exact Pareto front. In the future, we will (1) extend EMNH to other MOCOPs with complex constraints; (2) investigate other advanced meta-learning algorithms and neural solvers as the base model to further improve the performance.

Acknowledgments and disclosure of funding

This work is supported by the National Natural Science Foundation of China (62072483), and the Guangdong Basic and Applied Basic Research Foundation (2022A1515011690, 2021A1515012298).

References

- [1] Matthias Ehrgott and Xavier Gandibleux. A survey and annotated bibliography of multiobjective combinatorial optimization. *OR-spektrum*, 22(4):425–460, 2000.
- [2] 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.

- [3] Arne Herzel, Stefan Ruzika, and Clemens Thielen. Approximation methods for multiobjective optimization problems: A survey. *INFORMS Journal on Computing*, 33(4):1284–1299, 2021.
- [4] Azalia Mirhoseini, Anna Goldie, Mustafa Yazgan, Joe Wenjie Jiang, Ebrahim Songhori, Shen Wang, Young-Joon Lee, Eric Johnson, Omkar Pathak, Azade Nazi, Jiwoo Pak, Andy Tong, Kavya Srinivasa, William Hang, Emre Tuncer, Quoc V. Le, James Laudon, Richard Ho, Roger Carpenter, and Jeff Dean. A graph placement methodology for fast chip design. *Nature*, 594:207–212, 2021.
- [5] 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.
- [6] Nina Mazyavkina, Sergey Sviridov, Sergei Ivanov, and Evgeny Burnaev. Reinforcement learning for combinatorial optimization: A survey. *Computers & Operations Research*, 134:105400, 2021.
- [7] Qi Wang and Chunlei Tang. Deep reinforcement learning for transportation network combinatorial optimization: A survey. *Knowledge-Based Systems*, 233:107526, 2021.
- [8] Kaiwen Li, Tao Zhang, and Rui Wang. Deep reinforcement learning for multiobjective optimization. *IEEE Transactions on Cybernetics*, 51(6):3103–3114, 2021.
- [9] Hong Wu, Jiahai Wang, and Zizhen Zhang. MODRL/D-AM: Multiobjective deep reinforcement learning algorithm using decomposition and attention model for multiobjective optimization. In *ISICA 2019*, volume CCIS 1205, pages 575–589, 2020.
- [10] Yinan Shao, Jerry Chun-Wei Lin, Gautam Srivastava, Dongdong Guo, Hongchun Zhang, Hu Yi, and Alireza Jolfaei. Multi-objective neural evolutionary algorithm for combinatorial optimization problems. *IEEE Transactions on Neural Networks and Learning Systems*, 2021.
- [11] Yongxin Zhang, Jiahai Wang, Zizhen Zhang, and Yalan Zhou. MODRL/D-EL: Multiobjective deep reinforcement learning with evolutionary learning for multiobjective optimization. In *International Joint Conference on Neural Networks*, 2021.
- [12] Xi Lin, Zhiyuan Yang, and Qingfu Zhang. Pareto set learning for neural multi-objective combinatorial optimization. In *International Conference on Learning Representations*, 2022.
- [13] 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, 2023.
- [14] Alex Nichol, Joshua Achiam, and John Schulman. On first-order meta-learning algorithms, 2018.
- [15] Aniruddh Raghu, Maithra Raghu, Samy Bengio, and Oriol Vinyals. Rapid learning or feature reuse? towards understanding the effectiveness of maml. In *International Conference on Learning Representations*, 2020.
- [16] Kalyanmoy Deb, Amrit Pratap, Sameer Agarwal, and TAMT Meyarivan. A fast and elitist multiobjective genetic algorithm: NSGA-II. *IEEE Transactions on Evolutionary Computation*, 6(2):182–197, 2002.
- [17] Qingfu Zhang and Hui Li. MOEA/D: A multiobjective evolutionary algorithm based on decomposition. *IEEE Transactions on Evolutionary Computation*, 11(6):712–731, 2007.
- [18] Andrzej Jaskiewicz. Genetic local search for multi-objective combinatorial optimization. *European Journal of Operational Research*, 137(1):50–71, 2002.
- [19] Jialong Shi, Qingfu Zhang, and Jianyong Sun. PPLS/D: Parallel pareto local search based on decomposition. *IEEE transactions on Cybernetics*, 50(3):1060–1071, 2020.
- [20] Jialong Shi, Jianyong Sun, Qingfu Zhang, Haotian Zhang, and Ye Fan. Improving pareto local search using cooperative parallelism strategies for multiobjective combinatorial optimization. *IEEE Transactions on Cybernetics*, 2022.
- [21] Oriol Vinyals, Meire Fortunato, and Navdeep Jaitly. Pointer networks. In *Advances in Neural Information Processing Systems*, 2015.
- [22] Irwan Bello, Hieu Pham, Quoc V Le, Mohammad Norouzi, and Samy Bengio. Neural combinatorial optimization with reinforcement learning. In *International Conference on Learning Representations*, 2017.
- [23] Mohammadreza Nazari, Afshin Oroojlooy, Martin Takáč, and Lawrence V. Snyder. Reinforcement learning for solving the vehicle routing problem. In *Advances in Neural Information Processing Systems*, 2018.

- [24] Wouter Kool, Herke Van Hoof, and Max Welling. Attention, learn to solve routing problems! In *International Conference on Learning Representations*, 2019.
- [25] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in Neural Information Processing Systems*, 2017.
- [26] Yeong-Dae Kwon, Jinho Choo, Byoungjip Kim, Iljoo Yoon, Youngjune Gwon, and Seungjai Min. POMO: Policy optimization with multiple optima for reinforcement learning. In *Advances in Neural Information Processing Systems*, 2020.
- [27] Minsu Kim, Jinkyoo Park, et al. Learning collaborative policies to solve np-hard routing problems. In *Advances in Neural Information Processing Systems*, 2021.
- [28] Yeong-Dae Kwon, Jinho Choo, Iljoo Yoon, Minah Park, Duwon Park, and Youngjune Gwon. Matrix encoding networks for neural combinatorial optimization. In *Advances in Neural Information Processing Systems*, 2021.
- [29] Liang Xin, Wen Song, Zhiguang Cao, and Jie Zhang. Multi-decoder attention model with embedding glimpse for solving vehicle routing problems. In *AAAI Conference on Artificial Intelligence*, pages 12042–12049, 2021.
- [30] Bo Peng, Jiahai Wang, and Zizhen Zhang. A deep reinforcement learning algorithm using dynamic attention model for vehicle routing problems. In *ISICA 2019*, volume CCIS 1205, pages 636–650, 2020.
- [31] Elias Khalil, Hanjun Dai, Yuyu Zhang, Bistra Dilkina, and Le Song. Learning combinatorial optimization algorithms over graphs. In *Advances in Neural Information Processing Systems*, 2017.
- [32] Sahil Manchanda, Akash Mittal, Anuj Dhawan, Sourav Medya, Sayan Ranu, and Ambuj Singh. GCOMB: Learning budget-constrained combinatorial algorithms over billion-sized graphs. In *Advances in Neural Information Processing Systems*, 2020.
- [33] Lu Duan, Yang Zhan, Haoyuan Hu, Yu Gong, Jiangwen Wei, Xiaodong Zhang, and Yinghui Xu. Efficiently solving the practical vehicle routing problem: A novel joint learning approach. In *ACM SIGKDD Conference on Knowledge Discovery & Data Mining*, pages 3054–3063, 2020.
- [34] Jinbiao Chen, Huanhuan Huang, Zizhen Zhang, and Jiahai Wang. Deep reinforcement learning with two-stage training strategy for practical electric vehicle routing problem with time windows. In *International Conference on Parallel Problem Solving from Nature*, 2022.
- [35] Xinyun Chen and Yuandong Tian. Learning to perform local rewriting for combinatorial optimization. In *Advances in Neural Information Processing Systems*, 2019.
- [36] Hao Lu, Xingwen Zhang, and Shuang Yang. A learning-based iterative method for solving vehicle routing problems. In *International Conference on Learning Representations*, 2020.
- [37] André Hottung and Kevin Tierney. Neural large neighborhood search for the capacitated vehicle routing problem. In *European Conference on Artificial Intelligence*, 2020.
- [38] 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, 2022.
- [39] Yining Ma, Jingwen Li, Zhiguang Cao, Wen Song, Le Zhang, Zhenghua Chen, and Jing Tang. Learning to iteratively solve routing problems with dual-aspect collaborative transformer. In *Advances in Neural Information Processing Systems*, 2021.
- [40] Xiaoliang Ma, Qingfu Zhang, Guangdong Tian, Junshan Yang, and Zexuan Zhu. On tchebycheff decomposition approaches for multiobjective evolutionary optimization. *IEEE Transactions on Evolutionary Computation*, 22(2):226–244, 2017.
- [41] Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adaptation of deep networks. In *International Conference on Machine Learning*, pages 1126–1135, 2017.
- [42] Ronald J Williams. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine Learning*, 8(3):229–256, 1992.
- [43] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *International Conference on Learning Representations*, 2015.

