

NEURAL MULTI-OBJECTIVE COMBINATORIAL OPTIMIZATION VIA GRAPH-IMAGE MULTIMODAL FUSION

Anonymous authors

Paper under double-blind review

ABSTRACT

Existing neural multi-objective combinatorial optimization (MOCO) methods still exhibit an optimality gap since they fail to fully exploit the intrinsic features of problem instances. A significant factor contributing to this shortfall is their reliance solely on graph-modal information. To overcome this, we propose a novel graph-image multimodal fusion (GIMF) framework that enhances neural MOCO methods by integrating graph and image information of the problem instances. Our GIMF framework comprises three key components: (1) a constructed coordinate image to better represent the spatial structure of the problem instance, (2) a problem-size adaptive resolution strategy during the image construction process to improve the cross-size generalization of the model, and (3) a multimodal fusion mechanism with modality-specific bottlenecks to efficiently couple graph and image information. We demonstrate the versatility of our GIMF by implementing it with two state-of-the-art neural MOCO backbones. Experimental results on classic MOCO problems show that our GIMF significantly outperforms state-of-the-art neural MOCO methods and exhibits superior generalization capability.

1 INTRODUCTION

Multi-objective combinatorial optimization (MOCO) involves optimizing multiple conflicting objectives within a discrete solution space, resulting in trade-offs between different criteria for informative decision making. It holds significant importance due to its wide range of real-world applications in logistics, scheduling, resource allocation and so on (Ehrgott & Gandibleux, 2000; Liu et al., 2020; Türkyılmaz et al., 2020). MOCO aims to identify a *Pareto-optimal* set of solutions, where any improvement in one objective necessitates a compromise in another. Since an MOCO problem is more complex than its single-objective counterpart, exact methods are generally impractical for MOCO, due to the NP-hard complexity of typical CO problems. As an alternative, heuristic methods have been developed to efficiently search for approximate Pareto-optimal solutions. Nevertheless, the conventional heuristics heavily hinge on problem-specific expertise and instance-specific tuning work for achieving desirable performance, posing challenges in automatic algorithm development.

With the rapid progress of *neural CO* methods for single-objective CO problems (Kool et al., 2019; Kwon et al., 2020; Grinsztajn et al., 2023; Chalumeau et al., 2023; Drakulic et al., 2023; Son et al., 2023; Xiao et al., 2024; Goh et al., 2024; Liu et al., 2024; Zhou et al., 2024; Wang et al., 2024; Kong et al., 2024), recent years have witnessed a surge in the development of *neural MOCO* methods. Differing from conventional heuristics, the neural MOCO employs deep neural models to autonomously learn constructive policies from problem instances in a data-driven manner. It allows for the automatic search of promising end-to-end solutions, bypassing the labor-intensive algorithm development in conventional heuristics. However, current neural MOCO methods still exhibit an optimality gap, as they fail to fully exploit intrinsic features of problem instances. They commonly represent a problem instance as a graph that lacks the informed representations inherent in other modalities such as images, thus hindering neural models from achieving a comprehensive encoding.

In this work, we construct images of MOCO problems with the aim to offer the complementary instance information from different modalities, i.e., graphs and images. To advance neural models, we develop a *graph-image multimodal fusion* (GIMF) framework to synergize the multimodal representation learning for enhancing the MOCO performance. The difference between our framework and current neural MOCO methods, and the images used in this paper are illustrated in Figure 1.

054
055
056
057
058
059
060
061
062
063
064
065
066
067
068
069
070
071
072
073
074
075
076
077
078
079
080
081
082
083
084
085
086
087
088
089
090
091
092
093
094
095
096
097
098
099
100
101
102
103
104
105
106
107




Figure 1: Typical neural MOCO methods (a) vs. the proposed GIMF framework (b). An example of a coordinate image (c) for an M -objective 5-node travelling salesman problem is illustrated, with blank and black pixels filled with 1 and 0, respectively.

Our contributions are summarized as follows. (1) We propose a transformation from an MOCO instance to construct a *coordinate image*, which with the graph representation, provides complementary multimodal information to facilitate comprehensive representation learning in neural models. (2) During the image construction, we present a *problem-size adaptive resolution* (PSAR) strategy that effectively enhances the model’s generalization capacity for out-of-distribution problem sizes. (3) We design a multimodal fusion mechanism with *modality-specific bottlenecks* (MSB) to favorably synergize the multimodal information from the graph and image. (4) We demonstrate the versatility of our GIMF framework by deploying it with two state-of-the-art neural MOCO methods. Experimental results on classic MOCO problems show the significant superiority of our GIMF. The advantageous effects of the multimodal fusion and generalizability are corroborated as well.

2 RELATED WORKS

Conventional MOCO methods Conventional MOCO methods are categorized into exact and heuristic approaches. Exact methods exhaustively identify Pareto-optimal solutions, resulting in prohibitive exponential computational complexity (Ehrgott et al., 2016; Figueira et al., 2017). In contrast, heuristic methods, particularly multi-objective evolutionary algorithms (MOEAs) (Deb et al., 2002; Zhang & Li, 2007; Deb & Jain, 2013; Deng et al., 2022; Qi et al., 2014; Yuan et al., 2016), are competent to efficiently search for near-optimal solutions. Furthermore, the problem-specific local search algorithms can be integrated into general MOEAs (Jaszkiewicz, 2002; Shi et al., 2020; 2024) to enhance the performance. Despite considerable domain knowledge invested in algorithm design, these conventional heuristics still require intensive intrinsic searches from scratch when solving each instance, underscoring a significant limitation in their effectiveness. For more detailed survey on heuristics for MOCO, please refer to Verma et al. (2021); Liu et al. (2020).

Neural MOCO methods Most neural MOCO methods decompose the MOCO problem into a series of scalarized subproblems, addressing each through a single-objective neural CO method, such as the notable policy optimization with multiple optima (POMO) (Kwon et al., 2020). Based on the number of trained models, neural MOCO methods can be grouped into three categories: multi-model, single-model, and unified-model methods. Multi-model methods train a set of neural models via transfer learning (Li et al., 2021; Zhang et al., 2021) or meta-learning (Zhang et al., 2023; Chen et al., 2023a), with each model specialized for each specific subproblem. In contrast, single-model methods (Lin et al., 2022; Wang et al., 2024) employ a single model for all subproblems, typically represented by preference-conditioned multi-objective combinatorial optimization (PMOCO) (Lin et al., 2022), but they still require substantial efforts in the separate training for different problem sizes. The most recent unified-model method, the conditional neural heuristic (CNH) (Fan et al., 2024), trains only a generic model that can generalize across various sizes, achieving the state-of-the-art performance in neural MOCO. Distinct from the mainstream decomposition-based neural methods, some orthogonal research attempts to enhance diversity (Chen et al., 2023b) or accelerate heuristics in exact methods (Wu et al., 2022), which significantly increase computational overhead for limited improvements. This paper focuses on improving the optimality of subproblems in decomposition-based neural methods by employing the graph-image multimodal fusion.

3 PRELIMINARY

An MOCO problem can be defined as $\min_{\boldsymbol{\pi} \in \Omega} \mathbf{f}(\boldsymbol{\pi}) = (f_1(\boldsymbol{\pi}), f_2(\boldsymbol{\pi}), \dots, f_M(\boldsymbol{\pi}))$, where \mathbf{f} is the objective vector with M objective functions, $\boldsymbol{\pi}$ is the decision variable, and Ω is a discrete feasible solution space. The Pareto properties of solutions for an MOCO problem are provided as below.

Definition 1 (Pareto dominance) A solution $\boldsymbol{\pi}^1$ dominates another $\boldsymbol{\pi}^2$ (denoted as $\boldsymbol{\pi}^1 \prec \boldsymbol{\pi}^2$) if $f_i(\boldsymbol{\pi}^1) \leq f_i(\boldsymbol{\pi}^2), \forall i \in \{1, \dots, M\}$ and $f_j(\boldsymbol{\pi}^1) < f_j(\boldsymbol{\pi}^2), \exists j \in \{1, \dots, M\}$.

Definition 2 (Pareto optimality) A solution $\boldsymbol{\pi}^*$ is Pareto optimal if it is not dominated by any other solution $\boldsymbol{\pi}$. The *Pareto set* refers to all Pareto optimal solutions, i.e., $\mathcal{P} = \{\boldsymbol{\pi}^* \in \Omega \mid \nexists \boldsymbol{\pi}' \in \Omega : \boldsymbol{\pi}' \prec \boldsymbol{\pi}^*\}$. Its image in the objective space is known as the *Pareto front*, i.e., $\mathcal{F} = \{\mathbf{f}(\boldsymbol{\pi}) \in \mathcal{R}^M \mid \boldsymbol{\pi} \in \mathcal{P}\}$.

3.1 DECOMPOSITION-BASED NEURAL MOCO

The decomposition is a widely used technique for MOCO problems. An MOCO problem can be decomposed into N subproblems, each of which is a scalarized CO problem associated with a weight vector $\boldsymbol{\lambda} \in \mathcal{R}^M$ satisfying $\lambda_i \geq 0, \forall i \in \{1, \dots, M\}$ and $\sum_{i=1}^M \lambda_i = 1$. The scalarized objective $g(\boldsymbol{\pi}|\boldsymbol{\lambda})$ for a subproblem can be derived by different scalarization functions, e.g., the most straightforward weighted sum (WS) represents the objective $\min_{\boldsymbol{\pi} \in \Omega} g_{ws}(\boldsymbol{\pi}|\boldsymbol{\lambda}) = \sum_{i=1}^M \lambda_i f_i(\boldsymbol{\pi})$. After decomposition with N weight vectors, the N derived subproblems are solved by neural CO methods.

