

iMTSP: Solving Min-Max Multiple Traveling Salesman Problem with Imperative Learning

Yifan Guo, Zhongqiang Ren, and Chen Wang[✉]

Abstract—This paper considers a Min-Max Multiple Traveling Salesman Problem (MTSP), where the goal is to find a set of tours, one for each agent, to collectively visit all the cities while minimizing the length of the longest tour. Though MTSP has been widely studied, obtaining near-optimal solutions for large-scale problems is still challenging due to its NP-hardness. Recent efforts in data-driven methods face challenges of the need for hard-to-obtain supervision and issues with high variance in gradient estimations, leading to slow convergence and highly sub-optimal solutions. We address these issues by reformulating MTSP as a bilevel optimization problem, using the concept of imperative learning (IL). This involves introducing an allocation network that decomposes the MTSP into multiple single-agent traveling salesman problems (TSPs). The longest tour from these TSP solutions is then used to self-supervise the allocation network, resulting in a new self-supervised, bilevel, end-to-end learning framework, which we refer to as imperative MTSP (iMTSP). Additionally, to tackle the high-variance gradient issues during the optimization, we introduce a control variate-based gradient estimation algorithm. Our experiments showed that these innovative designs enable our gradient estimator to converge $20\times$ faster than the advanced reinforcement learning baseline, and find up to 80% shorter tour length compared with Google OR-Tools MTSP solver, especially in large-scale problems (e.g. 1000 cities and 15 agents).

I. INTRODUCTION

The multiple traveling salesman problem (MTSP) seeks tours for multiple agents such that all cities are visited exactly once while minimizing an objective function defined over the tours. MTSP arises in numerous robotics topics that require a team of robots to collectively visit many target locations such as unmanned aerial vehicles path planning [1], [2], automated agriculture [3], warehouse logistics [4]. As their names suggest, Min-Sum MTSP minimizes the sum of tour lengths, while Min-Max MTSP minimizes the longest tour length, both of which are NP-hard [5] to solve to optimality.¹ In this paper, we focus on Min-Max MTSP while our method will also be applicable to Min-Sum MTSP. Min-Max MTSP

is more often used when the application seeks to minimize the overall completion time [7]. Due to the NP-hardness of MTSP, one must consider the trade-off between the solution optimality and the computational efficiency, since finding an optimal solution is often intractable for large problems.

A variety of classic (non-learning) approaches, including exact [8], [9], approximation [10], and heuristic [2], [11], [12] algorithms, have been developed to handle MTSP, but most of them consider minimizing the sum of tour lengths (Min-Sum), rather than the maximum tour length (Min-Max). In recent years, there has been a notable shift towards employing deep learning techniques to tackle TSP and MTSP [13]–[22]. However, these methods still have fundamental limitations, particularly in their ability to generalize to unseen problem sizes, and in consistently finding high-quality solutions for large-scale problems. For deep learning-based methods that supervise the model with heuristic or exact solutions [15]–[17], they struggle with limited supervision on small-scale problems and lack feasible supervision for large-scale instances, leading to poor generalization. Reinforcement learning (RL)-based approaches [14], [18]–[20] usually exploit implementations of the policy gradient algorithm, such as the REINFORCE [23] algorithm and its variants. These RL approaches face the issue of high-variance gradient estimations, which can result in slow convergence and highly sub-optimal solutions. Researchers have also developed strategies like training greedy policy networks [18], [19], [21] or iteratively improving solutions [22], which often get stuck at local optima.

To enhance generalization and accelerate the training process, we reformulate the Min-Max MTSP as a bilevel optimization problem. This comprises an upper-level optimization that focuses on training a city allocation network, and a lower-level optimization that solves multiple TSPs, inspired by imperative learning (IL) [24]–[26]. Specifically, the allocation network is to assign each city to an agent, which decomposes the MTSP into multiple single-agent TSPs. Then, a classic TSP solver, which can quickly produce a near-optimal tour for a single-agent TSP with up to a few hundred cities, is employed to find a tour for each agent based on the allocation. A key aspect is that the upper-level optimization, namely the training of the allocation network, is supervised by the lower-level TSP solver. This metric-based supervision from the lower-level optimization results in an end-to-end, self-supervised framework for solving MTSP, which we refer to as imperative MTSP (iMTSP).

One technical obstacle when developing iMTSP is that the space of possible city allocations is discrete. As a result,

[✉]Corresponding author. E-mail: chenw@sairlab.org. The source code, pre-trained model, and relevant data of this paper will be released at <https://github.com/sair-lab/iMTSP>

Yifan Guo is with the Flight Dynamics & Control/Hybrid Systems Lab, Purdue University, West Lafayette, IN 47907. Email: guo781@purdue.edu

Zhongqiang Ren is with Robotics Institute, Carnegie Mellon University, Pittsburgh, PA 15213, USA. Email: zhongqir@andrew.cmu.edu

Chen Wang is with Spatial AI & Robotics (SAIR) Lab, Department of Computer Science and Engineering, University at Buffalo, NY 14260, USA.