- [44] I Das and JE Dennis. Normal-boundary intersection: A new method for generating pareto-optimal points in multieriteria optimization problems. *SIAM Journal on Optimization*, 8(3):631–657, 1998.
- [45] Thibaut Lust and Jacques Teghem. The multiobjective traveling salesman problem: A survey and a new approach. In *Advances in Multi-Objective Nature Inspired Computing*, pages 119–141, 2010.
- [46] Nicolas Jozefowiez, Frédéric Semet, and El-Ghazali Talbi. Multi-objective vehicle routing problems. *European Journal of Operational Research*, 189(2):293–309, 2008.
- [47] Cristina Bazgan, Hadrien Hugot, and Daniel Vanderpooten. Solving efficiently the 0–1 multi-objective knapsack problem. *Computers & Operations Research*, 36(1):260–279, 2009.
- [48] Hisao Ishibuchi, Naoya Akedo, and Yusuke Nojima. Behavior of multiobjective evolutionary algorithms on many-objective knapsack problems. *IEEE Transactions on Evolutionary Computation*, 19(2):264–283, 2015.
- [49] Keld Helsgaun. An effective implementation of the lin–kernighan traveling salesman heuristic. *European Journal of Operational Research*, 126(1):106–130, 2000.
- [50] Renato Tinós, Keld Helsgaun, and Darrell Whitley. Efficient recombination in the lin–kernighan–helsgaun traveling salesman heuristic. In *International Conference on Parallel Problem Solving from Nature*, pages 95–107, 2018.
- [51] Charles Audet, Jean Bignon, Dominique Cartier, Sébastien Le Digabel, and Ludovic Salomon. Performance indicators in multiobjective optimization. *European Journal of Operational Research*, 292(2):397–422, 2021.
- [52] Gerhard Reinelt. TSPLIB—a traveling salesman problem library. *ORSA Journal on Computing*, 3(4): 376–384, 1991.

Efficient Meta Neural Heuristic for Multi-Objective Combinatorial Optimization (Appendix)

A Model architecture

The architecture of the base model in meta-learning is the same as POMO [26], composed of an encoder and a decoder (see Figure 2(a)). For node features $\mathbf{x}_1, \dots, \mathbf{x}_n$, the encoder first computes initial node embeddings $\mathbf{h}_1^0, \dots, \mathbf{h}_n^0 \in \mathcal{R}^d$ ($d = 128$) by a linear projection (LP). The final node embeddings $\mathbf{h}_1^{\mathcal{N}}, \dots, \mathbf{h}_n^{\mathcal{N}}$ are further computed by $\mathcal{N} = 6$ attention layers. Each attention layer is composed of a multi-head attention (MHA) with $M = 8$ heads and fully connected feed-forward (FF) sublayer. Each sublayer adds a skip-connection (ADD) and batch normalization (BN).

The decoder sequentially chooses a node according to a probability distribution produced by the node embeddings to construct a solution. The total decoding step T is determined by the specific problem. At step t in the decoding procedure, the *query* $\mathbf{q}_c \in \mathcal{R}^d$ is computed by an MHA layer using the node embeddings and problem-specific *context* embedding \mathbf{h}_c (see Appendix D). The *key* $\mathbf{k}_1, \dots, \mathbf{k}_n \in \mathcal{R}^d$ is computed by $\mathbf{k}_{i'} = W^K \mathbf{h}_{i'}$ for node i' . Then, in the final single-head attention layer, the *compatibility* \mathbf{u} is computed by \mathbf{q}_c and $\mathbf{k}_{i'}$ as follows,

$$u_{i'} = \begin{cases} -\infty, & \text{node } i' \text{ is masked} \\ C \cdot \tanh\left(\frac{\mathbf{q}_c^T \mathbf{k}_{i'}}{\sqrt{d/M}}\right), & \text{otherwise} \end{cases} \quad (7)$$

where $C = 10$ is adopted to clip the result. Finally, the probability distribution $P_{\theta}(\pi|s)$ to select node i' is computed by the softmax function,

$$P_{i'} = P_{\theta}(\pi_t = i' | \pi_{1:t-1}, s) = \frac{e^{u_{i'}}}{\sum_{j=1}^n e^{u_j}}. \quad (8)$$

For the meta-model θ , θ_{head} can be defined as W^K for the final single-head attention layer, and θ_{body} is composed of the whole encoder θ_{en} and the decoder body $\theta_{\text{de-body}}$.

In each meta-iteration, we adopt a multi-task model $\tilde{\theta}$, as shown in Figure 2(b), to learn \tilde{N} sampled tasks in parallel. The multi-task model $\tilde{\theta}$ consists of $\tilde{\theta}_{\text{body}}$ and $\tilde{\theta}_{\text{head}_1}, \dots, \tilde{\theta}_{\text{head}_{\tilde{N}}}$, where $\tilde{\theta}_{\text{body}}$ and $\tilde{\theta}_{\text{head}_i}$ have the same architecture as θ_{body} and θ_{head} , respectively. $\tilde{\theta}_{\text{head}_i}$, i.e., W_i^K , is individually updated for subproblem i , while $\tilde{\theta}_{\text{body}}$ is shared across \tilde{N} tasks. Specifically, the shared node embeddings are first computed by $\tilde{\theta}_{\text{en}}$. Then, at step t in the decoding procedure, for subproblem i , the *query* $\mathbf{q}_{c,i}$ is computed using the node embeddings and *context* embedding $\mathbf{h}_{c,i}$. The *key* $\mathbf{k}_{1,i}, \dots, \mathbf{k}_{n,i}$ is computed by $\mathbf{k}_{i',i} = W_i^K \mathbf{h}_{i'}$. The *compatibility* \mathbf{u}_i is computed as follows,

$$u_{i',i} = \begin{cases} -\infty, & \text{node } i' \text{ is masked} \\ C \cdot \tanh\left(\frac{\mathbf{q}_{c,i}^T \mathbf{k}_{i',i}}{\sqrt{d/M}}\right), & \text{otherwise} \end{cases} \quad (9)$$

Finally, the probability $P_{\tilde{\theta}_i}(\pi|s)$ for subproblem i to select node i' is computed as follows,

$$P_{i',i} = P_{\tilde{\theta}_i}(\pi_t = i' | \pi_{1:t-1}, s) = \frac{e^{u_{i',i}}}{\sum_{j=1}^n e^{u_{j,i}}}. \quad (10)$$

B Scaled symmetric sampling method

The scaled symmetric sampling method is shown in Algorithm 2. The scaled factor f'_m is first estimated (Lines 1 – 3). Then $\lfloor \tilde{N}/M \rfloor$ weight vectors are randomly sampled (Lines 5 – 6). For each of them, $M - 1$ scaled symmetric weight vectors are generated (Lines 7 – 16).