Neural CO methods for subproblems Treating the solution $\boldsymbol{\pi}$ as a sequence $\boldsymbol{\pi} = \{\pi_1, \dots, \pi_T\}$ of length T , the solution construction process for a scalarized subproblem is a Markov decision process, with following definitions: (1) The *state* at step $t \in \{1, \dots, T\}$ consists of the weight vector $\boldsymbol{\lambda}$, the current partial solution $\boldsymbol{\pi}_{1:t-1}$, and the instance \mathcal{G} . (2) The *action* is selecting a node π_t to add to $\boldsymbol{\pi}_{1:t-1}$. (3) The *state transition* is represented as $\boldsymbol{\pi}_{1:t} = \{\boldsymbol{\pi}_{1:t-1}, \pi_t\}$. (4) The *reward* is defined as the negative of the scalarized objective, i.e., $R = -g(\boldsymbol{\pi}|\mathcal{G}, \boldsymbol{\lambda})$. (5) The stochastic *policy* is parameterized by a neural model $\boldsymbol{\theta}$ and used to sequentially construct the solution, with the process denoted by $P(\boldsymbol{\pi}|\boldsymbol{\lambda}, \mathcal{G}) = \prod_{t=1}^T P_{\boldsymbol{\theta}}(\pi_t|\boldsymbol{\pi}_{1:t-1}, \boldsymbol{\lambda}, \mathcal{G})$. The policy network is typically trained by the REINFORCE algorithm (Williams, 1992), with the gradient $\nabla \mathcal{L}(\boldsymbol{\theta}) = \frac{1}{B} \sum_{i=1}^B [(g(\boldsymbol{\pi}_i|\boldsymbol{\lambda}, \mathcal{G}_i) - b) \nabla_{\boldsymbol{\theta}} \log P(\boldsymbol{\pi}_i|\boldsymbol{\lambda}, \mathcal{G}_i)]$, where B is the batch size, and b is a baseline that is often derived from the average of multiple optima as done in POMO (Kwon et al., 2020), for reducing the variance.

Graph Transformer An MOCO instance can be defined over a *graph* $\mathcal{G} = \{\mathcal{V}, \mathcal{E}\}$, where $\mathcal{V} = \{v_1, \dots, v_n\}$ denotes the node set and $\mathcal{E} = \{e(v_i, v_j) \mid v_i, v_j \in \mathcal{V}, i \neq j\}$ denotes the edge set. Given an instance graph with n nodes featured by z -dimensional vectors $\mathbf{u}_1, \dots, \mathbf{u}_n \in \mathcal{R}^z$ (see Appendix A), the majority of existing neural MOCO methods all adopt a vanilla *graph Transformer* (Kool et al., 2019) (details in Appendix B) to handle the graph input. Generally, L self-attention layers in encoder evolve node embeddings to $\mathbf{h}_1^{(L)}, \dots, \mathbf{h}_n^{(L)} \in \mathcal{R}^d$ ($d = 128$). Then, the decoder uses the attention to autoregressively infer the probability of node selection with T steps.

4 METHODOLOGY

Our graph-image multimodal fusion (GIMF) framework integrates complementary information from both graph and image modalities. The two key challenges are the construction of informative images from MOCO instances and the effective fusion of graph and image data. Correspondingly, we propose an image construction approach with a problem-size adaptive resolution (PSAR) strategy, and design a multimodal fusion mechanism with modality-specific bottlenecks (MSB).

4.1 IMAGE CONSTRUCTION

Before stepping into the image construction approach, we first introduce the definition of the image.

Definition 3 (Image) An *image* is defined as a discrete function $\mathcal{I} : \mathcal{D} \rightarrow \mathcal{C}^K$, where $\mathcal{D} = \{(x, y) \mid x, y \in \mathcal{Z}, 1 \leq x \leq W, 1 \leq y \leq H\}$ represents the set of *pixels*, with (x, y) denoting a pixel, W denoting the image width, and H denoting the image height. The parameter K refers to the number of *channels*, and \mathcal{C} is the set of possible *pixel values* for each channel. The image size

(a.k.a., *resolution*) is defined as $W \times H$. The most ubiquitous example is the 3-channel RGB image, where $\mathcal{C}^3 = [0, 255]^3$ represents the intensities in the red, green, and blue channels.

Coordinate image for MOCO Given an MOCO instance, we construct its image by specifying node features as *pixels* and *pixel values* separately, referred to as the *coordinate image*. Our image construction is generally applicable to different MOCO problems. In this paper, we consider classic MOCO problems that are extensively studied in the neural MOCO literature (Lin et al., 2022; Chen et al., 2023a;b; Fan et al., 2024), including the bi/tri-objective traveling salesman problem (Bi/Tri-TSP) (Lust & Teghem, 2010), bi-objective capacitated vehicle routing problem (Bi-CVRP) (Zajac & Huber, 2021), and bi-objective knapsack problem (Bi-KP) (Ishibuchi et al., 2015) (Please see Appendix A for detailed problem statements). Their coordinate images are described below.

For the M -objective TSP, each node $i, \forall i \in \{1, \dots, n\}$, is featured by $\mathbf{u}_i \in \mathcal{R}^{2M}$ representing M groups of coordinates. Each pair $(v_{i,2j-1}, v_{i,2j}), \forall j \in \{1, \dots, M\}$, corresponds to the j -th group of coordinates. An M -channel coordinate image $\mathcal{I}^{M\text{-TSP}}$ of the M -objective TSP, can be constructed by treating each coordinate as a pixel, and defined as follows,

$$\mathcal{I}_j^{M\text{-TSP}}(x, y) = \begin{cases} 0, & \text{if } (x, y) = (v'_{i,2j-1}, v'_{i,2j}) \\ 1, & \text{otherwise} \end{cases}, \quad (1)$$

where $(v'_{i,2j-1}, v'_{i,2j}) \in [1, W] \times [1, H]$ is a coordinate derived by normalizing the original $(v_{i,2j-1}, v_{i,2j}) \in [0, 1]^2$, that is, $(v'_{i,2j-1}, v'_{i,2j}) = (\lfloor Wv_{i,2j-1} \rfloor + 1, \lfloor Hv_{i,2j} \rfloor + 1)$. An example of $\mathcal{I}^{M\text{-TSP}}$ for a 5-node M -objective TSP instance is illustrated in Figure 1(c). For Bi-CVRP, each node feature $\mathbf{u}_i \in \mathcal{R}^3$ represents a coordinate $(v_1, v_2) \in [0, 1]^2$ and a demand $v_3 \in [0, 1]$. The single-channel coordinate image $\mathcal{I}^{\text{Bi-CVRP}}$ can be constructed by taking the normalized coordinate as a pixel and the demand as the pixel value. Formally, the image can be defined by $\mathcal{I}^{\text{Bi-CVRP}}(x, y) = v_3$, if $(x, y) = (v'_1, v'_2)$. For Bi-KP, the node feature $\mathbf{u}_i \in \mathcal{R}^3$ represents two values $(v_1, v_2) \in [0, 1]^2$ and a weight $v_3 \in (0, 1)$ of an item. Similarly, the single-channel coordinate image $\mathcal{I}^{\text{Bi-KP}}$ is constructed by taking two normalized values as a pixel and the weight as the pixel value, resulting in $\mathcal{I}^{\text{Bi-KP}}(x, y) = v_3$, if $(x, y) = (v'_1, v'_2)$.

4.2 PROBLEM-SIZE ADAPTIVE RESOLUTION

The constructed images are processed by a *vision Transformer* (Dosovitskiy et al., 2021), which is a mainstream neural model to handle image data. Concretely, a K -channel image with the resolution $W \times H$ is divided into non-overlapping patches with a uniform size $w \times h$. Each patch is flattened into a vector and linearly projected to a d -dimensional embedding. The positional encodings (e.g., the sinusoidal encoding) are added to the patch embeddings, enabling the model to be aware of the spatial relations between patches. Then, a standard Transformer encoder with L self-attention layers evolves the embeddings to capture global features and contextual relationships within the image.

When our coordinate image adheres to the fixed resolution, a common setting in ordinary image-related tasks, it could pose challenges in the generalization across problem sizes. Specifically, we define the *node density* ρ as the ratio of the non-one pixels in our coordinate image (Note: A low node density means that the number of non-one pixels approximately equals to the problem size n). Intuitively, the node density increases as the problem size n grows with a fixed resolution, since more pixels are occupied by non-one values. The density at the patch level varies proportionally due to the fixed patch size. Consequently, the neural model is subject to the out-of-distribution generalization issue when the patch-level density significantly differs from those of the training data.

Problem-size adaptive resolution (PSAR) To address the above generalization issue, we propose a PSAR strategy to maintain a relatively stable density for both the image and patches. Generally, we can adaptively adjust the resolution by setting an approximately linear relationship between $W \times H$ and n when constructing our coordinate image, so that the density is roughly uniform with varied problem sizes. In this paper, we set $w = h$ and $W = H = \lceil 10\sqrt{n}/w \rceil * w$ to maintain the node density ρ of approximately 0.01. By doing so, the number of patches (i.e., the length of the patch sequence) varies with the changing problem size. However, the vision Transformer benefits from the stable density among patches, and hence can gain better generalization performance.