¹MTSP is challenging and renowned for its NP-hardness [5]. Intuitively, MTSP involves many decision variables: both assigning the cities to the agents and determining the visiting order of the assigned cities. As an example, an MTSP with 100 cities and 10 agents involves $\frac{100! \times 99!}{10! \times 89!} \approx 10^{20000}$ possible solutions, while the number of atoms in the observable universe is estimated to be “merely” in the range from 10^{78} to 10^{82} [6].

the objective function, i.e., the maximum tour length given by the lower-level TSP solver, is non-differentiable to the city allocation, which prevents the back-propagation to the allocation network’s parameters. One of the solutions is to explore the policy gradient algorithm or the reparameterization trick [27], which often use Monte-Carlo methods to estimate the gradient. However, these methods can lead to gradient estimations with high variance since the space of city allocation is extremely large. To handle this difficulty, we introduce a control variate [28] to reduce the variance of the gradient estimations. Parameterized by a trainable surrogate network running in parallel with the TSP solver, the control variate can be efficiently computed, and the overall algorithm can provide low variance gradient information through the non-differentiable solver and the discrete decision space. The main contributions of this paper are summarized as follows.

- 1) **Framework:** We formulate the MTSP as a bilevel optimization problem, introduce a new end-to-end, self-supervised framework for Min-Max MTSP. This decomposes a large-scale Min-Max MTSP into several smaller TSPs, each of which can be solved efficiently.
- 2) **Methodology:** To tackle the problem of bilevel optimization in discrete space, we introduce a control variate-based technique for back-propagation. This produces low-variance gradient estimations despite the non-differentiability of the lower-level optimization problem.
- 3) **Experiments:** We corroborate iMTSP in terms of efficiency and solution quality against both an RL-based and a classic (non-learning) methods. On large-scale problems, iMTSP achieves at most 80% shorter tours and with only 1.6% inference time compared with Google OR-Tools routing library [29]. Also, iMTSP’s gradient estimator converges about $20\times$ faster and produces $3.2 \pm 0.01\%$ shorter tours, compared with the RL-based approach [14].

II. RELATED WORKS

MTSP has been extensively studied. Analytical methods [8] use graph theory, game theory, and combinatorial optimization, etc. [30], and can provide guarantees on solution optimality. However, these methods usually struggle at solving large-scale MTSP due to its NP-hardness.

Heuristic approaches for MTSP handle the challenge by decomposing the original problem into several phases which in turn reduces the computational burden of each phase. Examples are particle swarm optimization [31], ant colony optimization [12], etc [32], [33].

Although heuristic algorithms usually run faster than analytical approaches, they lack theoretic guarantees on solution quality and could produce highly sub-optimal solutions especially for large-scale problems with hundreds and thousands of cities.

Recently, researchers tried to solve MTSP with machine learning (ML) methods using recurrent neural networks (RNN) [34], transformer [35], etc. These learning-based methods usually scale well as the size of the problem increases. However, they may not generalize well to problem sizes outside the training set. As an early attempt [17], an

RNN based method is developed and can handle TSP with up to 50 cities. Hu et al. [14] propose a hybrid approach within the RL framework, which shows impressive generalization ability. However, their gradient estimator has high variance, which leads to slow convergence and sub-optimal solutions especially when dealing with large-scale problems. Other ML-based methods include [18], [19], [21], [22]. Note that most of them consider TSP instead of MTSP.

Imperative Learning As an emerging learning framework, imperative learning (IL) has inspired some pioneering works in simultaneous localization and mapping (SLAM) [24], path planning [25], and visual feature matching [26]. The most significant difference between our work and theirs is that the lower-level optimization problems in their scenarios are all in continuous space, while iMTSP needs to deal with an extremely large discrete decision space.

Control Variate Control variate has been investigated by the ML community. In general, any policy gradient algorithm with a baseline or an actor-critic can be interpreted as an additive control variate method [36]. Additionally, Grathwohl et al. [37] propose gradient estimators for a non-differentiable function with continuous and discrete inputs, respectively. Their insight is evaluating the control variate directly with a single sample estimation of the gradient variance. Our work leverages this insight. Gu et al. [38] produce a family of policy gradient methods using control variate to merge gradient information from both on-policy and off-policy updates. Control variate technique has also been used for meta-RL [39], federated learning [40], etc [41]–[43].

III. PROBLEM FORMULATION

Given a set of N cities, each represented by a 2-D coordinate $\mathbf{x}_i \in \mathbb{R}^2$ and collectively denoted as $X = \{\mathbf{x}_1, \dots, \mathbf{x}_N\}$, along with M agents starting from a common city, i.e., the depot $\mathbf{x}_d \in X$, the MTSP is characterized by the tuple (X, \mathbf{x}_d, M) . The objective is to find, for each agent j , a tour $T_j = \{\mathbf{x}_{j,1}, \dots, \mathbf{x}_{j,k}, \dots\}$ such that all cities are visited exactly once, each tour returns to the depot, and the length of the longest tour reaches the minimum. The overall optimization objective for Min-Max MTSP can be defined as:

$$\min_{T_1, \dots, T_M} \max_j D(T_j), \quad (1)$$

where the tour length D is the distance between two cities:

$$D(T_j) = \sum_{k=1}^{\|T_j\|} \|\mathbf{x}_{j,k+1} - \mathbf{x}_{j,k}\|_2, \quad j = 1, 2, \dots, M, \quad (2)$$

where $\mathbf{x}_{j,k}$ denotes the city coordinate that is the k -th destination of the j -th agent.