Algorithm 2 Scaled symmetric sampling method

Input: meta-model θ , problem size n , weight vector distribution Λ , number of objectives M , number of symmetric sampled weight vectors \tilde{N} , validation dataset \mathcal{V}

- 1: $\{\pi^k | \mathcal{V}_j\} \sim \mathbf{GreedyRollout}(P_\theta(\cdot | \mathcal{V}_j)), \quad \forall j \in \{1, \dots, |\mathcal{V}|\}, \forall k \in \{1, \dots, n\}$
- 2: $\pi_j \leftarrow \operatorname{argmax}_k g(\pi^k | (\mathcal{V}_j, \mathbf{1}/M))$
- 3: $f'_m \leftarrow \frac{1}{|\mathcal{V}|} \sum_{j=1}^{|\mathcal{V}|} f_m(\pi_j), \quad \forall m \in \{1, \dots, M\}$
- 4: **for** $i = 1$ to \tilde{N} **do**
- 5: **if** $i \leq \lfloor \tilde{N}/M \rfloor$ **then**
- 6: $\lambda_i \sim \mathbf{SampleWeight}(\Lambda)$
- 7: **else if** $\lfloor \tilde{N}/M \rfloor < i \leq M \times \lfloor \tilde{N}/M \rfloor$ **then**
- 8: **for** $m = 1$ to M **do**
- 9: $\lambda'_i \leftarrow \lambda_i$
- 10: **if** $m = 1$ **then**
- 11: $\lambda_{i,m} \leftarrow \lambda'_{i - \lfloor \tilde{N}/M \rfloor, M} \times f'_M / f'_m$
- 12: **else**
- 13: $\lambda_{i,m} \leftarrow \lambda'_{i - \lfloor \tilde{N}/M \rfloor, m-1} \times f'_{m-1} / f'_m$
- 14: **end if**
- 15: **end for**
- 16: $\lambda_i \leftarrow \lambda_i / \sum_{m=1}^M \lambda_{i,m}$
- 17: **else**
- 18: $\lambda_i \sim \mathbf{SampleWeight}(\Lambda)$
- 19: **end if**
- 20: **end for**

Output: $\{\lambda_1, \dots, \lambda_{\tilde{N}}\}$

C Hierarchical fine-tuning method

We consider an L -level a -section hierarchy. The whole weight space is uniformly divided into $N^{(l)}$ subspaces in level l . $N^{(l)}$ weight vectors are the centers of these subspaces. In level $l + 1$, the $N^{(l)}$ submodels are fine-tuned to derive $N^{(l+1)} = aN^{(l)}$ submodels with $K^{(l+1)}$ fine-tuning steps, which are the centers of $N^{(l+1)}$ subspaces. The j -th submodel in level $l + 1$ is fine-tuned from the i -th submodel in level l , where the j -th weight vector belongs to the i -th subspace in level l .

The uniform division of the weight space is illustrated as follows. According to the Das and Dennis [44] method, C_{H+M-1}^{M-1} vertices can be generated, where M is the number of objectives and H is a user-defined hyper-parameter. Then, the whole weight space is uniformly divided by C_{H+M-1}^{M-1} vertices. The division with $H^{(l)} = 2^l$ can be seen in Figure 5. There are 2^l subspaces for $M = 2$ and 4^l subspaces for $M = 3$ in level l .

Note that the final N weights are given beforehand. Thus, we set $N^{(L)} = N$, where $N^{(L-1)} = a^{L-1} < N$ and $a^L \geq N$, i.e., $L = 7$ for $M = 2$ and $L = 4$ for $M = 3$. The final N submodels are fine-tuned by $N^{(L-1)}$ submodels.

For the meta-model or a coarse-tuned submodel in level l ($l < L$), its fine-tuning process with a given weight vector and K fine-tuning steps is shown in Algorithm 3.

D Details of MOCOPs

D.1 MOTSP

D.1.1 Problem definition

The multi-objective traveling salesman problem (MOTSP) is defined on a complete graph with n nodes and M cost matrices. Node i has M 2D coordinates $\{\mathbf{x}_i^1, \dots, \mathbf{x}_i^M\}$, where the m -th cost between node i and j is given by the Euclidean distance $c_{ij}^m = \|\mathbf{x}_i^m - \mathbf{x}_j^m\|_2$. The goal is to find a




Figure 5: The hierarchy of fine-tuning. (a) $M = 2$. (b) $M = 3$.

Algorithm 3 Fine-tuning process in each level

Input: the meta-model or a coarse-tuned submodel θ , a given weight vector λ , fine-tuning steps K , batch size B , problem size n

1: **for** $k = 1$ to K **do**

2: $s_j \sim \text{SampleInstance}(\mathcal{S}), \quad \forall j \in \{1, \dots, B\}$

3: $\{\pi^k | s_j, \lambda\} \sim \text{SampleRollout}(P_{\theta_i}(\cdot | s_j)), \quad \forall j \in \{1, \dots, B\}, \forall k \in \{1, \dots, n\}$

4: $b_j \leftarrow \frac{1}{n} \sum_{k=1}^n g(\pi^k | s_j, \lambda)$

5: $\nabla \mathcal{L}(\tilde{\theta}) \leftarrow \frac{1}{Bn} \sum_{j=1}^B \sum_{k=1}^n [(g(\pi^k | s_j, \lambda) - b_j) \nabla \log P_{\theta}(\pi^k | s_j)]$

6: $\theta \leftarrow \text{Adam}(\tilde{\theta}, \nabla \mathcal{L}(\tilde{\theta}))$

7: **end for**

Output: The fine-tuned submodel θ

permutation π to minimize all the M costs simultaneously, as follows,

$$\min \mathbf{f}(\pi) = (f_1(\pi), f_2(\pi), \dots, f_M(\pi)), \quad (11)$$

$$\text{where } f_m(\pi) = c_{\pi_n, \pi_1}^m + \sum_{j=1}^{n-1} c_{\pi_j, \pi_{j+1}}^m. \quad (12)$$

For the M -objective TSP type 2, each node has $(M - 1)$ 2D coordinates and one number interpreted as its altitude. The single-objective counterpart, TSP, is a well-known NP-hard problem. It appears that MOTSP is even harder. Thus, its approximate Pareto optimal solutions are commonly pursued.

D.1.2 Instance

For the M -objective TSP type 1, each node has M 2D coordinates. The random instances with n nodes are sampled from uniform distribution on $[0, 1]^{2M}$. For the M -objective TSP type 2, each node has $(M - 1)$ 2D coordinates and one number. The random instances with n nodes are sampled from uniform distribution on $[0, 1]^{2M-1}$.

D.1.3 Model details

The input dimension of the M -objective TSP type 1 is $2M$ for the encoder. The input dimension of the M -objective TSP type 2 is $2M - 1$ for the encoder.