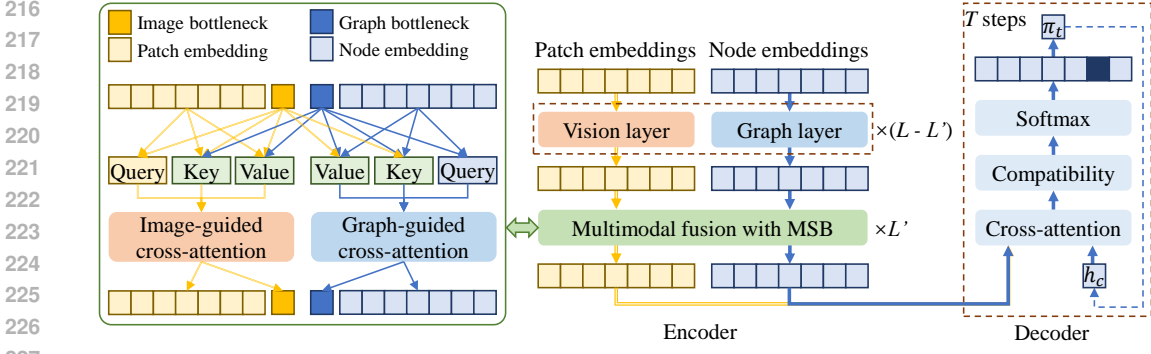


Figure 2: Our multimodal neural model based on an encoder-decoder (middle-right) architecture. The multimodal fusion layer with modality-specific bottlenecks (MSB) (left) leverages modality-specific bottlenecks to efficiently couple the graph and image information. **The dashed line pointing from the selected node π_t to the context embedding h_c indicates that h_c is determined by π_t .**

Scalable positional encoding With the fixed patch size, the PSAR results in varied lengths of the patch sequence. To indicate the positions of patches, we propose to learn scalable positional encodings. We utilize a multi-layer perceptron (MLP) with a d -dimensional hidden layer and ReLU activation to map the patch coordinates $(x^p, y^p) \in [0, 1]^2$ into position embeddings, which are then added to the patch embeddings. The MLP is applicable to varied problem sizes due to the fixed size of input (i.e., patches).

4.3 MULTIMODAL FUSION

The simple graph structures (e.g., fully-connected graphs for TSP, CVRP, and KP) may obscure spatial relationships among nodes. Our coordinate image offers a more informative view of the spatial structure, which maps a more explicit node distribution to a image. Meanwhile, the graph remains essential for fine-grained node-level encoding and decoding. To effectively integrate these two complementary modalities, we develop a multimodal fusion model with modality-specific bottlenecks.

Our multimodal fusion model, built in an encoder-decoder structure, is illustrated in Figure 2. Given the graph and coordinate image of an instance, the L -layer encoder processes them to generate high-dimensional node and patch embeddings. After that, the decoder computes the probabilities for node selection, which are used to sample nodes for constructing the solution.

Single-modal layer In the encoder, the node and patch embeddings are first passed through $L - L'$ single-modal graph and vision layers, respectively. Specifically, the initial node embeddings $\mathbf{h}_1^{(0)}, \dots, \mathbf{h}_n^{(0)} \in \mathcal{R}^d$ are obtained via linear projection with a trainable matrix W^u and bias \mathbf{b}^u , which is formulated by $\mathbf{h}_i^{(0)} = W^u \mathbf{u}_i + \mathbf{b}^u, \forall i \in \{1, \dots, n\}$. The initial patch embeddings $\mathbf{h}'_1^{(0)}, \dots, \mathbf{h}'_{n'}^{(0)} \in \mathcal{R}^d$ are derived from the flattened patch vectors $\mathbf{u}'_1, \dots, \mathbf{u}'_{n'} \in \mathcal{R}^{wh}$ via linear projection and position encoding, which is formulated by $\mathbf{h}'_i^{(0)} = W^{u'} \mathbf{u}'_i + \mathbf{b}^{u'} + \text{MLP}(x_i^p, y_i^p), \forall i \in \{1, \dots, n'\}$, where n' is the number of patches. Then, each single-modal graph and vision layer consists of a multi-head self-attention (MHSA) block and a feed-forward (FF) block. Each block is followed by the residual connection (He et al., 2016) and instance normalization (IN). The embeddings are updated at each layer $l, \forall l \in \{1, \dots, L - L'\}$, as follows,

$$\begin{aligned} \hat{\mathbf{H}} &= \text{IN}(\mathbf{H}^{(l-1)} + \text{MHSA}(\mathbf{H}^{(l-1)})), & \mathbf{H}^{(l)} &= \text{IN}(\hat{\mathbf{H}} + \text{FF}(\hat{\mathbf{H}})), \\ \hat{\mathbf{H}}' &= \text{IN}(\mathbf{H}'^{(l-1)} + \text{MHSA}(\mathbf{H}'^{(l-1)})), & \mathbf{H}'^{(l)} &= \text{IN}(\hat{\mathbf{H}}' + \text{FF}(\hat{\mathbf{H}}')), \end{aligned} \quad (2)$$

where \mathbf{H} and \mathbf{H}' denote the concatenation of the node embeddings $\{\mathbf{h}_1, \dots, \mathbf{h}_n\}$ and the patch embeddings $\{\mathbf{h}'_1, \dots, \mathbf{h}'_{n'}\}$, respectively, which are iteratively evolved by single-modal layers.

Multimodal layer After $L - L'$ single-modal layers, L' multimodal fusion layers are used to efficiently synergize the graph and image information. We introduce a small set of bottleneck tokens

to manage cross-modal interaction for condensing essential information and alleviating the computational burden of full attention calculations. Unlike the vanilla bottleneck approach (Nagrani et al., 2021), which uses shared bottlenecks across modalities, we propose modality-specific bottlenecks (MSB) to more precisely concentrate the information within each modality.

Specifically, we introduce $n_b (n_b \ll n)$ graph bottlenecks $\mathbf{B} = \{\mathbf{b}_1, \dots, \mathbf{b}_{n_b}\}$ and $n'_b (n'_b \ll n')$ image bottlenecks $\mathbf{B}' = \{\mathbf{b}'_1, \dots, \mathbf{b}'_{n'_b}\}$ into the graph and vision layer, respectively. The initial graph and image bottlenecks $\mathbf{B}_i^{(L-L'+1)} \in \mathcal{R}^d$ and $\mathbf{B}'_i^{(L-L'+1)} \in \mathcal{R}^d$ are random learnable parameters and updated through L' multimodal fusion layers. Using the bottlenecks, the efficient fusion is achieved by two types of multi-head cross-attention (MHCA). As depicted in Figure 2, the graph-guided and image-guided cross-attentions symmetrically fuse information between the two modalities. The multimodal cross-attentions at layer $l, \forall l \in \{L - L' + 1, \dots, L\}$ are given below,

$$\begin{aligned} \{\mathbf{H}^{(l)}, \mathbf{B}^{(l)}\} &= \text{MHCA}(\{\mathbf{H}^{(l-1)}, \mathbf{B}^{(l-1)}\}, \{\mathbf{H}^{(l-1)}, \mathbf{B}^{(l-1)}, \mathbf{B}'^{(l-1)}\}), \\ \{\mathbf{H}'^{(l)}, \mathbf{B}'^{(l)}\} &= \text{MHCA}(\{\mathbf{H}'^{(l-1)}, \mathbf{B}'^{(l-1)}\}, \{\mathbf{H}'^{(l-1)}, \mathbf{B}'^{(l-1)}, \mathbf{B}^{(l-1)}\}), \end{aligned} \quad (3)$$

where $\text{MHCA}(\mathbf{X}, \mathbf{Y})$ denotes the multi-head cross-attention with \mathbf{X} representing the *queries* and \mathbf{Y} representing the *keys* and *values*. The modal-specific bottlenecks improve performance of multimodal fusion, while avoiding the computational burden by self-attention between all embeddings in modalities. We use the FF block, residual connection and IN similarly as in the single-modal layer.

Probability for node selection In the decoder, we apply the node and patch embeddings processed by multimodal layers to autoregressively compute the probability for node selection. At each decoding step $t \in \{1, \dots, T\}$, a *context* embedding \mathbf{h}_c , which depends on π_{t-1} (see Appendix C), is used to compute a *glimpse* \mathbf{q}_c via an MHCA block. This glimpse is then used to calculate the *compatibility* score α , as follows,

$$\begin{aligned} \mathbf{q}_c &= \text{MHCA}(\mathbf{h}_c, \{\mathbf{H}^{(L)}, \mathbf{H}'^{(L)}\}), \\ \alpha_i &= \begin{cases} -\infty, & \text{if node } i \text{ is masked} \\ C \cdot \tanh\left(\frac{\mathbf{q}_c^T (W^K \mathbf{h}_i^{(L)})}{\sqrt{d/Y}}\right), & \text{otherwise} \end{cases} \end{aligned} \quad (4)$$

where $Y = 8$ represents the number of attention heads and the result is clipped to $C = 10$ (Kool et al., 2019). Finally, the probabilities for selecting eligible nodes are normalized using softmax.

Deployment onto neural MOCO methods Our GIMF is a generic framework that can be integrated with existing neural MOCO methods. We implement it with the state-of-the-art CNH to further enhance the performance in solving MOCO problems. Additionally, we implement GIMF with the well-known PMOCO, a type of method different from CNH, to demonstrate its flexibility. The resulting methods are referred to as GIMF-C and GIMF-P, respectively. GIMF-C employs a dual-attention to associate the instance with the weight vector of the scalarized subproblem, while GIMF-P utilizes a hypernetwork to tackle the weight vector of the scalarized subproblem. Further details about GIMF-C and GIMF-P are given in Appendix D.

5 EXPERIMENTS

5.1 EXPERIMENTAL SETTINGS

Problems We assess our GIMF framework on four well-established MOCO problems commonly examined in neural MOCO research, including Bi-TSP, Bi-CVRP, Bi-KP, and Tri-TSP. Detailed descriptions are provided in Appendix A. For the experiments, we use three standard instance sizes: $n = 20/50/100$ for Bi-TSP, Bi-CVRP, and Tri-TSP, and $n = 50/100/200$ for Bi-KP. While CNH and GIMF-C are trained across problem sizes $n \in \{20, 21, \dots, 100\}$, other neural methods are trained separately for each problem size. Moreover, we evaluate the model’s out-of-distribution generalization capability on larger instances of Bi-TSP150/200 and three widely used TSPLIB (Reinelt, 1991) benchmark instances, i.e., KroAB100/150/200.

Hyperparameters Most hyperparameters for GIMF-P and GIMF-C are configured in line with the original PMOCO and CNH, respectively. For our model, $L = 6$, $L' = 3$, and $n_b = n'_b = 10$.