IV. METHODS

This section first elaborates on the reformulation of MTSP as a bilevel optimization problem, and then introduces the forward computation of the our framework and the control variate-based optimization algorithm.




Fig. 1: The framework of our self-supervised MTSP network. The allocation network uses supervision from the TSP solver, and the surrogate network is supervised by the single sample variance estimator, reducing the gradient variance.

A. Reformulation of MTSP

The MTSP formulation in Section III is reformulated as the following bilevel optimization problem: For each instance (X, \mathbf{x}_d, M) , the upper-level problem is to assign each city to an agent, and the lower-level problem is to compute the visiting order of assigned cities for each agent. We define f as the allocation network with parameters θ , and g as the single-agent TSP solver with parameters μ . The imperative MTSP (iMTSP) is defined as:

$$\min_{\theta} U_{\theta}(L(\mu^*), f(\theta); \gamma); \quad \min_{\gamma} U_{\gamma}(L(\mu^*), s(\gamma); \theta) \quad (3a)$$

$$\text{s.t. } \mu^* = \arg \min_{\mu} L(\mu), \quad (3b)$$

$$L \doteq \max_j D(g(a_j; \mu)), \quad j = 1, 2, \dots, M, \quad (3c)$$

$$a_j \sim f(\theta), \quad (3d)$$

where a_j is the allocation for agent j sampling the probability matrix by row, and L returns the maximum route length among all agents. s is a surrogate network with parameter γ , which will be explained in detail in Section IV-C along with the specific definitions of upper-level objectives. Specifically, the upper-level optimization consists of two separate steps: the upper-level cost U_{θ} is to optimize the allocation network $f(\theta)$, while U_{γ} is to optimize the surrogate network $s(\gamma)$. The framework of iMTSP is shown in Fig. 1.

B. Forward Propagation Details

We next present the details of solving iMTSP (3). For each instance with coordinates X , we first embed agents and cities into feature vectors that contain the agent-city and city-city topological information. The embedded instance X' is the input to the allocation network $P = f(X', \mathbf{x}_d, M; \theta)$, where θ is the network parameters. Each element $p_{i,j}$ in the output $P \in \mathbb{R}^{N \times M}$ is the probability that the i -th city is allocated to the j -th agent. This probability matrix is then sampled in each row so that each city is assigned to an agent. The sampling result is then reorganized into allocations $A = \{a_1, a_2, \dots, a_M\}$, where a_j represents the set of cities assigned to the j -th agent. Subsequently, a TSP solver $g(a_j; \mu)$ for each agent j is invoked to order the cities in the allocation a_j and return a tour. Finally, the cost of each tour is computed with (2) and the length of the longest tour is the cost of that instance, i.e., $L = \max_j D(g(a_j; \mu))$. The

internal process of embedding and the allocation network is below.

1) *Embedding*: To extract the topological information from the coordinate data, we use a compositional message passing neural network (CMPNN) [44] to embed each city and agent as a feature vector. First, the feature of the city i is iteratively updated using its feature \mathbf{f}_i^t , features of its neighbors \mathbf{f}_k^t , and the weights $\mathbf{e}_{i,k}$ on the connecting edges. Here, i and k indicate different cities, and t is the number of iterations of graph aggregation (instead of the training loop). The embedding of the entire graph \mathbf{f}_g comes from all cities' features. Finally, a context feature vector \mathbf{f}_c is formed as a concatenation of the graph feature \mathbf{f}_g and the feature of the depot \mathbf{f}_d . Formally speaking, the embedding process is:

$$\mathbf{m}_{i,k}^t = \Omega(\mathbf{f}_i^t, \mathbf{f}_k^t, \mathbf{e}_{i,k}), \quad (4a)$$

$$\mathbf{l}_i^t = \Phi_{k \in \mathcal{N}_i}(\mathbf{m}_{i,k}^t), \quad (4b)$$

$$\mathbf{f}_i^{t+1} = \Psi(\mathbf{f}_i^t, \mathbf{l}_i^t), \quad (4c)$$

$$\mathbf{f}_g^t = \Phi'(\mathbf{f}_i^t), \quad i = 2, 3, \dots, n, \quad (4d)$$

$$\mathbf{f}_c = [\mathbf{f}_g; \mathbf{f}_d], \quad (4e)$$

where functions Ω and Ψ are represented by neural networks; Φ and Φ' are element-wise nonlinear functions; \mathcal{N}_i contains all the neighbours of city i ; $\mathbf{m}_{i,k}^t$ is the processed information from city k to city i ; and \mathbf{l}_i^t represents aggregated information from all the neighbours of city i . After obtaining the context feature \mathbf{f}_c , the embedding of agent j is calculated using the attention mechanism as follows:

$$\mathbf{k}_{j,i} = \theta_{\mathbf{k}} \mathbf{f}_i, \quad \mathbf{q}_j = \theta_{\mathbf{q}} \mathbf{f}_c, \quad i = 2, 3, \dots, n \quad (5a)$$

$$\mathbf{v}_{j,i} = \theta_{\mathbf{v}} \mathbf{f}_i, \quad i = 2, 3, \dots, n \quad (5b)$$

$$u_{j,i} = \frac{\mathbf{q}_j^T \mathbf{k}_{j,i}}{\sqrt{d_{\mathbf{k}}}}, \quad i = 2, 3, \dots, n \quad (5c)$$

$$w_{j,i} = \frac{e^{u_{j,i}}}{\sum_k e^{u_{j,k}}}, \quad i = 2, 3, \dots, n; \quad k = 2, 3, \dots, n \quad (5d)$$

$$\mathbf{h}_j = \sum_i w_{j,i} \mathbf{v}_{j,i}, \quad (5e)$$

where \mathbf{k} , \mathbf{q} and \mathbf{v} are respectively the keys, queries, and values of the attention layer. $d_{\mathbf{k}}$ is the dimension of keys, and i and k are used to index cities. All the θ are trainable parameters. The keys and values are first computed from the city embedding, and the queries are from the context feature \mathbf{f}_c . Then, we calculate the attention weight $w_{j,i}$ of each agent-city pair. Finally, the agent embedding \mathbf{h}_j is the sum of all the city values, weighted by the attention weights.

2) *Allocation Network*: As we have extracted the topological information into vector representations, we can now calculate the allocation matrix using the attention mechanism again. The keys \mathbf{k}' and queries \mathbf{q}' are computed from the city and agent embeddings, respectively, and the relative

importance of each city to each agent is computed as:

$$\mathbf{k}'_{j,i} = \theta_{\mathbf{k}'} \mathbf{f}_i, \quad \mathbf{q}'_j = \theta_{\mathbf{k}'} \mathbf{h}_j, \quad i = 2, 3, \dots, n \quad (6a)$$

$$u'_{j,i} = \frac{\mathbf{q}'_j{}^\top \mathbf{k}'_{j,i}}{\sqrt{d_{\mathbf{k}'}'}}, \quad i = 2, 3, \dots, n \quad (6b)$$

$$\beta_{j,i} = \alpha \tanh(u'_{j,i}), \quad (6c)$$

where $d_{\mathbf{k}'}$ is the dimension of new keys. Note that queries only depend on the agent embedding but not the city's. The notation $u'_{j,i}$ represents the non-clipped importance scores, α is a clip constant, and $\beta_{j,i}$ is the clipped importance score of city i to agent j , which will be used to decide which agent is assigned to visit the city. The probability of the city i is visited by agent j is calculated by applying the softmax function on the importance over all agents:

$$p_{j,i} = \frac{e^{\beta_{j,i}}}{\sum_j e^{\beta_{j,i}}}. \quad (7)$$

The allocation matrix has a shape of $N \times M$ with N and M are the number of cities and agents, respectively. The model produces the final allocation by sampling each row of the allocation matrix. With the allocation, the MTSP is decomposed into M TSPs. A classic TSP solver (such as Google OR-Tools) is invoked on each of the TSP to find the tour length of each TSP. The final return L of the MTSP is the maximum tour length among all the TSPs. Because the TSP solver is non-differentiable and the allocation variables are discrete, network optimization is very challenging.

C. Optimization

Now we demonstrate how the non-differentiability prevents us from using the analytical gradient, and how our control-variate approach can pass low-variance gradient estimations through the classic solver and the discrete allocation space.

1) *Gradient*: Directly taking the derivative of the objective $L = \max_j D(g(a_j; \mu))$ w.r.t. the neural parameter θ , we have the analytical form of the gradient:

$$\nabla_\theta L = \frac{\partial L}{\partial f} \frac{\partial f}{\partial \theta} + \frac{\partial L}{\partial g} \frac{\partial g}{\partial \mu^*} \frac{\partial \mu^*}{\partial \theta}. \quad (8)$$

There are two major difficulties prevent us from directly deploying this equation to update θ . First, the term $\frac{\partial L}{\partial f} \frac{\partial f}{\partial \theta}$ is always zero since the objective L does not directly depend on f . Thus, this term cannot provide any information to update θ . The second term cannot be directly computed because: (i) for $\frac{\partial L}{\partial g}$, the max function used to select the longest tour length is not differentiable; (ii) for $\frac{\partial g}{\partial \mu^*}$ and $\frac{\partial \mu^*}{\partial \theta}$, the optimization process in $g(\cdot)$ is not differentiable; and (iii) μ^* and θ related implicitly. To deal with the non-differentiability, we can use the Log-derivative trick, which has been widely used in RL, and estimates the above gradient as

$$\nabla_\theta L = L \frac{\partial}{\partial \theta} \log P(\theta). \quad (9)$$

However, our experiments showed that this estimator provides a gradient with a high variance, which slows down the convergence rate and compromises the solution quality.