For MOTSP, POMO [26] uses n context embedding $\mathbf{h}_c^1, \dots, \mathbf{h}_c^n$ to calculate the probability of node selection in the decoder. At the decoding step t , \mathbf{h}_c^k is defined as follows,

$$\mathbf{h}_c^k = \begin{cases} [\bar{\mathbf{h}}^k; \mathbf{h}_{\pi_{t-1}}^k; \mathbf{h}_{\pi_1}^k], & t > 1 \\ \text{none}, & t = 1 \end{cases} \quad (13)$$

where $[\cdot; \cdot]$ is the concatenation and the graph embedding $\bar{\mathbf{h}}^k = \sum_{i=1}^n \mathbf{h}_i^k$. For $t = 1$, the first node is not selected by the decoder. Instead, it is defined as $\mathbf{h}_{\pi_1}^k = \mathbf{h}_k$. In the decoding procedure, the nodes already visited need to be masked.

D.2 MOCVRP

D.2.1 Problem definition

The capacitated vehicle routing problem (CVRP), a classical extension of TSP, contains n customer nodes and a depot node. Each node has a 2D coordinate. In addition, customer i has a demand d_i to be satisfied. A fleet of homogeneous vehicles with identical capacity D is initially placed at the depot. Vehicles must serve all the customers and finally return to the depot. The capacity constraints must be satisfied, i.e., the remaining capacity of vehicles for serving customer i must be no less than d_i .

For the multi-objective capacitated vehicle routing problem (MOCVRP), we consider two conflicting objectives, i.e., the total traveling distance and the traveling distance of the longest route (makespan).

D.2.2 Instance

For MOCVRP, the coordinates of the depot and customers are sampled from uniform distribution on $[0, 1]^2$. Following the previous work [26, 12], the demand d_i is randomly chosen from $\{1, \dots, 9\}$ and the capacity is set to $D = 30/40/50$ for $n = 20/50/100$. Without loss of generality, the demand and capacity are normalized as $\hat{d}_i = d_i/D$ and $\hat{D} = D/D = 1$, respectively.

D.2.3 Model details

The inputs of MOCVRP are a 2D vector of the depot and n 3D vectors of the customers for the encoder. Their embeddings with both 128 dimensions are obtained by two linear projections with two separate parameter matrices.

For MOCVRP, the context embedding in the decoder is defined as the concatenation of the graph embedding, the embedding of the last node, and the remaining capacity. Since the first node of CVRP must be the depot node, POMO assigns the second node, i.e., the first customer node, to produce n various solutions. In the decoding procedure, the nodes already visited and the nodes with a larger demand than the remaining capacity need to be masked.

D.3 MOKP

D.3.1 Problem definition

Knapsack problem (KP) is also a representative combinatorial optimization problem. The multi-objective 0-1 knapsack problem (MOKP) with M objectives and n items is defined as follows,

$$\max \mathbf{f}(\mathbf{x}) = (f_1(\mathbf{x}), f_2(\mathbf{x}), \dots, f_M(\mathbf{x})), \quad (14)$$

$$\text{where } f_m(\mathbf{x}) = \sum_{i=1}^n v_i^m x_i, \quad (15)$$

$$\text{subject to } \sum_{i=1}^n w_i x_i \leq W, \quad (16)$$

$$x_i \in \{0, 1\}, \quad (17)$$

where item i has a weight w_i and M different values $\{v_i^1, v_i^2, \dots, v_i^M\}$. W is the weight capacity of the knapsack. The goal is to find a solution \mathbf{x} to maximize all the M total values simultaneously. The single-objective KP is also an NP-hard problem, so MOKP is even harder.




Table 3: Hypervolume illustration.

Table 4: Reference points and ideal points.

Problem	Size	r^*	z^*
Bi-TSP-1	20	(20, 20)	(0, 0)
	50	(35, 35)	(0, 0)
	100	(65, 65)	(0, 0)
	150	(85, 85)	(0, 0)
	200	(115, 115)	(0, 0)
Bi-KP	50	(5, 5)	(30, 30)
	100	(20, 20)	(50, 50)
	200	(30, 30)	(75, 75)
Tri-TSP-1	20	(20, 20, 20)	(0, 0)
	50	(35, 35, 35)	(0, 0)
	100	(65, 65, 65)	(0, 0)
Bi-CVRP	20	(30, 4)	(0, 0)
	50	(45, 4)	(0, 0)
	100	(80, 4)	(0, 0)
Bi-TSP-2	20	(20, 12)	(0, 0)
	50	(35, 25)	(0, 0)
	100	(65, 45)	(0, 0)
Tri-TSP-2	20	(20, 20, 12)	(0, 0)
	50	(35, 35, 25)	(0, 0)
	100	(65, 65, 45)	(0, 0)

D.3.2 Instance

As in the previous work [26, 12], the values and weight for each item of MOKP are all sampled from uniform distribution on $[0, 1]$. The knapsack capacity is set to $W = 12.5/25/25$ for $n = 50/100/200$.

D.3.3 Model details

For the M -objective MOKP, as each item has M values and one weight, the input dimension is $M + 1$ for the encoder. The policy network of MOKP is as the same as that of MOTSP. In the decoding procedure, the context embedding is defined as the concatenation of the graph embedding and the remaining capacity. The items already selected and the items with a larger weight than the remaining capacity need to be masked.

E Hypervolume indicator

The hypervolume (HV) indicator is widely used to evaluate the performance of the methods for MOCOPs, since HV can comprehensively measure the convergence and diversity of \mathcal{PF} without the ground truth Pareto front [51]. For a given M -objective \mathcal{PF} and a reference point r^* , $HV(\mathcal{PF}, r^*)$ is defined as follows,

$$HV(\mathcal{PF}, r^*) = \mu(S), \quad (18)$$

$$S = \{r \in \mathcal{R}^M | \exists r \in \mathcal{PF} \text{ such that } \mathbf{y} \prec r \prec r^*\}, \quad (19)$$

where $r_i^* > \max\{f_i(x) | f(x) \in \mathcal{PF}\}$ (or $r_i^* < \min\{f_i(x) | f(x) \in \mathcal{PF}\}$ for the maximization problem), $\forall i \in \{1, \dots, M\}$, and μ is the Lebesgue measure. An example with $M = 2$ is presented in Figure 3, where $\mathcal{PF} = \{f(x_1), f(x_2), f(x_3), f(x_4)\}$. $HV(\mathcal{PF}, r^*)$ is equal to the size of the grey area in this example.

Since HV value would considerably vary with the objective domain, we record the normalized HV ratio $HV'(\mathcal{PF}, r^*) = HV(\mathcal{PF}, r^*) / \prod_{i=1}^M |r_i^* - z_i^*|$, where z^* is an ideal point satisfying $z_i^* < \min\{f_i(x) | f(x) \in \mathcal{PF}\}$ (or $z_i^* > \max\{f_i(x) | f(x) \in \mathcal{PF}\}$ for the maximization problem), $\forall i \in \{1, \dots, M\}$. For a problem, all methods share the same r^* and z^* , as shown in Table 4.

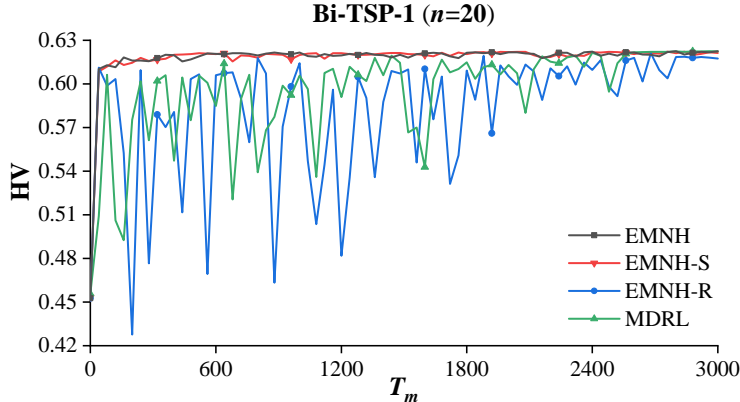


Figure 6: Training Stability on MOCOPs with balanced objective domains.