The patch dimensions are fixed at $w = h = 16$. The model undergoes training for 200 epochs, with each epoch processing 100,000 randomly selected instances and a batch size of $B = 64$. The Adam optimizer (Kingma & Ba, 2015) is used with a learning rate of 10^{-4} (except 10^{-5} for Bi-KP) and weight decay of 10^{-6} . The N weight vectors for the decomposition are generated according to (Das & Dennis, 1998), with $N = 101$ for $M = 2$ and $N = 105$ for $M = 3$.

Baselines We compare our method against state-of-the-art neural MOCO methods, the widely used MOEAs, and strong problem-specific heuristics, as outlined below. (1) Neural MOCO methods: This includes the state-of-the-art **CNH** (Fan et al., 2024) and **PMOCO** (Lin et al., 2022). Additionally, we include multi-model methods such as **DRL-MOA** (Li et al., 2021), **MDRL** (Zhang et al., 2023), and **EMNH** (Chen et al., 2023a). In particular, DRL-MOA trains N POMO models for the N subproblems, starting with 200 epochs for the first model and using 5-epoch parameter transfer for each subsequent model. Both MDRL and EMNH fine-tune the N POMO models from a shared pretrained meta-model with the same structure, following the training and fine-tuning configurations as described in (Chen et al., 2023a). (2) Widely used MOEAs: Specifically, **MOEA/D** (Zhang & Li, 2007) and **NSGA-II** (Deb et al., 2002) are implemented with 4,000 iterations, representing decomposition-based and dominance-based MOEAs, respectively. **MOGLS** (Jaszkiewicz, 2002), also with 4,000 iterations and 100 local search steps per iteration, and **PPLS/D-C** (Shi et al., 2024), which runs for 200 iterations, are both tailored for MOCO with a 2-opt heuristic for TSP and CVRP and a greedy transformation heuristic (Ishibuchi et al., 2015) for KP. (3) Strong problem-specific heuristics: For the multi-objective TSP and KP, **WS-LKH** and **WS-DP** combine the weighted sum (WS) scalarization with the strong LKH (Tinós et al., 2018) and dynamic programming (DP) solvers to handle scalarized subproblems. All methods adopt WS scalarization for fair comparisons, and are executed on a machine equipped with an RTX 3090 GPU and an Intel Xeon 4216 CPU. Our codes will be made publicly available.

Metrics To assess the performance of the MOCO methods, we use the widely recognized hypervolume (HV) indicator (Audet et al., 2021), where a higher HV value indicates a superior solution set (see Appendix E for more details). We report the average HV, the gaps relative to GIMF-C-Aug, and the total computation time for 200 instances. Methods with “-Aug” incorporate instance augmentation (Lin et al., 2022) (details in Appendix F) to enhance performance. A Wilcoxon rank-sum test at a 1% significance level is conducted to assess statistical differences. The best result and those not significantly different from it are highlighted in **bold**, while the second-best result and those not significantly different from it are underlined. The names of our methods are also presented in **bold**.

5.2 EXPERIMENTAL RESULTS

Comparison analysis The comparison results are presented in Tables 1 and 2. GIMF-C consistently outperforms CNH across all scenarios, establishing itself as the new state-of-the-art neural MOCO method. Similarly, GIMF-P shows significant improvement over PMOCO in every case, and it even surpasses PMOCO-Aug, which employs instance augmentation, on Bi-CVRP100, achieving a gap of 0.66% vs 2.67%. These results confirm the effectiveness of GIMF in synergizing the complementary strengths of graph and image information. Moreover, when enhanced with instance augmentation, GIMF’s performance improves further. Compared with the cumbersome multi-model methods that require training or fine-tuning numerous models, GIMF-C maintains superiority in most cases, with only slight underperformance relative to EMNH on Bi-CVRP and Bi-KP. Notably, GIMF-C still manifests significant superiority over EMNH on Bi-TSP and Tri-TSP, such as a gap of 1.00% vs 3.01% on Tri-TSP100. Additionally, GIMF significantly reduces computational time compared with conventional methods, as demonstrated by GIMF-C-Aug, which requires just 21 minutes vs 6.0 hours for WS-LKH, while delivering competitive results.

Out-of-distribution size generalization analysis We evaluate the model’s generalization capability on the out-of-distribution larger-size Bi-TSP150/200 instances and benchmark instances KroAB100/150/200. The results of the neural methods are obtained based on models trained on Bi-TSP100, except CNH and GIMF-C trained across sizes $n \in \{20, 21, \dots, 100\}$, as shown in Table 3 and Figure 3. Compared with all neural baselines and well-known MOEAs, GIMF-C achieves the best generalization performance across all out-of-distribution cases. Similarly, GIMF-P consistently outperforms PMOCO in all scenarios. The visualized Pareto fronts of KroAB100/150/200 in Figure

Table 1: Comparison results on Bi-TSP and Bi-CVRP both with 200 random instances.

Method	Bi-TSP20			Bi-TSP50			Bi-TSP100		
	HV↑	Gap↓	Time↓	HV↑	Gap↓	Time↓	HV↑	Gap↓	Time↓
WS-LKH	0.6270	0.00%	10m	0.6415	-0.05%	1.8h	0.7090	-0.31%	6.0h
MOEA/D	0.6241	0.46%	1.7h	0.6316	1.50%	1.8h	0.6899	2.39%	2.2h
NSGA-II	0.6258	0.19%	6.0h	0.6120	4.55%	6.1h	0.6692	5.32%	6.9h
MOGLS	0.6279	-0.14%	1.6h	0.6330	1.28%	3.7h	0.6854	3.03%	11h
PPLS/D-C	0.6256	0.22%	26m	0.6282	2.03%	2.8h	0.6844	3.17%	11h
DRL-MOA	0.6257	0.21%	6s	0.6360	0.81%	9s	0.6970	1.39%	21s
MDRL	<u>0.6271</u>	<u>-0.02%</u>	5s	0.6364	0.75%	9s	0.6969	1.40%	17s
EMNH	<u>0.6271</u>	<u>-0.02%</u>	5s	0.6364	0.75%	9s	0.6969	1.40%	16s
PMOCO	0.6259	0.18%	6s	0.6351	0.95%	10s	0.6957	1.57%	19s
GIMF-P	0.6266	0.06%	7s	0.6374	0.59%	12s	0.7006	0.88%	24s
CNH	0.6270	0.00%	14s	0.6387	0.39%	17s	0.7019	0.69%	29s
GIMF-C	<u>0.6271</u>	<u>-0.02%</u>	15s	0.6399	0.20%	19s	0.7041	0.38%	33s
MDRL-Aug	<u>0.6271</u>	<u>-0.02%</u>	33s	0.6408	0.06%	1.7m	0.7022	0.65%	14m
EMNH-Aug	<u>0.6271</u>	<u>-0.02%</u>	33s	0.6408	0.06%	1.7m	0.7023	0.64%	14m
PMOCO-Aug	0.6270	0.00%	1.1m	0.6395	0.27%	3.2m	0.7016	0.74%	15m
GIMF-P-Aug	<u>0.6271</u>	<u>-0.02%</u>	1.5m	0.6403	0.14%	4.7m	0.7043	0.35%	20m
CNH-Aug	<u>0.6271</u>	<u>-0.02%</u>	1.5m	0.6410	0.03%	4.1m	0.7054	0.20%	16m
GIMF-C-Aug	0.6270	0.00%	2.0m	<u>0.6412</u>	<u>0.00%</u>	5.5m	<u>0.7068</u>	<u>0.00%</u>	21m
Method	Bi-CVRP20			Bi-CVRP50			Bi-CVRP100		
	HV↑	Gap↓	Time↓	HV↑	Gap↓	Time↓	HV↑	Gap↓	Time↓
MOEA/D	0.4255	1.07%	2.3h	0.4000	2.49%	2.9h	0.3953	3.07%	5.0h
NSGA-II	0.4275	0.60%	6.4h	0.3896	5.02%	8.8h	0.3620	11.23%	9.4h
MOGLS	0.4278	0.53%	9.0h	0.3984	2.88%	20h	0.3875	4.98%	72h
PPLS/D-C	0.4287	0.33%	1.6h	0.4007	2.32%	9.7h	0.3946	3.24%	38h
DRL-MOA	0.4287	0.33%	10s	0.4076	0.63%	12s	0.4055	0.56%	33s
MDRL	0.4291	0.23%	8s	0.4082	0.49%	13s	0.4056	0.54%	32s
EMNH	0.4299	0.05%	7s	0.4098	0.10%	13s	0.4072	0.15%	31s
PMOCO	0.4267	0.79%	7s	0.4036	1.61%	12s	0.3913	4.05%	27s
GIMF-P	0.4287	0.33%	7s	0.4076	0.63%	12s	0.4051	0.66%	29s
CNH	0.4287	0.33%	15s	0.4087	0.37%	17s	0.4065	0.32%	31s
GIMF-C	0.4292	0.21%	15s	0.4089	0.32%	18s	0.4068	0.25%	36s
MDRL-Aug	0.4294	0.16%	11s	0.4092	0.24%	36s	0.4072	0.15%	2.8m
EMNH-Aug	0.4302	-0.02%	11s	0.4106	-0.10%	35s	0.4079	-0.02%	2.8m
PMOCO-Aug	0.4294	0.16%	13s	0.4080	0.54%	36s	0.3969	2.67%	2.7m
GIMF-P-Aug	0.4298	0.07%	17s	0.4098	0.10%	48s	0.4075	0.07%	3.0m
CNH-Aug	0.4299	0.05%	22s	0.4101	0.02%	45s	0.4077	0.02%	2.5m
GIMF-C-Aug	<u>0.4301</u>	<u>0.00%</u>	25s	<u>0.4102</u>	<u>0.00%</u>	54s	<u>0.4078</u>	<u>0.00%</u>	3.1m

3 clearly show that the solutions identified by GIMF-C (or GIMF-P) are superior to those found by CNH (or PMOCO). Detailed results on these benchmark instances are provided in Appendix G.