Algorithm 1 The iMTSP algorithm

Require: The coordinate data X and embedded data X' , allocation network $f(X'; \theta)$, TSP solver $g(A; \mu)$, surrogate network $s(P(\theta); \gamma)$, learning rates α_1, α_2

while not converged **do**

$P \leftarrow f(X'; \theta)$

$A \leftarrow \text{Sampling } P \text{ by row}$

$C \leftarrow g(A; \mu), C' \leftarrow s(P(\theta); \gamma)$

$\frac{\partial}{\partial \theta} U_\theta \leftarrow [L - L'] \frac{\partial}{\partial \theta} \log P(\theta) + \frac{\partial}{\partial \theta} s(P(\theta); \gamma)$

$\frac{\partial}{\partial \phi} U_\gamma \leftarrow \frac{\partial}{\partial \gamma} (\frac{\partial}{\partial \theta} U_\theta)^2$

$\theta \leftarrow \theta - \alpha_1 \frac{\partial}{\partial \theta} U_\theta$

$\gamma \leftarrow \gamma - \alpha_2 \frac{\partial}{\partial \gamma} U_\gamma$

end while

return θ

2) *Control Variate*: To address the above issue, we introduce a control variate [28] to reduce the variance of a stochastic estimator. Suppose d (as in derivative) is a random variable (RV), $h(d)$ is an estimator of $\mathbb{E}(d)$, the control variate $c(d)$ is a function of d whose mean value is known. Define a new estimator with a control variate as:

$$h(d)_{\text{new}} = h(d) + \gamma(c(d) - \mathbb{E}[c(d)]). \quad (10)$$

When constant γ is properly chosen, $h(d)_{\text{new}}$ has the same expected value but lower variance than $h(d)$, as long as the control variate $c(d)$ is correlated to $h(d)$. We briefly summarize the underlying principle of control variate technique [45] as follows and omit the dependency on d for simplicity.

Lemma 1: Let h denotes an RV with an **unknown** expected value $\mathbb{E}(h)$, and c denotes an RV with a known expected value $\mathbb{E}(c)$. The new RV $h_{\text{new}} = h + \gamma(c - \mathbb{E}(c))$ has same expected value but smaller variance, when the constant γ is properly chosen and c is correlated with h .

Proof: The variance of the new estimator is

$$\text{Var}(h_{\text{new}}) = \text{Var}(h) + \gamma^2 \text{Var}(c) + 2\gamma \text{Cov}(h, c). \quad (11)$$

Differentiate the above equation w.r.t. γ and set the derivative to zero, the optimal choice of γ is

$$\gamma^* = -\frac{\text{Cov}(h, c)}{\text{Var}(c)}. \quad (12)$$

The minimum variance of h_{new} with γ^* is $(1 - \rho_{h,c}^2) \text{Var}(h)$ where ρ is the correlation coefficient of h and c . Since $\rho \in [-1, 1]$, the inequality $\text{Var}(h_{\text{new}}) \leq \text{Var}(h)$ always holds with $\gamma = \gamma^*$. When $\|\rho\| \rightarrow 1$, $\text{Var}(h_{\text{new}}) \rightarrow 0$. ■

In practice, γ^* is usually inaccessible, and thus we cannot arbitrarily choose c . Without losing generality, consider the case where h and c are positively correlated. By solving the inequality $\gamma^2 \text{Var}(c) + 2\gamma \text{Cov}(h, c) < 0$, we get $\gamma \in (\frac{-2\text{Cov}(h, c)}{\text{Var}(c)}, 0)$, which indicates that as long as k is set in this range, $\text{Var}(h_{\text{new}})$ is less than $\text{Var}(h)$. When h and c are highly correlated and are of similar ranges, $\text{Cov}(h, c) \approx \text{Var}(c)$, thus $\gamma^* \approx -1$, which is our empirical choice of γ .

Inspired by the control variant gradient estimator (10) and to have a lower gradient variance, we design the upper-level

TABLE I: Performance comparison on test MTSPs. The boldface denotes the best results on each training size. iMTSP provides solutions with overall $3.2 \pm 0.01\%$ shorter maximum route length comparing with the RL baseline, and has evidently advantages over the Google OR-Tools. The last column indicates how much longer are the baseline’s routes than iMTSP’s routes, averaging on all test size.

Training Setting			# Test Cities							Avg gap
Model	# Agents	# Cities	400	500	600	700	800	900	1000	
Google OR-Tools [29]	10	N.A.	10.482	10.199	12.032	13.198	14.119	18.710	18.658	277.1%
Reinforcement Learning [14]	10	50	3.058	3.301	3.533	3.723	3.935	4.120	4.276	0.2%
iMTSP (Ours)	10	50	3.046	3.273	3.530	3.712	3.924	4.122	4.283	0
Google OR-Tools [29]	10	N.A.	10.482	10.199	12.032	13.198	14.119	18.710	18.658	290.1%
Reinforcement Learning [14]	10	100	3.035	3.271	3.542	3.737	3.954	4.181	4.330	4.3%
iMTSP (Ours)	10	100	3.019	3.215	3.424	3.570	3.763	3.930	4.042	0
Google OR-Tools [29]	15	N.A.	10.293	10.042	11.640	13.145	14.098	18.471	18.626	349.9%
Reinforcement Learning [14]	15	100	2.718	2.915	3.103	3.242	3.441	3.609	3.730	6.3%
iMTSP (Ours)	15	100	2.614	2.771	2.944	3.059	3.221	3.340	3.456	0