F More details of learning efficiency

F.1 Training efficiency

EMNH uses almost the equal number of training instances to MDRL [13] and PMOCO [12]. Attributed to the accelerated training method, EMNH only consumes approximately $1/\tilde{N}$ training time of MDRL. For EMNH, we set $T_m = 3000$, $T_u = 100$, and $B = 64$ to train the meta-model, i.e., overall 1.92×10^7 training instances. MDRL adopts the same training hyper-parameters as EMNH, i.e., the same amount of training instances. According to the settings of PMOCO, it trains the model by 200 epochs, with 1×10^5 instances at each epoch, i.e., overall 2×10^7 training instances, a little more than that of EMNH. The total training time of PMOCO is close to that of EMNH.

As a multi-model method, DRL-MOA [8] needs to train multiple networks with more training instances to deal with multiple subproblems. DRL-MOA first trains a submodel for the first weight vector by 200 epochs with 1×10^5 instances at each epoch, and then transfers its parameter by 5 epochs to derive another submodel for its neighbor subproblem. With a sequential parameter-transfer process, a set of submodels are obtained for various pre-given weight vectors. For the bi/tri-objective problems, the number of subproblems N is set to 101/105. Hence, DRL-MOA needs to train 101/105 submodels with total 700/720 training epochs.

F.2 Training stability

Figure 6 shows the results on Bi-TSP-1 with the balanced objective domains. EMNH-S and EMNH-R denote EMNH with the symmetric sampling method and the random sampling method, respectively. HV is computed by fine-tuning the meta-model with $K = 1$ step at each level. The results show that EMNH achieves the stablest training process. The performance of EMNH and EMNH-S are close, while the training processes of other two methods with random sampling are unstable. However, for Bi-CVRP with the imbalanced objective domains, EMNH, which adjusts the sampled weight vectors by the objective domains, outperforms other methods, as shown in Figure 4(b).

Table 5 reports the final HV indicators of the trained model, which also exhibits higher solution quality of EMNH. The best result and its statistically indifferent results using a Wilcoxon rank-sum test at 1% level are highlighted as **bold**. The second-best result and its statistically indifferent results are highlighted as underline. The method with “-Aug” represents the inference results using instance augmentation (see Appendix G) of POMO.

F.3 Fine-tuning efficiency

To study the fine-tuning efficiency, we compare the hierarchical fine-tuning method of EMNH with the vanilla fine-tuning method of MDRL. For K steps at each level with $M = 2$ in EMNH, the total fine-tuning step is $\sum_{l=1}^{L-1} 2^l K + NK = 227K$ with $N = 101$, where L satisfies $2^{L-1} < N$ and $2^L \geq N$, i.e., $L = 7$. For MDRL, each submodel is fine-tuned directly from the meta-model with \tilde{K} steps,

Table 5: Solution quality of the scaled symmetric sampling method.

Method		$n=20$		$n=50$	
		HV \uparrow	Gap \downarrow	HV \uparrow	Gap \downarrow
Bi-TSP-1	MDRL	0.6271	0.00%	0.6374	0.53%
	EMNH-R	0.6269	0.04%	0.6339	1.07%
	EMNH-S	<u>0.6270</u>	<u>0.02%</u>	0.6362	0.72%
	EMNH	0.6271	0.00%	0.6362	0.72%
	MDRL-Aug	0.6271	0.00%	0.6408	0.00%
	EMNH-R-Aug	0.6269	0.04%	<u>0.6391</u>	<u>0.26%</u>
	EMNH-S-Aug	<u>0.6270</u>	<u>0.02%</u>	0.6408	0.00%
	EMNH-Aug	0.6271	0.00%	0.6408	0.00%
Bi-CVRP	MDRL	0.4291	0.26%	0.4082	0.58%
	EMNH-R	0.3774	12.26%	0.3743	8.83%
	EMNH-S	0.4158	3.34%	0.3825	6.84%
	EMNH	<u>0.4299</u>	<u>0.07%</u>	<u>0.4098</u>	<u>0.19%</u>
	MDRL-Aug	<u>0.4294</u>	<u>0.19%</u>	<u>0.4092</u>	<u>0.34%</u>
	EMNH-R-Aug	0.3849	10.53%	0.3786	7.78%
	EMNH-S-Aug	0.4179	2.87%	0.3870	5.76%
	EMNH-Aug	0.4302	0.00%	0.4106	0.00%

so the total fine-tuning step is $N\tilde{K} = 101\tilde{K}$. Thus, for $K = 1/5/10/15/20$, $\tilde{K} = 2/11/22/34/45$ ensures that the total fine-tuning step of MDRL is approximately equal to that of EMNH. For $M = 3$ with $N = 105$, K is set to 25 for EMNH, while \tilde{K} is set to 45 for MDRL to make the equal total of fine-tuning steps.

Figure 7 shows that, on various MOCOPs, EMNH attains better solution quality than MDRL with equal total fine-tuning steps, especially a larger gap with a few fine-tuning steps, which means that EMNH has higher fine-tuning efficiency. Furthermore, PMOCO is equipped with the hierarchical fine-tuning method for a fair comparison. However, it could hardly get improved, which indicates that EMNH has more potential to derive performant submodels to deal with specific tasks. Besides, in most cases, EMNH outperforms PMOCO only with a few fine-tuning steps ($K = 5$).

G Instance augmentation

Instance augmentation exploits multiple efficient transformations for the original instance that share the same optimal solution. Then, all transformed problems are solved and the best solution among them are finally selected. According to POMO [26], a 2D coordinate (x, y) has eight different transformations, $\{(x, y), (y, x), (x, 1-y), (y, 1-x), (1-x, y), (1-y, x), (1-x, 1-y), (1-y, 1-x)\}$.

For the M -objective TSP type 1, each node has M 2D coordinates, so it has 8^M different transformations. For the M -objective TSP type 2, each node has $(M - 1)$ 2D coordinates and a 1D coordinate. The 1D coordinate x has two different transformations, $\{x, 1 - x\}$. Thus, the M -objective TSP type 2 has $2 \times 8^{M-1}$ different transformations. For MOCVRP, each node has a 2D coordinate, so it has 8 different transformations. MOKP has no transformation.

H More results on MOCOPs with imbalanced objective domains

H.1 Scaled weight sampling method for training

For the problems with imbalanced objective domains, including Bi-CVRP, Bi-TSP-2, and Tri-TSP-2, we further study PMOCO with a scaled sampling method in the training phase, denoted as PMOCO-S. Specifically, in each sampling for PMOCO, a sampled weight vector λ is scaled by f' as $\lambda_m^s = \lambda_m / f'_m$, where f'_m is dynamically estimated by the model on a validation dataset associated with $\lambda = (1/M, \dots, 1/M)$ during training. Table 6 shows that PMOCO-S could certainly improve the performance of PMOCO for Bi-CVRP and Tri-TSP-2, but it is still inferior to EMNH.