Effectiveness of the multimodal fusion mechanism with MSB To verify the effectiveness of our MSB, we compare GIMF with models utilizing other multimodal fusion mechanisms. The baselines include the vanilla self-attention (VSA) using a unified Transformer stream, the vanilla cross-attention (VCA) with two dedicated Transformer streams, a variant of MSB using fully learnable bottlenecks (FLB) in each layer, and the vanilla shared bottlenecks (SB) (Nagrani et al., 2021) with the same total number of bottlenecks as MSB. As shown in Figure 4, the results on in-distribution Bi-TSP100 demonstrate that all multimodal methods outperform CNH, although the differences among these methods are not significant. However, the results differ on the out-of-distribution Bi-TSP200, where GIMF performs the best, while VCA, the worst-performing model, even falls short of CNH. In conclusion, while incorporating image information improves in-distribution performance, our multimodal fusion mechanism with MSB is critical for superior out-of-distribution performance.

Effectiveness of the PSAR strategy To assess the effectiveness of the PSAR strategy, we remove it from both GIMF and VSA, resulting in the variants GIMF w/o PSAR and VSA w/o PSAR. For

Table 2: Comparison results on Bi-KP and Tri-TSP both with 200 random instances.

Method	Bi-KP50			Bi-KP100			Bi-KP200		
	HV \uparrow	Gap \downarrow	Time \downarrow	HV \uparrow	Gap \downarrow	Time \downarrow	HV \uparrow	Gap \downarrow	Time \downarrow
WS-DP	0.3561	-0.06%	22m	0.4532	0.02%	2.0h	0.3601	0.00%	5.8h
MOEA/D	0.3540	0.53%	1.6h	0.4508	0.55%	1.7h	0.3581	0.56%	1.8h
NSGA-II	0.3547	0.34%	7.8h	0.4520	0.29%	8.0h	0.3590	0.31%	8.4h
MOGLS	0.3540	0.53%	5.8h	0.4510	0.51%	10h	0.3582	0.53%	18h
PPLS/D-C	0.3528	0.87%	18m	0.4480	1.17%	47m	0.3541	1.67%	1.5h
DRL-MOA	0.3559	0.00%	9s	0.4531	0.04%	18s	0.3601	0.00%	1.0m
MDRL	0.3530	0.81%	6s	0.4532	0.02%	21s	0.3601	0.00%	1.2m
EMNH	0.3561	-0.06%	6s	0.4535	-0.04%	21s	0.3603	-0.06%	1.2m
PMOCO	0.3552	0.20%	6s	0.4523	0.22%	19s	0.3595	0.07%	1.3m
GIMF-P	<u>0.3560</u>	<u>-0.03%</u>	8s	<u>0.4533</u>	<u>0.00%</u>	21s	<u>0.3602</u>	<u>-0.03%</u>	1.3m
CNH	0.3556	0.08%	18s	0.4527	0.13%	27s	0.3598	0.08%	1.2m
GIMF-C	0.3559	0.00%	18s	<u>0.4533</u>	<u>0.00%</u>	29s	0.3601	0.00%	1.4m
Method	Tri-TSP20			Tri-TSP50			Tri-TSP100		
	HV \uparrow	Gap \downarrow	Time \downarrow	HV \uparrow	Gap \downarrow	Time \downarrow	HV \uparrow	Gap \downarrow	Time \downarrow
WS-LKH	0.4712	-0.15%	12m	0.4440	-0.54%	1.9h	0.5076	-1.18%	6.6h
MOEA/D	0.4702	0.06%	1.9h	0.4314	2.31%	2.2h	0.4511	10.09%	2.4h
NSGA-II	0.4238	9.93%	7.1h	0.2858	35.28%	7.5h	0.2824	43.71%	9.0h
MOGLS	0.4701	0.09%	1.5h	0.4211	4.64%	4.1h	0.4254	15.21%	13h
PPLS/D-C	0.4698	0.15%	1.4h	0.4174	5.48%	3.9h	0.4376	12.78%	14h
DRL-MOA	0.4699	0.13%	6s	0.4303	2.56%	9s	0.4806	4.21%	19s
MDRL	0.4699	0.13%	5s	0.4317	2.24%	9s	0.4852	3.29%	16s
EMNH	0.4699	0.13%	5s	0.4324	2.08%	9s	0.4866	3.01%	16s
PMOCO	0.4693	0.26%	5s	0.4315	2.29%	8s	0.4858	3.17%	18s
GIMF-P	0.4702	0.06%	6s	0.4354	1.40%	10s	0.4927	1.79%	23s
CNH	0.4698	0.15%	10s	0.4358	1.31%	14s	0.4931	1.71%	26s
GIMF-C	0.4701	0.09%	12s	0.4382	0.77%	16s	0.4967	1.00%	29s
MDRL-Aug	0.4712	-0.15%	2.6m	0.4408	0.18%	25m	0.4958	1.18%	1.7h
EMNH-Aug	0.4712	-0.15%	2.6m	<u>0.4418</u>	<u>-0.05%</u>	25m	0.4973	0.88%	1.7h
PMOCO-Aug	0.4712	-0.15%	5.1m	0.4409	0.16%	28m	0.4956	1.22%	1.7h
GIMF-P-Aug	0.4712	-0.15%	13m	0.4415	0.02%	42m	0.5001	0.32%	2.7h
CNH-Aug	0.4704	0.02%	8.0m	0.4409	0.16%	33m	0.4996	0.42%	2.1h
GIMF-C-Aug	<u>0.4705</u>	<u>0.00%</u>	16m	0.4416	0.00%	44m	<u>0.5017</u>	<u>0.00%</u>	2.9h




Figure 3: Pareto fronts of benchmark instances, KroAB100/150/200 (left/middle/right).

these variants, we set $W = H = \lceil 100/w \rceil * w$, consistent with the size when $n = 100$. As shown in Figure 5, PSAR leads to a slight performance improvement on the in-distribution Bi-TSP100. More importantly, it significantly boosts the out-of-distribution performance for both GIMF and VSA. Notably, without PSAR, GIMF and VSA even fall below CNH’s generalization capability. These results underscore the critical role of the PSAR strategy in enhancing out-of-distribution performance.

Table 3: Out-of-distribution generalization on larger-size problems with 200 random instances.

Method	Bi-TSP150			Bi-TSP200		
	HV↑	Gap↓	Time↓	HV↑	Gap↓	Time↓
WS-LKH	0.7149	-1.63%	13h	0.7490	-1.77%	22h
MOEA/D	0.6809	3.20%	2.4h	0.7139	3.00%	2.7h
NSGA-II	0.6659	5.33%	6.8h	0.7045	4.28%	6.9h
MOGLS	0.6768	3.78%	22h	0.7114	3.34%	38h
PPLS/D-C	0.6784	3.55%	21h	0.7106	3.45%	32h
DRL-MOA	0.6901	1.89%	45s	0.7219	1.92%	1.5m
MDRL	0.6922	1.59%	40s	0.7251	1.48%	1.4m
EMNH	0.6930	1.48%	40s	0.7260	1.36%	1.4m
PMOCO	0.6910	1.76%	45s	0.7231	1.75%	1.5m
GIMF-P	0.6958	1.08%	60s	0.7267	1.26%	2.1m
CNH	0.6985	0.70%	1.1m	0.7292	0.92%	1.9m
GIMF-C	0.6995	0.55%	1.2m	0.7321	0.53%	2.2m
MDRL-Aug	0.6976	0.82%	47m	0.7299	0.83%	1.6h
EMNH-Aug	0.6983	0.73%	47m	0.7307	0.72%	1.6h
PMOCO-Aug	0.6967	0.95%	47m	0.7283	1.05%	1.6h
GIMF-P-Aug	0.7003	0.44%	1.0h	0.7311	0.67%	2.1h
CNH-Aug	0.7025	0.13%	52m	0.7343	0.23%	1.7h
GIMF-C-Aug	0.7034	0.00%	1.1h	0.7360	0.00%	2.2h




Figure 4: Effects of multimodal fusion

Figure 5: Effects of the PSAR strategy

Hyperparameter Study We conducted experiments to investigate the impact of hyperparameters. The detailed results, provided in Appendix H, indicate that both the number of multimodal fusion layers and the configuration of bottlenecks can influence the performance of multimodal fusion. For the problems we study, $L' = 3$ and $n_b = n'_b = 10$ are the desirable settings.

6 CONCLUSION

This paper proposes a generic GIMF framework that integrates complementary graph and image information for neural MOCO. We first construct a coordinate image to introduce structured image data, and the PSAR strategy during image construction is designed to enhance out-of-distribution generalization. Then, we develop a multimodal fusion mechanism with MSB to efficiently fuse graph and image information. Our GIMF framework is deployed with two state-of-the-art neural MOCO methods. Experimental results confirm its effectiveness, and the ablation study highlights the necessity of both PSAR and MSB, particularly for improving out-of-distribution performance. A limitation of our current approach is its generalization performance on real-world scenarios and various MOCO problems. In future work, we plan to address this by exploring the hierarchical approach (Goh et al., 2024) and multi-task learning technique (Zhou et al., 2024).

Another potential limitation of our method is that constructing images requires problem-specific designs to a certain extent, and incorporating images slightly increases the demand for computational resources. However, the resulting performance improvements are substantial, making this trade-off worthwhile. Furthermore, while the problems studied in this paper are well-suited for image representation, there may be some other MOCO problems for which generating simple yet meaningful images is challenging. In these cases, the applicability of our method remains an open question.