costs U_θ for the allocation network as:

$$U_\theta \doteq [L(\mu^*) - L'] \log P(\theta) + s(P(\theta); \gamma), \quad (13)$$

where L' is output of the surrogate network (control variate) but is taken as a constant by cutting its gradients. In turn, our gradient for the allocation network parameters θ is:

$$\frac{\partial U_\theta}{\partial \theta} = [L(\mu^*) - L'] \frac{\partial}{\partial \theta} \log P(\theta) + \frac{\partial}{\partial \theta} s(P(\theta); \gamma). \quad (14)$$

3) *Surrogate Network*: We implement the idea of control variate with a surrogate network $s(\gamma)$ in parallel with the TSP solver. Note that the surrogate network takes the allocation matrix P as its input and predicts the maximum tour length, while the inputs of the actual TSP solver are allocations A and coordinate data X . Naively optimizing the surrogate network to imitate the input-output map of the TSP solver is ineffective because learning such a map does not guarantee learning correct gradient information. Since we require the surrogate network (control variate) to reduce the gradient variance of allocation parameters, we directly optimize the surrogate parameters to minimize a single sample estimation of the allocation gradient variance. Thus, the upper-level objective U_γ for the surrogate network is

$$U_\gamma \doteq \text{Var}(\frac{\partial U_\theta}{\partial \theta}). \quad (15)$$

Using the fact $\text{Var}(\frac{\partial U_\theta}{\partial \theta}) = \mathbb{E}[(\frac{\partial U_\theta}{\partial \theta})^2] - \mathbb{E}[\frac{\partial U_\theta}{\partial \theta}]^2$, the gradient of the surrogate network is

$$\frac{\partial U_\gamma}{\partial \gamma} = \frac{\partial}{\partial \gamma} \mathbb{E}[(\frac{\partial U_\theta}{\partial \theta})^2] - \frac{\partial}{\partial \gamma} \mathbb{E}[\frac{\partial U_\theta}{\partial \theta}]^2 = \frac{\partial}{\partial \gamma} \mathbb{E}[(\frac{\partial U_\theta}{\partial \theta})^2], \quad (16)$$

The term $\frac{\partial}{\partial \gamma} \mathbb{E}[(\frac{\partial U_\theta}{\partial \theta})^2]$ is always zero because the surrogate network has no explicit influence on the return L . For ease of understanding, we summarize our control variate-based bilevel optimization method in Algorithm 1.

V. EXPERIMENTS

We next conduct comprehensive experiments to demonstrate the generalization ability, computational efficiency, and convergence speed of our iMTSP framework.

TABLE II: Computing time of iMTSP dealing with 1000 cities. Due to a similar network structure, the RL baseline [14] has roughly the same runtime efficiency compared with ours, while Google OR-Tools [29] cannot find feasible solutions within a similar period.

Number of agents	5	10	15
Computing time	4.85	1.98	1.35

A. Baselines

To demonstrate the advantages of our method, we compare it with both classic methods and learning-based methods. Specifically, we select a state-of-the-art reinforcement learning (RL)-based model from [14] as a baseline. Their framework can efficiently solve MTSPs with up to a few thousand cities. Using S-batched REINFORCE, it has shown outstanding generalization ability and has a reduced variance for the gradient estimator than the vanilla REINFORCE [23]. The other potential baseline methods [13], [20], [46] either don’t provide source codes or cannot be directly applied to our settings. We also compare iMTSP with a well-known meta-heuristic method implemented by the Google OR-Tools routing library [14], [18], [22]. It provides near-optimal solutions to MTSPs with a few hundred cities and is a popular MTSP baseline [29]. For abbreviation, we will refer to the two methods as an RL baseline and Google OR-Tools.

B. Experiment Setting

We learn from [17], [18], [22] that, existing classic and learning-based approaches can handle MTSPs with about 150 cities, but their performance will significantly compromise with 400 or more cities. To demonstrate iMTSP’s ability to generalization and handle large-scale problems, we conduct experiments on training sets with 50 to 100 cities but test the models on problems with 400 to 1000 cities. We believe this challenging setting can reflect the generalization ability of the proposed models.

In our tests, all the instances are generated by uniformly sampling N coordinates in a unit rectangular so that both x and y coordinates are in the range of 0 to 1. We employ a batch size of 512 and a training-to-validation data size ratio of 10 : 1. The test set consists of i.i.d. batches but with a larger number of cities (400 to 1000) as aforementioned.