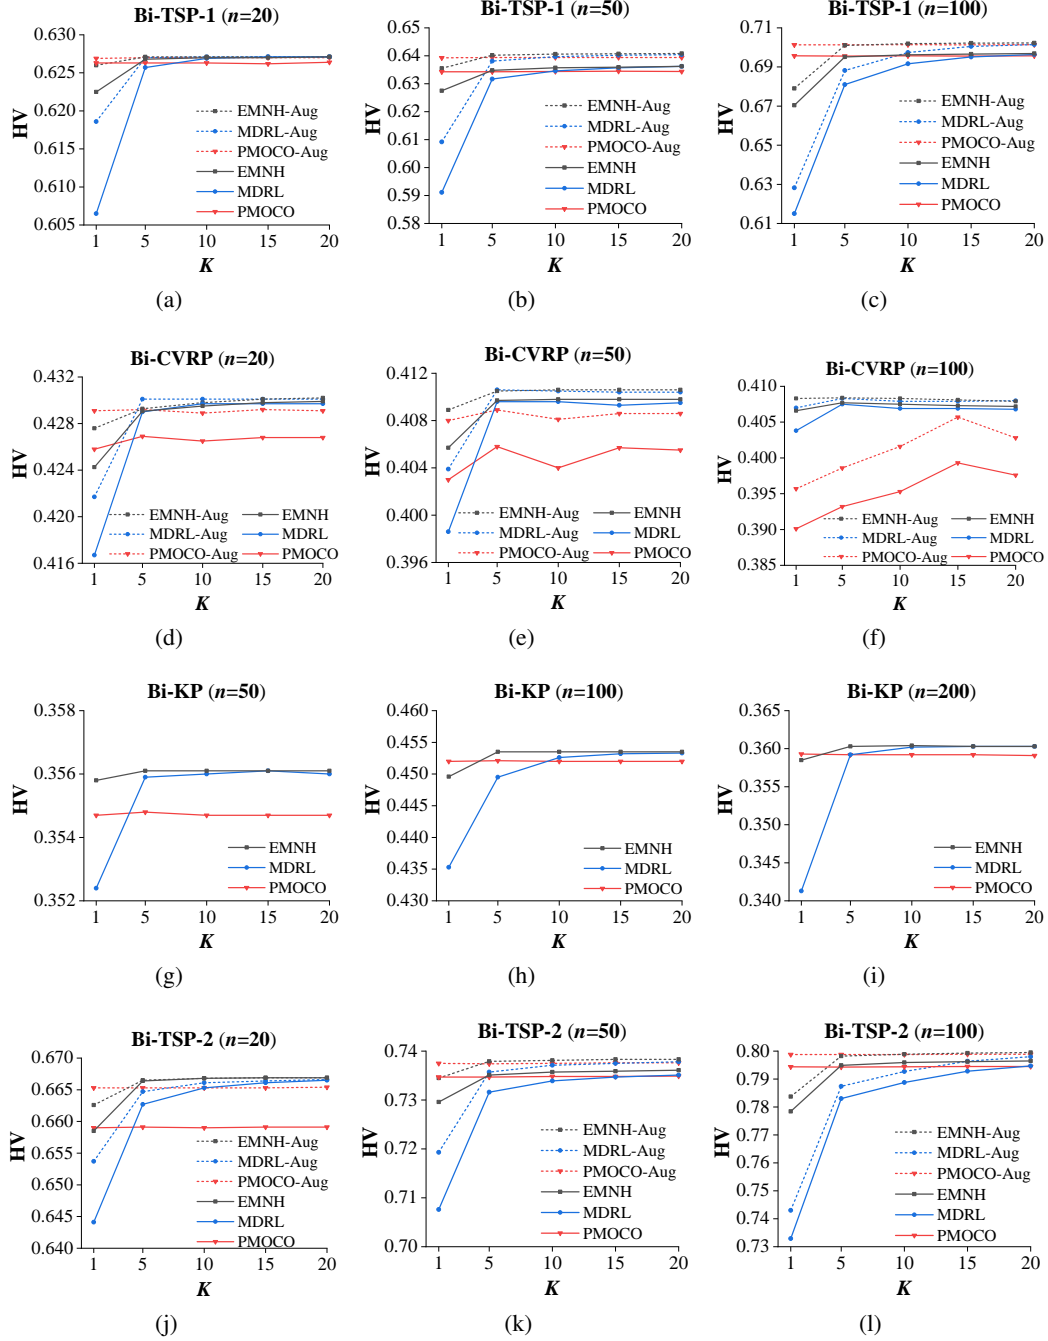


Figure 7: Fine-tuning efficiency on various MOCOPs.

H.2 Scaled weight assignment method for inference

Similar to other decomposition-based methods like PMOCO and MDRL, EMNH also faces the challenge of non-uniformly solution distributions. However, this issue can be mitigated by employing appropriate weight assignment methods during inference, as EMNH offers the flexibility to handle arbitrary weight vectors. Specifically, when the approximate scales of different objectives are known, we can normalize them to $[0,1]$ to achieve a more uniform Pareto front. Alternatively, we can adjust the weight assignment to generate a more uniform Pareto front.

Table 6: PMOCO with the scaled sampling method.

Method		$n=20$		$n=50$		$n=100$	
		HV \uparrow	Gap \downarrow	HV \uparrow	Gap \downarrow	HV \uparrow	Gap \downarrow
Bi-CVRP	PMOCO	0.4267	0.81%	0.4036	1.70%	0.3913	4.07%
	PMOCO-S	0.4274	0.65%	0.4057	1.19%	0.4042	0.91%
	EMNH	0.4299	0.07%	0.4098	0.19%	0.4072	0.17%
	PMOCO-Aug	0.4294	0.19%	0.4080	0.63%	0.3969	2.70%
	PMOCO-S-Aug	0.4295	0.16%	0.4090	0.39%	0.4063	0.39%
	EMNH-Aug	0.4302	0.00%	0.4106	0.00%	0.4079	0.00%
Bi-TSP-2	PMOCO	0.6590	1.18%	0.7347	0.51%	0.7944	0.65%
	PMOCO-S	0.6520	2.23%	0.7333	0.70%	0.7927	0.86%
	EMNH	0.6669	0.00%	0.7361	0.32%	0.7965	0.39%
	PMOCO-Aug	0.6653	0.24%	0.7375	0.14%	0.7988	0.10%
	PMOCO-S-Aug	0.6624	0.67%	0.7366	0.26%	0.7974	0.28%
	EMNH-Aug	0.6669	0.00%	0.7385	0.00%	0.7996	0.00%
Tri-TSP-2	PMOCO	0.5020	0.30%	0.5176	1.95%	0.5777	2.02%
	PMOCO-S	0.4917	2.34%	0.5175	1.97%	0.5778	2.00%
	EMNH	0.5022	0.26%	0.5205	1.40%	0.5813	1.41%
	PMOCO-Aug	0.5035	0.00%	0.5258	0.40%	0.5862	0.58%
	PMOCO-S-Aug	0.5029	0.12%	0.5259	0.38%	0.5863	0.56%
	EMNH-Aug	0.5035	0.00%	0.5279	0.00%	0.5896	0.00%

A scaled weight assignment (SWA) method can be directly applied to alleviate this issue. Specifically, each uniform weight vector λ is scaled by f' as $\lambda_m^s = \lambda_m / f'_m$ and normalized to $[0, 1]^M$. Here, f' is estimated using a validation dataset associated with $\lambda = (1/M, \dots, 1/M)$. The advantage of this SWA method is that it does not require prior problem information.