REFERENCES

- 540
541
542 Charles Audet, Jean Bignon, Dominique Cartier, Sébastien Le Digabel, and Ludovic Salomon. Per-
543 formance indicators in multiobjective optimization. *European Journal of Operational Research*,
544 292(2):397–422, 2021.
- 545
546 Felix Chalumeau, Shikha Surana, Clément Bonnet, Nathan Grinsztajn, Arnu Pretorius, Alexandre
547 Laterre, and Tom Barrett. Combinatorial optimization with policy adaptation using latent space
548 search. In *Advances in Neural Information Processing Systems*, 2023.
- 549
550 Jinbiao Chen, Jiahai Wang, Zizhen Zhang, Zhiguang Cao, Te Ye, and Siyuan Chen. Efficient meta
551 neural heuristic for multi-objective combinatorial optimization. In *Advances in Neural Informa-
tion Processing Systems*, 2023a.
- 552
553 Jinbiao Chen, Zizhen Zhang, Zhiguang Cao, Yaixin Wu, Yining Ma, Te Ye, and Jiahai Wang. Neural
554 multi-objective combinatorial optimization with diversity enhancement. In *Advances in Neural
Information Processing Systems*, 2023b.
- 555
556 I Das and JE Dennis. Normal-boundary intersection: A new method for generating pareto-optimal
557 points in multieriteria optimization problems. *SIAM Journal on Optimization*, 8(3):631–657,
558 1998.
- 559
560 Kalyanmoy Deb and Himanshu Jain. An evolutionary many-objective optimization algorithm using
561 reference-point-based nondominated sorting approach, part i: solving problems with box con-
562 straints. *IEEE transactions on evolutionary computation*, 18(4):577–601, 2013.
- 563
564 Kalyanmoy Deb, Amrit Pratap, Sameer Agarwal, and TAMT Meyarivan. A fast and elitist multi-
565 objective genetic algorithm: NSGA-II. *IEEE Transactions on Evolutionary Computation*, 6(2):
182–197, 2002.
- 566
567 Wu Deng, Xiaoxiao Zhang, Yongquan Zhou, Yi Liu, Xiangbing Zhou, Huiling Chen, and Huimin
568 Zhao. An enhanced fast non-dominated solution sorting genetic algorithm for multi-objective
569 problems. *Information Sciences*, 585:441–453, 2022.
- 570
571 Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas
572 Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An im-
573 age is worth 16x16 words: Transformers for image recognition at scale. In *International Confer-
ence on Learning Representations*, 2021.
- 574
575 Darko Drakulic, Sofia Michel, Florian Mai, Arnaud Sors, and Jean-Marc Andreoli. BQ-NCO:
576 Bisimulation quotienting for efficient neural combinatorial optimization. In *Advances in Neu-
ral Information Processing Systems*, 2023.
- 577
578 Matthias Ehrgott and Xavier Gandibleux. A survey and annotated bibliography of multiobjective
579 combinatorial optimization. *OR-spektrum*, 22(4):425–460, 2000.
- 580
581 Matthias Ehrgott, Xavier Gandibleux, and Anthony Przybylski. Exact methods for multi-objective
582 combinatorial optimisation. *Multiple criteria decision analysis: State of the art surveys*, pp. 817–
583 850, 2016.
- 584
585 Mingfeng Fan, Yaixin Wu, Zhiguang Cao, Wen Song, Guillaume Sartoretti, Huan Liu, and Guohua
586 Wu. Conditional neural heuristic for multiobjective vehicle routing problems. *IEEE Transactions
on Neural Networks and Learning Systems*, 2024.
- 587
588 José Rui Figueira, Carlos M Fonseca, Pascal Halffmann, Kathrin Klamroth, Luís Paquete, Stefan
589 Ruzika, Britta Schulze, Michael Stiglmayr, and David Willems. Easy to say they are hard, but
590 hard to see they are easy—towards a categorization of tractable multiobjective combinatorial op-
591 timization problems. *Journal of Multi-Criteria Decision Analysis*, 24(1-2):82–98, 2017.
- 592
593 Yong Liang Goh, Zhiguang Cao, Yining Ma, Yanfei Dong, Mohammed Haroon Dupty, and Wee Sun
Lee. Hierarchical neural constructive solver for real-world tsp scenarios. In *Proceedings of the
30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pp. 884–895, 2024.

- 594 Nathan Grinsztajn, Daniel Furelos-Blanco, Shikha Surana, Clément Bonnet, and Tom Barrett. Winner takes it all: Training performant rl populations for combinatorial optimization. In *Advances in Neural Information Processing Systems*, 2023.
- 595
- 596
- 597 Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *IEEE Conference on Computer Vision and Pattern Recognition*, pp. 770–778, 2016.
- 598
- 599
- 600 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.
- 601
- 602
- 603 Andrzej Jaszkiewicz. Genetic local search for multi-objective combinatorial optimization. *European Journal of Operational Research*, 137(1):50–71, 2002.
- 604
- 605
- 606 Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *International Conference on Learning Representations*, 2015.
- 607
- 608
- 609 Detian Kong, Yining Ma, Zhiguang Cao, Tianshu Yu, and Jianhua Xiao. Efficient neural collaborative search for pickup and delivery problems. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2024.
- 610
- 611
- 612 Wouter Kool, Herke Van Hoof, and Max Welling. Attention, learn to solve routing problems! In *International Conference on Learning Representations*, 2019.
- 613
- 614
- 615 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.
- 616
- 617
- 618 Kaiwen Li, Tao Zhang, and Rui Wang. Deep reinforcement learning for multiobjective optimization. *IEEE Transactions on Cybernetics*, 51(6):3103–3114, 2021.
- 619
- 620
- 621 Xi Lin, Zhiyuan Yang, and Qingfu Zhang. Pareto set learning for neural multi-objective combinatorial optimization. In *International Conference on Learning Representations*, 2022.
- 622
- 623
- 624 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.
- 625
- 626
- 627
- 628 Qi Liu, Xiaofeng Li, Haitao Liu, and Zhaoxia Guo. Multi-objective metaheuristics for discrete optimization problems: A review of the state-of-the-art. *Applied Soft Computing*, 93:106382, 2020.
- 629
- 630
- 631
- 632 Thibaut Lust and Jacques Teghem. The multiobjective traveling salesman problem: A survey and a new approach. In *Advances in Multi-Objective Nature Inspired Computing*, pp. 119–141, 2010.
- 633
- 634
- 635 Arsha Nagrani, Shan Yang, Anurag Arnab, Aren Jansen, Cordelia Schmid, and Chen Sun. Attention bottlenecks for multimodal fusion. In *Advances in neural information processing systems*, volume 34, pp. 14200–14213, 2021.
- 636
- 637
- 638 Yutao Qi, Xiaoliang Ma, Fang Liu, Licheng Jiao, Jianyong Sun, and Jianshe Wu. Moea/d with adaptive weight adjustment. *Evolutionary computation*, 22(2):231–264, 2014.
- 639
- 640
- 641 Gerhard Reinelt. TSPLIB—a traveling salesman problem library. *ORSA Journal on Computing*, 3(4):376–384, 1991.
- 642
- 643 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.
- 644
- 645
- 646 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*, 54(4):2369–2382, 2024.
- 647

- 648 Jiwoo Son, Minsu Kim, Hyeonah Kim, and Jinkyoo Park. Meta-SAGE: Scale meta-learning sched-
649 uled adaptation with guided exploration for mitigating scale shift on combinatorial optimization.
650 In *International Conference on Machine Learning*, pp. 32194–32210. PMLR, 2023.
- 651 Renato Tinós, Keld Helsgaun, and Darrell Whitley. Efficient recombination in the lin-kernighan-
652 helsgaun traveling salesman heuristic. In *International Conference on Parallel Problem Solving*
653 *from Nature*, pp. 95–107, 2018.
- 654 Alper Türkyılmaz, Özlem Şenvar, İrem Ünal, and Serol Bulkan. A research survey: heuristic ap-
655 proaches for solving multi objective flexible job shop problems. *Journal of Intelligent Manufac-*
656 *turing*, 31:1949–1983, 2020.
- 657 Shanu Verma, Millie Pant, and Vaclav Snasel. A comprehensive review on nsga-ii for multi-objective
658 combinatorial optimization problems. *IEEE access*, 9:57757–57791, 2021.
- 659 Chenguang Wang, Zhouliang Yu, Stephen McAleer, Tianshu Yu, and Yaodong Yang. ASP: Learn
660 a universal neural solver! *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 46
661 (06):4102–4114, 2024.
- 662 Ronald J Williams. Simple statistical gradient-following algorithms for connectionist reinforcement
663 learning. *Machine learning*, 8:229–256, 1992.
- 664 Yaoxin Wu, Wen Song, Zhiguang Cao, Jie Zhang, Abhishek Gupta, and Mingyan Lin. Graph learn-
665 ing assisted multi-objective integer programming. In *Advances in Neural Information Processing*
666 *Systems*, volume 35, pp. 17774–17787, 2022.
- 667 Yubin Xiao, Di Wang, Boyang Li, Mingzhao Wang, Xuan Wu, Changliang Zhou, and You Zhou.
668 Distilling autoregressive models to obtain high-performance non-autoregressive solvers for ve-
669 hicle routing problems with faster inference speed. In *Proceedings of the AAAI Conference on*
670 *Artificial Intelligence*, volume 38, pp. 20274–20283, 2024.
- 671 Yuan Yuan, Hua Xu, Bo Wang, Bo Zhang, and Xin Yao. Balancing convergence and diversity in
672 decomposition-based many-objective optimizers. *IEEE Transactions on Evolutionary Computa-*
673 *tion*, 20(2):180–198, 2016.
- 674 Sandra Zajac and Sandra Huber. Objectives and methods in multi-objective routing problems: a
675 survey and classification scheme. *European Journal of Operational Research*, 290(1):1–25, 2021.
- 676 Qingfu Zhang and Hui Li. MOEA/D: A multiobjective evolutionary algorithm based on decompo-
677 sition. *IEEE Transactions on Evolutionary Computation*, 11(6):712–731, 2007.
- 678 Yongxin Zhang, Jiahai Wang, Zizhen Zhang, and Yalan Zhou. MODRL/D-EL: Multiobjective deep
679 reinforcement learning with evolutionary learning for multiobjective optimization. In *Internat-*
680 *ional Joint Conference on Neural Networks*, 2021.
- 681 Zizhen Zhang, Zhiyuan Wu, Hang Zhang, and Jiahai Wang. Meta-learning-based deep reinforce-
682 ment learning for multiobjective optimization problems. *IEEE Transactions on Neural Networks*
683 *and Learning Systems*, 34(10):7978–7991, 2023.
- 684 Jianan Zhou, Zhiguang Cao, Yaoxin Wu, Wen Song, Yining Ma, Jie Zhang, and Chi Xu. MV-
685 MoE: Multi-task vehicle routing solver with mixture-of-experts. In *International Conference on*
686 *Machine Learning*, 2024.
- 687
688
689
690
691
692
693
694
695
696
697
698
699
700
701