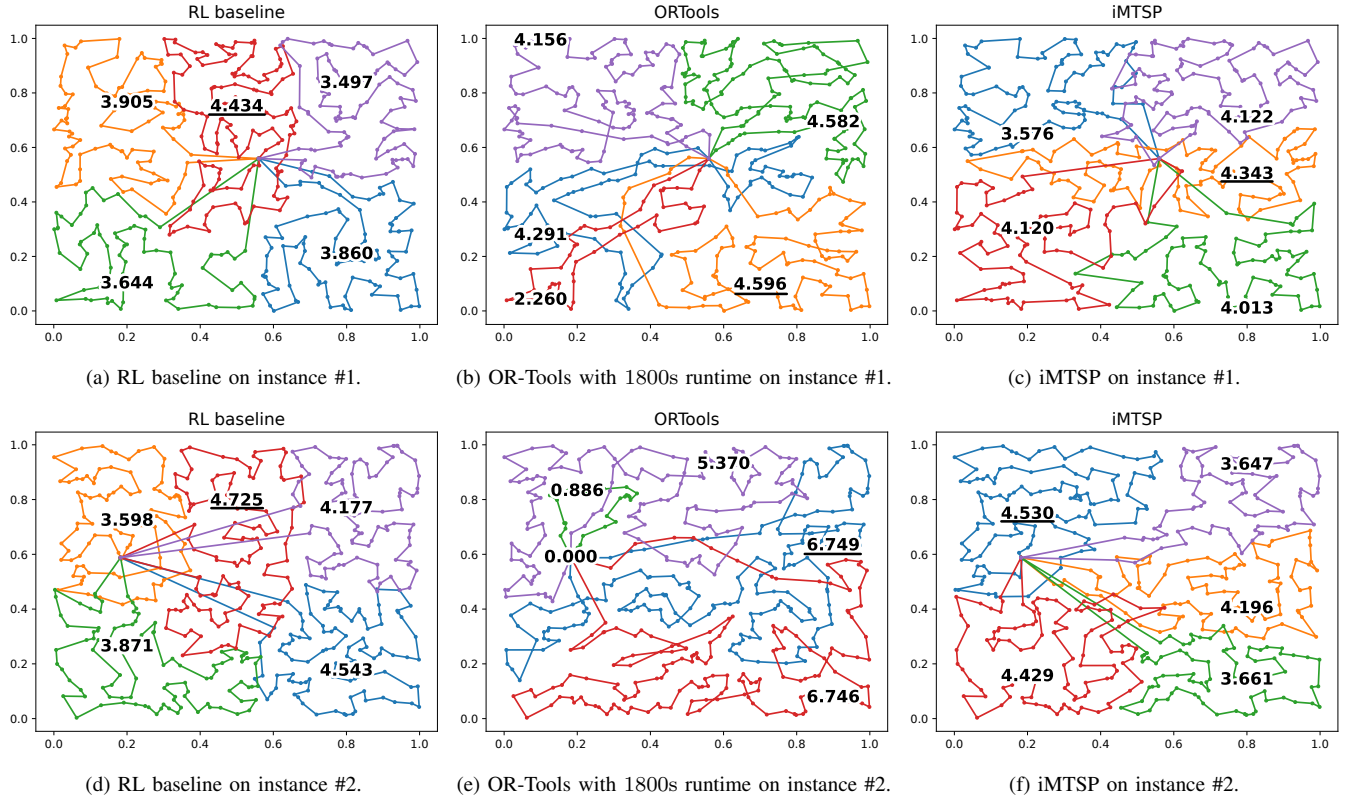


Fig. 2: The figures visualize the performance of the baselines and the proposed model on two example MTSPs with 5 agents and 500 cities, where the first instance has a central depot and the second instance has an off-center depot. The numbers denote the length of the route, and different colors represent different agents. Longest tour lengths are underlined in each figure. iMTSP always produces the best solution and has fewer sub-optimal patterns like circular partial routes or long straight partial routes.

Specifically, agents and cities are embedded into 64 dimensional vectors. The allocation network contains one attention layer, where keys, queries, and values are of 128, 16, and 16 dimensions, respectively. The clip constant α in (6c) is set to be 10. The surrogate network consists of three fully connected layers, whose hidden dimension is 256 and activation function is \tanh . Google OR-Tools provides the TSP solver in iMTSP and the MTSP solver as a baseline. It uses the “Global Cheapest Arc” strategy to generate its initial solutions and uses “Guided local search” for local improvements. All the strategies are defined in [29] and we selected the best options in the experiments.

C. Quantitative Performance

We next conduct a quantitative performance comparison with the two baseline methods and show it in Table I. We focus on large-scale MTSPs since many approaches can well solve small-size MTSPs but very few algorithms (both classic or learning-based) can deal with large-scale problems.

1) *Comparison with Google OR-Tools*: By decomposing a large MTSP into several small TSPs, iMTSP is stronger to handle large-scale problems. We demonstrate such a merit by comparing our model with the classic routing module in OR-Tools, which is given 300 seconds time budget. The results are shown in Table I. Specially, these route lengths of OR-Tools are averaged over 10 test instances. OR-Tools appears

to be very inferior because it cannot converge to a local minimum given 300 seconds time budget. In our experiments, OR-Tools can handle 500 or less cities with 1800 seconds budget. However, it does not converge even computing for 7200 seconds with 600 or more cities, indicating that its computational complexity grows steeply w.r.t. the problem size. Also, the difference of the route length increases as the problem size becomes larger and larger. Remember that the final routes of iMTSP are also solved by Google routing module, while it can easily handle MTSP with up to a few thousand cities with the help of the allocation network.