The results on Tri-TSP instances with asymmetric Pareto fronts are presented in Figure 8. For these instances, the coordinates for the three objectives are randomly sampled from $[0, 1]^2$, $[0, 0.5]^2$, $[0, 0.1]^2$, respectively. The results demonstrate that EMNH-SWA effectively produces a more uniform Pareto front. Compared to a scaling weight method with prior knowledge used in PMOCO [12], where uniform weight vectors are element-wise multiplied by (1,2,10) and then normalized back to $[0, 1]^3$, EMNH-SWA achieves desirable performance.

I Detailed results of generalization capability

We test the generalization capability of our EMNH on larger-scale random instances ($n=150/200$) and three commonly used MOTSP benchmark instances (KroAB100/150/200) in TSPLIB [52], as shown in Tables 7 and 8, respectively. EMNH, which is trained and fine-tuned both on the instances with $n=100$, has a superior generalization capability compared with the state-of-the-art MOEA and other neural heuristics for larger problem sizes.

J Hyper-parameter study

J.1 Number of sampled weight vectors

Figure 9 shows the results for various \tilde{N} . For bi-objective problems, $\tilde{N} = 1$ leads to a significantly unstable training process, since the sample number is quite small and it has no symmetric sample. $\tilde{N} = 3$ also causes a slightly unstable training process. $\tilde{N} = kM$ with $k \in \{1, 2, \dots\}$ can effectively stabilize the training process. In summary, $\tilde{N} = M$ is a more favorable setting.

J.2 Scalarization method

EMNH is generic for solving MOCOPs based on decomposition [17], which can employ various scalarization methods, including weighted sum (WS) and Tchebycheff (TCH).

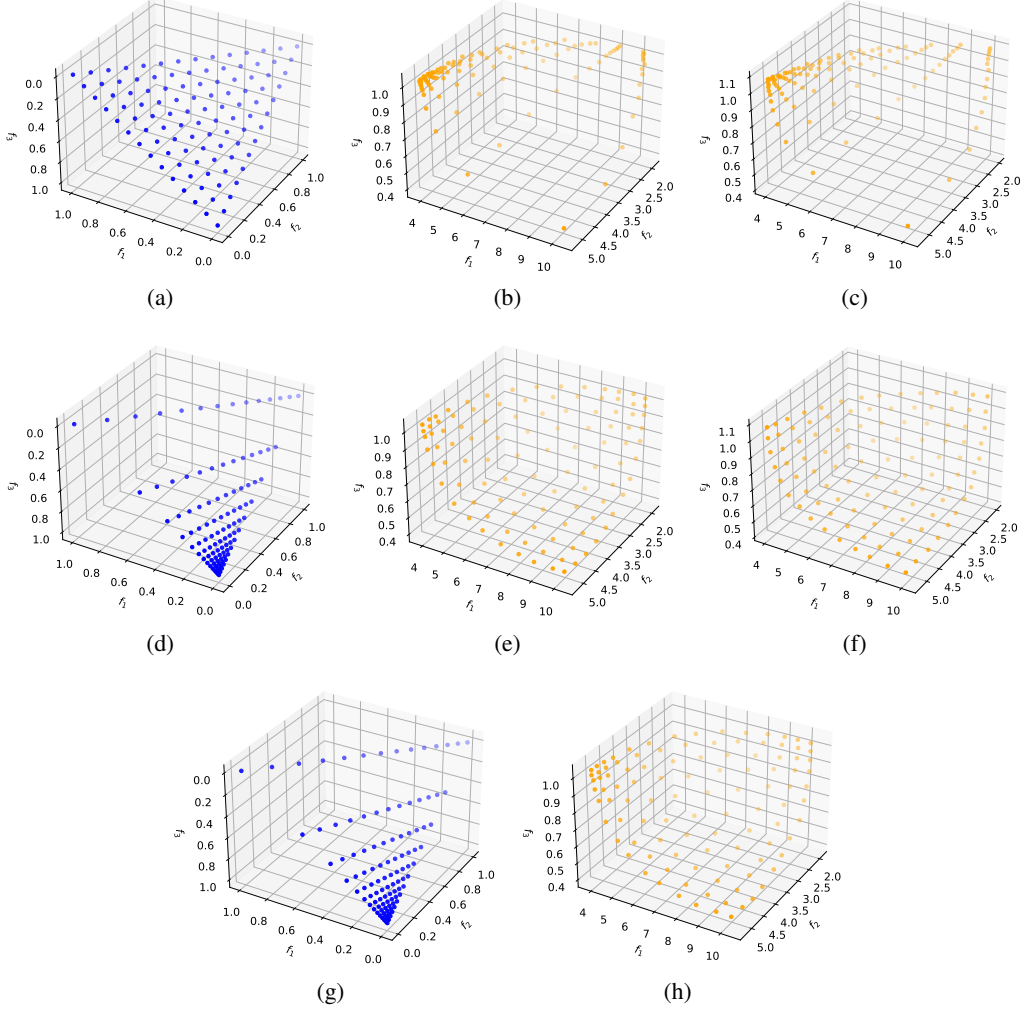


Figure 8: Solutions generated by using 105 uniform/non-uniform distributed weights on Tri-TSP-1 ($n=20$) with asymmetric Pareto front. (a) Uniform weights. (b) EMNH with uniform weights. (c) PMOCO with uniform weights. (d) A priori non-uniform weights. (e) EMNH with a priori non-uniform weights. (f) PMOCO with a priori non-uniform weights. (g) Weights generated by the scaled weight assignment (SWA) method. (h) EMNH with SWA.

Table 7: Results of generalization capability on 200 random instances .

Method	Bi-TSP-1 ($n=150$)			Bi-TSP-1 ($n=200$)		
	HV \uparrow	Gap \downarrow	Time	HV \uparrow	Gap \downarrow	Time
WS-LKH	0.7149	-2.38%	13h	0.7490	-2.50%	22h
PPLS/D-C	0.6784	3.25%	21h	0.7106	3.11%	32h
DRL-MOA	0.6901	1.17%	45s	0.7219	1.20%	87s
PMOCO	0.6910	1.05%	45s	0.7231	1.04%	87s
MDRL	0.6922	0.87%	40s	0.7251	0.77%	84s
EMNH	0.6930	0.76%	40s	0.7260	0.64%	83s
PMOCO-Aug	0.6967	0.23%	47m	0.7283	0.33%	1.6h
MDRL-Aug	0.6976	0.10%	47m	0.7299	0.11%	1.6h
EMNH-Aug	<u>0.6983</u>	<u>0.00%</u>	47m	<u>0.7307</u>	<u>0.00%</u>	1.6h

Table 8: Results of generalization capability on benchmark instances.

Method	KroAB100			KroAB150			KroAB200		
	HV \uparrow	Gap \downarrow	Time	HV \uparrow	Gap \downarrow	Time	HV \uparrow	Gap \downarrow	Time
WS-LKH	0.7022	-0.92%	2.3m	0.7017	-1.81%	4.0m	0.7430	-2.20%	5.6m
PPLS/D-C	0.6785	2.77%	31m	0.6659	3.84%	1.1h	0.7100	2.69%	3.1h
DRL-MOA	0.6903	0.79%	10s	0.6794	1.42%	18s	0.7185	1.17%	23s
PMOCO	0.6878	1.15%	9s	0.6819	1.06%	17s	0.7193	1.06%	23s
MDRL	0.6881	1.11%	10s	0.6831	0.89%	17s	0.7209	0.84%	23s
EMNH	0.6900	0.83%	9s	0.6832	0.87%	16s	0.7217	0.73%	23s
PMOCO-Aug	0.6937	0.30%	12s	0.6886	0.09%	19s	0.7251	0.26%	27s
MDRL-Aug	0.6950	0.11%	13s	0.6890	0.03%	19s	0.7261	0.12%	28s
EMNH-Aug	<u>0.6958</u>	<u>0.00%</u>	12s	<u>0.6892</u>	<u>0.00%</u>	18s	<u>0.7270</u>	<u>0.00%</u>	27s