A DETAILED STATEMENTS OF THE STUDIED MOCO PROBLEMS

Multi-objective traveling salesman problem (MOTSP) In the M -objective traveling salesman problem with n nodes, each node $i \in 1, \dots, n$ is associated with M groups of 2-dimensional coordinates. The Euclidean distance c_{ij}^m between nodes i and j is determined by their coordinates for objective m . The task is to find a tour π that visits all nodes while minimizing the total distances for each of the M objectives. Specifically, the objective is to minimize $\mathbf{f}(\pi) = (f_1(\pi), f_2(\pi), \dots, f_M(\pi))$, where $f_m(\pi) = \sum_{i=1}^{n-1} c_{\pi_i, \pi_{i+1}}^m + c_{\pi_n, \pi_1}^m, \forall m \in \{1, \dots, M\}$. The instances are generated by uniformly sampling coordinates within the range $[0, 1]^{2M}$.

Multi-objective capacitated vehicle routing problem (MOCVRP) We study the bi-objective capacitated vehicle routing problem (Bi-CVRP), which involves n customer nodes and one depot node. Each node has a 2-dimensional coordinate, and each customer is associated with a demand. A fleet of homogeneous vehicles, all with the same capacity, is based at the depot and must serve all customers before returning. When serving a customer, the vehicle’s remaining capacity must be at least equal to the customer’s demand. The two conflicting objectives are to minimize the total tour length and the makespan, defined as the longest route. For Bi-CVRP instances, the coordinates of both the depot and customers are uniformly sampled from the range $[0, 1]^2$, while demands are drawn from the set $\{1, \dots, 9\}$. Vehicle capacity is set at 30, 40, and 50 for $20 \leq n < 40$, $40 \leq n < 70$, and $70 \leq n \leq 100$, respectively. All demands are normalized by the vehicle capacity for consistency.

Multi-objective knapsack problem (MOKP) In the multi-objective knapsack problem with M objectives and n items, each item has a weight and M distinct values. The items are represented as nodes in the instance graph. The objective is to select items that maximize all M objectives simultaneously, while ensuring the total weight does not exceed the knapsack capacity. The instances are generated by sampling the weight and values of each item from a uniform distribution over $[0, 1]$. The knapsack capacity is set to 12.5 for $50 \leq n < 100$ and 25 for $100 \leq n \leq 200$, respectively.

B ARCHITECTURE OF THE GRAPH TRANSFORMER

The MOCO instance graph can be processed by a graph Transformer (Kool et al., 2019) based on an encoder-decoder architecture. Given n nodes with z -dimensional features $\mathbf{u}_1, \dots, \mathbf{u}_n \in \mathcal{R}^z$, the encoder first transforms them to the initial node embeddings $\mathbf{h}_1^{(0)}, \dots, \mathbf{h}_n^{(0)} \in \mathcal{R}^d$ via linear projection as $\mathbf{h}_i^{(0)} = W^u \mathbf{u}_i + \mathbf{b}^u, \forall i \in \{1, \dots, n\}$, and then compute the final node embeddings $\mathbf{h}_1^{(L)}, \dots, \mathbf{h}_n^{(L)}$ via L Transformer layer. Each Transformer layer sequentially comprises a multi-head self-attention (MHSA) block, a residual connection and instance normalization (IN) block, a feed-forward (FF) block, and another residual connection and IN block. The node embeddings $\mathbf{H}^{(l)} = \{\mathbf{h}_1^{(l)}, \dots, \mathbf{h}_n^{(l)}\}$ are updated as follows,

$$\hat{\mathbf{H}} = \text{IN}(\mathbf{H}^{(l-1)} + \text{MHSA}(\mathbf{H}^{(l-1)})) \quad (5)$$

$$\mathbf{H}^{(l)} = \text{IN}(\hat{\mathbf{H}} + \text{FF}(\hat{\mathbf{H}})). \quad (6)$$

The decoder take the derived node embeddings as inputs to calculate the selection probabilities for candidate nodes in an autoregressive manner with T steps. Specifically, at decoding step $t \in \{1, \dots, T\}$, the *glimpse* \mathbf{q}_c is first computed using a problem-specific *context* embedding \mathbf{h}_c , as follows,

$$\mathbf{q}_c = \text{MHCA}(\mathbf{h}_c, \mathbf{H}^{(L)}), \quad (7)$$

where $\text{MHCA}(\mathbf{X}, \mathbf{Y})$ the multi-head cross-attention with \mathbf{X} as the *queries* and with \mathbf{Y} as the *keys* and *values*. The definition of \mathbf{h}_c is provided in Appendix C. Then, the *compatibility* is calculated by the attention mechanism, as follows,

$$\alpha_i = \begin{cases} -\infty, & \text{if node } i \text{ is masked} \\ C \cdot \tanh\left(\frac{\mathbf{q}_c^T (W^K \mathbf{h}_i^{(L)})}{\sqrt{d/Y}}\right), & \text{otherwise} \end{cases} \quad (8)$$

where $Y = 8$ is the number of attention heads and $C = 10$ is use to clip the result. Finally, the probabilities of node selection is obtained using a softmax function.

C CONTEXT EMBEDDING

For MOTSP, the context embedding \mathbf{h}_c at each decoding step is constructed by concatenating the embeddings of the first and last visited nodes, with all visited nodes masked when calculating node selection probabilities. In MOCVRP, the context embedding \mathbf{h}_c consists of the embedding of the last visited node and the remaining vehicle capacity, and nodes already visited or with demands exceeding the remaining capacity are masked during probability computation. For MOKP, the context embedding \mathbf{h}_c combines the graph embedding $\mathbf{h} = \sum_{i=1}^n \mathbf{h}_i/n$ with the remaining knapsack capacity, masking selected items and those with weights larger than the remaining capacity during probability calculations.

D DETAILS OF GIMF-C AND GIMF-P

D.1 GIMF-C

Following CNH (Fan et al., 2024), GIMF-C also employs a multi-head dual-attention mechanism to handle the given weight vector and uses a size-aware decoder to capture the feature of the problem size. Thus, we adopt the multi-head dual-attention (MHDA) mechanism for each modality. Specifically, the initial embedding of the weight vector $\mathbf{h}_{n+1}^{(0)}$, node embeddings $\mathbf{H} = \{\mathbf{h}_1^{(0)}, \dots, \mathbf{h}_n^{(0)}\}$, and patch embeddings $\mathbf{H}' = \{\mathbf{h}'_1^{(0)}, \dots, \mathbf{h}'_n^{(0)}\}$ are first obtained using separate linear projections. Then, in each single-modal Transformer layer $l, \forall l \in \{1, \dots, L - L'\}$, the multi-head self-attention (MHSA) is replaced by the MHDA, as follows,

$$\{\hat{\mathbf{H}}, \hat{\mathbf{h}}_{n+1}\} = \text{IN}(\{\mathbf{H}^{(l-1)}, \mathbf{h}_{n+1}^{(l-1)}\} + \text{MHDA}(\mathbf{H}^{(l-1)}, \mathbf{h}_{n+1}^{(l-1)})), \quad (9)$$

$$\{\mathbf{H}^{(l)}, \mathbf{h}_{n+1}^{(l)}\} = \text{IN}(\{\hat{\mathbf{H}}, \hat{\mathbf{h}}_{n+1}\} + \text{FF}(\{\hat{\mathbf{H}}, \hat{\mathbf{h}}_{n+1}\})), \quad (10)$$

$$\{\hat{\mathbf{H}}', \hat{\mathbf{h}}'_{n+1}\} = \text{IN}(\{\mathbf{H}'^{(l-1)}, \mathbf{h}'_{n+1}^{(l-1)}\} + \text{MHDA}(\mathbf{H}'^{(l-1)}, \mathbf{h}'_{n+1}^{(l-1)})), \quad (11)$$

$$\{\mathbf{H}'^{(l)}, \mathbf{h}'_{n+1}^{(l)}\} = \text{IN}(\{\hat{\mathbf{H}}', \hat{\mathbf{h}}'_{n+1}\} + \text{FF}(\{\hat{\mathbf{H}}', \hat{\mathbf{h}}'_{n+1}\})), \quad (12)$$

where $\mathbf{h}'_{n+1}^{(0)} = \mathbf{h}_{n+1}^{(0)}$ and $\text{MHDA}(\mathbf{X}, \mathbf{Y})$ is defined as follows,

$$\text{MHDA}(\mathbf{X}, \mathbf{Y}) = \{\text{MHSA}(\{\mathbf{X}, \mathbf{Y}\})_{:-1} + \text{MHCA}(\mathbf{X}, \mathbf{Y}), \text{MHSA}(\{\mathbf{X}, \mathbf{Y}\})_{-1}\}. \quad (13)$$