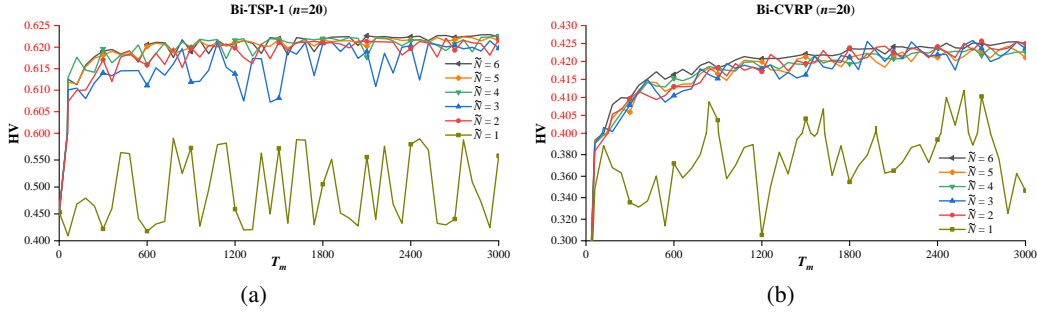


Figure 9: Effects of the number of sampled weight vectors.

Table 9: Effects of the scalarization method.

Method	$n=20$		$n=50$		$n=100$		
	HV \uparrow	Gap \downarrow	HV \uparrow	Gap \downarrow	HV \uparrow	Gap \downarrow	
Bi-TSP-1	EMNH-TCH	<u>0.6271</u>	0.00%	0.6331	1.20%	0.6927	1.37%
	EMNH-WS	<u>0.6271</u>	0.00%	0.6364	0.69%	0.6969	0.77%
Bi-TSP-1	EMNH-TCH-Aug	0.6294	-0.37%	<u>0.6401</u>	<u>0.11%</u>	<u>0.6995</u>	<u>0.40%</u>
	EMNH-WS-Aug	<u>0.6271</u>	0.00%	0.6408	0.00%	0.7023	0.00%
Tri-TSP-1	EMNH-TCH	0.4665	1.00%	0.4175	5.50%	0.4681	5.87%
	EMNH-WS	0.4698	0.30%	<u>0.4324</u>	<u>2.13%</u>	<u>0.4866</u>	<u>2.15%</u>
	EMNH-TCH-Aug	<u>0.4710</u>	<u>0.04%</u>	0.4295	2.78%	0.4814	3.20%
	EMNH-WS-Aug	0.4712	0.00%	0.4418	0.00%	0.4973	0.00%

WS is the simplest method and is effective for the convex \mathcal{PF} . It considers the linear combination of M objectives, as follows,

$$\min_{x \in \mathcal{X}} g_{\text{ws}}(x|\boldsymbol{\lambda}) = \sum_{m=1}^M \lambda_m f_m(x). \quad (20)$$

In theory, TCH can tackle the concave \mathcal{PF} , but it would lead to a more complex objective function. It is defined as follows,

$$\min_{x \in \mathcal{X}} g_{\text{tch}}(x|\boldsymbol{\lambda}) = \max_{1 \leq m \leq M} \{\lambda_m |f_m(x) - z_m^*|\}, \quad (21)$$

where $z_m^* < \min_{x \in \mathcal{X}} f_m(x)$ is an ideal value of $f_m(x)$.

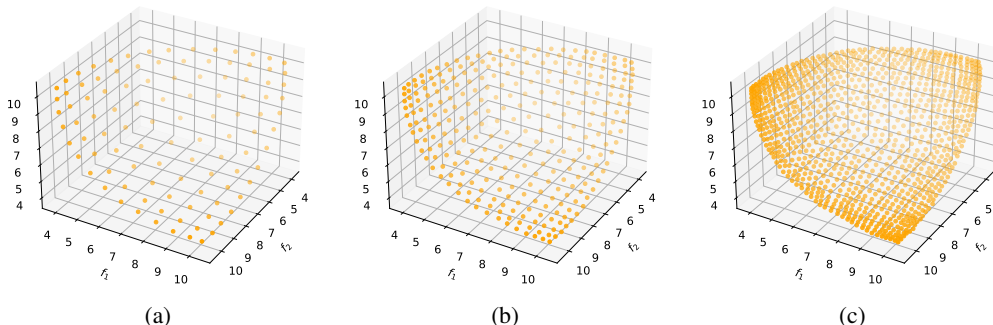
As shown in Table 9, the results show that WS is a simple yet effective scalarization method for the studied problems. In principle, EMNH can freely use any existing scalarizing function. Different scalarization methods have their own strengths and drawbacks, but the study with respect to scalarization methods is beyond the scope of this paper, which we will investigate in the future.

Table 10: Results of lightweight fine-tuning methods.

Method		$n=20$		$n=100$	
		HV	Gap	HV	Gap
Bi-CVRP	PMOCO-Aug	0.4294	0.19%	0.3966	2.77%
	MDRL-Aug	0.4292	0.23%	0.4072	0.17%
	EMDRL-Aug	0.4302	0.00%	0.4079	0.00%
	EMDRL-FD-Aug	0.4299	0.07%	0.4082	-0.07%
	EMDRL-FH-Aug	0.4298	0.09%	0.4082	-0.07%
Tri-TSP-1	PMOCO-Aug	0.4712	0.00%	0.4956	0.34%
	MDRL-Aug	0.4712	0.00%	0.4958	0.30%
	EMDRL-Aug	0.4712	0.00%	0.4973	0.00%
	EMDRL-FD-Aug	0.4710	0.04%	0.4925	0.97%
	EMDRL-FH-Aug	0.4707	0.11%	0.4906	1.35%

Table 11: Parameter numbers of various parts of models.

#(Parameters)	Bi-CVRP Model			Tri-TSP-1 Model		
	Whole Model	Decoder	Head	Whole Model	Decoder	Head
	1287K	98K	16K	1303K	115K	16K

Figure 10: Solutions of Tri-TSP-1 ($n=20$) generated with various numbers of weight vectors. (a) 105 weights. (b) 300 weights. (c) 1035 weights.

K Trade-off between lightweight fine-tuning and performance

For a given weight vector, EMNH fine-tunes the meta-model to derive a submodel to solve the corresponding subproblem. We study another two (relatively) lightweight fine-tuning methods, including only updating the head parameter (denoted as EMNH-FH) according to feature reuse [15] and only updating the decoder parameter (denoted as EMNH-FD) like PMOCO [12]. These two methods even allow us to only fine-tune and store parts of the original submodels, i.e., N heads or N decoders, thereby being more computationally efficient. Meanwhile, such benefit may bring about performance sacrifices in some cases. We report the results in Table 10 and the parameter numbers of various parts of the model in Table 11. The lightweight fine-tuning has slightly inferior performance compared with the original EMNH in most cases except on Bi-CVRP ($n=100$). Generally, the more lightweight of the fine-tuning, the more performance deterioration (i.e., EMNH-FH v.s. EMNH-FD as displayed in the table below, where FH is more light than FD). However, these lightweight fine-tuning methods can be used as alternatives when the computational and memory resources are limited.

Moreover, same as EMNH, EMNH-FH can also generate much more dense Pareto solutions to improve the performance via increasing weight vectors and corresponding fine-tuned heads. We have plotted the generated Pareto fronts with 105, 300 and 1035 weight vectors on Tri-TSP-1 which verified the above point, as shown in Figure 10.