After $L - L'$ single-modal Transformer layers, the embeddings are updated by L' multimodal fusion layers. The multimodal fusion layers involve additional n_b graph bottlenecks $\mathbf{B} = \{\mathbf{b}_1, \dots, \mathbf{b}_{n_b}\}$ and n'_b image bottlenecks $\mathbf{B}' = \{\mathbf{b}'_1, \dots, \mathbf{b}'_{n'_b}\}$. In each layer $l, \forall l \in \{L - L' + 1, \dots, L\}$, the embeddings are updated through the graph-guided and image-guided dual-attention, as follows,

$$\tilde{\mathbf{H}} = \text{MHCA}(\{\mathbf{B}^{(l-1)}, \mathbf{H}^{(l-1)}, \mathbf{h}_{n+1}^{(l-1)}\}, \{\mathbf{B}^{(l-1)}, \mathbf{H}^{(l-1)}, \mathbf{h}_{n+1}^{(l-1)}, \mathbf{B}'^{(l-1)}\}) \quad (14)$$

$$\{\mathbf{B}^{(l)}, \mathbf{H}^{(l)}, \mathbf{h}_{n+1}^{(l)}\} = \{\tilde{\mathbf{H}}_{:-1} + \text{MHCA}(\{\mathbf{B}^{(l-1)}, \mathbf{H}^{(l-1)}\}, \{\mathbf{h}_{n+1}^{(l-1)}, \mathbf{B}'^{(l-1)}\}), \tilde{\mathbf{H}}_{-1}\}, \quad (15)$$

$$\tilde{\mathbf{H}}' = \text{MHCA}(\{\mathbf{B}'^{(l-1)}, \mathbf{H}'^{(l-1)}, \mathbf{h}'_{n+1}^{(l-1)}\}, \{\mathbf{B}^{(l-1)}, \mathbf{H}'^{(l-1)}, \mathbf{h}'_{n+1}^{(l-1)}, \mathbf{B}^{(l-1)}\}) \quad (16)$$

$$\{\mathbf{B}'^{(l)}, \mathbf{H}'^{(l)}, \mathbf{h}'_{n+1}^{(l)}\} = \{\tilde{\mathbf{H}}'_{:-1} + \text{MHCA}(\{\mathbf{B}'^{(l-1)}, \mathbf{H}'^{(l-1)}\}, \{\mathbf{h}'_{n+1}^{(l-1)}, \mathbf{B}^{(l-1)}\}), \tilde{\mathbf{H}}'_{-1}\}. \quad (17)$$

In the decoder, the size-injected node embeddings $\bar{\mathbf{H}}^{(L)}$ are obtained by adding the problem size embedding (Fan et al., 2024) to the node embeddings. Then, the glimpse is derived as follows,

$$\mathbf{q}_c = \text{MHCA}(\mathbf{h}_c, \{\bar{\mathbf{H}}^{(L)}, \mathbf{H}'^{(L)}\}). \quad (18)$$

Finally, the probabilities of node selection can be computed as the same.

D.2 GIMF-P

Following PMOCO (Lin et al., 2022), GIMF-P directly uses the multimodal fusion model as the base model. To manage each scalarized subproblem associated with a weight vector, GIMF-P employs a hypernetwork with a multi-layer perceptron (MLP) structure to generate decoder parameters, using the weight vector as input.

Table 4: Reference points and ideal points for the MOCO problems.

Problem	Size	\mathbf{r}	\mathbf{z}
Bi-TSP	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-CVRP	20	(30, 4)	(0, 0)
	50	(45, 4)	(0, 0)
	100	(80, 4)	(0, 0)
Bi-KP	50	(5, 5)	(30, 30)
	100	(20, 20)	(50, 50)
	200	(30, 30)	(75, 75)
Tri-TSP	20	(20, 20, 20)	(0, 0)
	50	(35, 35, 35)	(0, 0)
	100	(65, 65, 65)	(0, 0)

E HYPERVOLUME

Hypervolume (HV) is a widely used indicator for evaluating the performance of MOCO methods, as it effectively measures both the convergence and diversity of the obtained Pareto front without requiring ground truth. The HV of a Pareto front \mathcal{F} with respect to a reference point $\mathbf{r} \in \mathcal{R}^M$, denoted as $\text{HV}_{\mathbf{r}}(\mathcal{F})$, is defined as:

$$\text{HV}_{\mathbf{r}}(\mathcal{F}) = \mu \left(\bigcup_{\mathbf{f}(\boldsymbol{\pi}) \in \mathcal{F}} [\mathbf{f}(\boldsymbol{\pi}), \mathbf{r}] \right), \quad (19)$$

where μ represents the Lebesgue measure, and $[\mathbf{f}(\boldsymbol{\pi}), \mathbf{r}]$ refers to an M -dimensional hypercube, i.e., $[\mathbf{f}(\boldsymbol{\pi}), \mathbf{r}] = [f_1(\boldsymbol{\pi}), r_1] \times \cdots \times [f_M(\boldsymbol{\pi}), r_M]$. The HV is normalized as $\text{HV}'_{\mathbf{r}}(\mathcal{F}) = \text{HV}_{\mathbf{r}}(\mathcal{F}) / \prod_{i=1}^M |r_i - z_i|$, where \mathbf{z} is an ideal point such that $z_i < \min\{f_i(\boldsymbol{\pi}) | \mathbf{f}(\boldsymbol{\pi}) \in \mathcal{F}\}$ (or $z_i > \max\{f_i(\boldsymbol{\pi}) | \mathbf{f}(\boldsymbol{\pi}) \in \mathcal{F}\}$ for maximization), $\forall i \in \{1, \dots, M\}$. According to prior literature (Chen et al., 2023a;b), the same \mathbf{r} and \mathbf{z} are used across all methods for a given MOCO problem, as summarized in Table 4.

F MULTI-OBJECTIVE INSTANCE AUGMENTATION

During the inference phase, instance augmentation Lin et al. (2022) can be used to improve performance by transforming an instance into multiple variations that all share the same optimal solution. Each transformed instance is then solved, and the best solution among them is chosen. For Bi-CVRP, there are 8 possible transformations based on the 2-dimensional coordinates, such as (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, this leads to 8^M transformations, as each of the M groups of coordinates can be permuted independently. For KP, this instance augmentation is not applicable.

G GENERALIZATION RESULTS ON BENCHMARK INSTANCES

The detailed out-of-distribution generalization results are presented in Table 5, further confirming the exceptional generalization ability of our GIMF.

H HYPERPARAMETER STUDY

Effects of the number of multimodal fusion layers We vary the number of multimodal fusion layers L' while keeping the total number of layers fixed at $L = 6$. The results in Figure 6 show

Table 5: Detailed results on benchmark instances.

Method	KroAB100			KroAB150			KroAB200		
	HV↑	Gap↓	Time↓	HV↑	Gap↓	Time↓	HV↑	Gap↓	Time↓
WS-LKH	0.7022	-0.34%	2.3m	0.7017	-0.99%	4.0m	0.7430	-1.42%	5.6m
MOEA/D	0.6836	2.31%	5.8m	0.6710	3.43%	7.1m	0.7106	3.00%	7.3m
NSGA-II	0.6676	4.60%	7.0m	0.6552	5.70%	7.9m	0.7011	4.30%	8.4m
MOGLS	0.6817	2.59%	52m	0.6671	3.99%	1.3h	0.7083	3.32%	1.6h
PPLS/D-C	0.6785	3.04%	38m	0.6659	4.16%	1.4h	0.7100	3.08%	3.8h
DRL-MOA	0.6903	1.36%	10s	0.6794	2.22%	18s	0.7185	1.92%	23s
MDRL	0.6881	1.67%	10s	0.6831	1.68%	17s	0.7209	1.60%	23s
EMNH	0.6900	1.40%	9s	0.6832	1.67%	16s	0.7217	1.49%	23s
PMOCO	0.6878	1.71%	9s	0.6819	1.86%	17s	0.7193	1.82%	23s
GIMF-P	0.6936	0.89%	11s	0.6878	1.01%	21s	0.7231	1.30%	26s
CNH	0.6947	0.73%	28s	0.6892	0.81%	37s	0.7250	1.04%	43s
GIMF-C	0.6958	0.57%	25s	0.6907	0.59%	35s	0.7285	0.56%	42s
MDRL-Aug	0.6950	0.69%	13s	0.6890	0.83%	19s	0.7261	0.89%	28s
EMNH-Aug	0.6958	0.57%	12s	0.6892	0.81%	18s	0.7270	0.76%	27s
PMOCO-Aug	0.6937	0.87%	12s	0.6886	0.89%	19s	0.7251	1.02%	27s
GIMF-P-Aug	0.6971	0.39%	15s	0.6924	0.35%	26s	0.7271	0.75%	45s
CNH-Aug	0.6980	0.26%	31s	0.6938	0.14%	37s	0.7303	0.31%	54s
GIMF-C-Aug	<u>0.6998</u>	<u>0.00%</u>	30s	<u>0.6948</u>	<u>0.00%</u>	39s	<u>0.7326</u>	<u>0.00%</u>	59s




Figure 6: Effects of the number of fusion layers

that the model performs best when $L' = 3$. If L' is too small, there may be insufficient fusion of multimodal information. Conversely, if L' is too large, the model’s ability to learn from single modalities may weaken, reducing its performance. Therefore, selecting an appropriate value for L' is essential to achieve a balance between multimodal fusion and single-modality learning, ultimately improving the model’s performance.

Effects of the number of bottlenecks For the number n_b (where we set $n'_b = n_b$) of bottlenecks introduced for multimodal fusion, the model performs best when n_b is set to 10, as shown in Figure 7. Introducing too many bottlenecks can increase the model’s complexity, making effective learning more difficult, while too few bottlenecks may hinder the proper fusion of multimodal information, limiting the model’s capabilities. Thus, selecting an appropriate number of bottlenecks ensures effective multimodal fusion and improves the model’s performance.

918
919
920
921
922
923
924
925
926
927
928
929
930
931
932
933
934
935
936
937
938
939
940
941
942
943
944
945
946
947
948
949
950
951
952
953
954
955
956
957
958
959
960
961
962
963
964
965
966
967
968
969
970
971




Figure 7: Effects of the number of bottlenecks