2) *Comparison with the RL Baseline*: To show the effectiveness of our control variate-based optimization algorithm, we also compare iMTSP with the RL baseline [14] which shares a similar forward process as iMTSP. In most cases, training on bigger problems leads to a better generalization ability on large-scale problems, but it also indicates larger parameter space, more complicated gradient information, and longer training time. This makes the quality of the gradient very crucial. As presented in the Table I, iMTSP outperforms the RL baseline in most cases, providing solutions with $3.2 \pm 0.01\%$ shorter maximum route length over the RL baseline on average. When the models are trained with 100 cities, the difference between the route lengths monotonically increases from 0.4% to 8.0% and from 3.4% to 8.9%, respectively with 10 and 15 agents. The results demonstrate that with




Fig. 3: The gradient variance history of our method and the RL baseline during the training process with 50 and 100 cities. The y -axis is the sum of the natural logarithm of the variance of all trainable parameters in the allocation network, and x -axis denotes the number of iterations. Our gradient estimator converges about $20\times$ faster than the RL baseline.

the lower-variance gradient provided by our control variate-based optimization algorithm, iMTSP usually converges to better solutions when being trained on large-scale instances.

D. Qualitative Performance

To find out the specific advantages in iMTSP’s solutions, we plot and compare the results of iMTSP and two baselines on two example instances. Fig. 2 provides a visualization of the solutions from both baselines and iMTSP on two instances with 5 agents and 500 cities. The depot of the first instance is close to the center of cities while it’s off the center in the second instance. In Fig. 2d, there are long straight partial routes (e.g., blue and purple), indicating some agents have travelled for a long distance before they visit their first destination. Also, some circular parts exist in the OR-Tools’ solutions (e.g., the green route in Fig. 2b) meaning that this agent travels for extra distance to visit the cities on the circle. Both of these are typical symptoms of sub-optimality. Such symptoms are not observed in the two solutions from iMTSP. In both instances, iMTSP produces the best performance. Note that both the RL baseline and iMTSP finish their computation within a few seconds while OR-Tools takes 1800 seconds to produce comparable solutions.

E. Efficiency Analysis

Since iMTSP contains both a learning-based allocation network and a classic TSP solver, it is important to identify the bottleneck of the architecture for future improvements. As in Table II, with 5, 10, and 15 agents, the computing time of our model are 4.85 seconds, 1.98 seconds, and 1.35 seconds, respectively, to solve one instance with 1000 cities. Note that the computing time decrease as the number of agents increases. This fact indicates that the major computational bottleneck is the TSP solvers rather than the allocation network, because more agents are required to run more times of TSP solvers. One of the promising solutions to further reduce our computing time is to create multiple threads to run the TSP solvers in parallel since the TSPs in the lower-level optimization of iMTSP are independent. Since this is not the focus of this paper, we put it as a future work.

TABLE III: Results of the ablation studies. The first row is the results from a model optimized by the baseline algorithm while the second line is our iMTSP approach. Our iMTSP approach always produces shorter maximum tour lengths.

Optimization	# Test Cities						
	400	500	600	700	800	900	1000
REINFORCE	3.200	3.477	3.779	4.007	4.249	4.488	4.699
Control Variate	3.046	3.273	3.530	3.712	3.924	4.122	4.283

F. Gradient Variance

As demonstrated in Section IV-C, with an optimized surrogate network, the variance of iMTSP’s gradient are expected to be smaller than the gradient from the vanilla REINFORCE algorithm [23]. We verify such a hypothesis by explicitly recording the mini-batch gradient variance during the training process. The experiment is conducted twice with 10 agents visiting 50 cities and 100 cities. Training data are divided into several mini-batches, each of which contains 32 instances. The gradient is computed and stored for every mini-batch, and the variance is later computed with these stored gradient. The network parameters are updated with the average gradient of the whole batch. The results are shown in Fig. 3, where the x -axis denotes the number of iterations and the y -axis represents the natural logarithm of the summed variance of all trainable parameters in the allocation network. It is observed that the gradient variance of our iMTSP is much smaller and converges $20\times$ faster than the RL baseline. This verifies the effectiveness of our control variate-based bilevel optimization process.

G. Ablation Study

To confirm that the improvements of solution quality are majorly contributed by the control variate-based optimization algorithm, we also track the performance of the fine-tuned network with the baseline optimization method. As shown in Table III, the route lengths from our control variate-based optimization algorithm are averagely $7.8\% \pm 2.1\%$ shorter than that from the baseline algorithm. We can also observe that better results always come from iMTSP, which indicates that the better solution quality is mainly contributed by our control variate-based optimization algorithm.

VI. CONCLUSION AND FUTURE WORK

In the paper, we reformulate the Min-Max MTSP as a bilevel optimization problem. We introduce a self-supervised framework iMTSP, which can efficiently handle large-scale MTSP with thousands of cities. With the control variate-based optimization algorithm, iMTSP produces a low-variance gradient through a non-differentiable TSP solver and a discrete allocation space to update the allocation network. Experimental results demonstrated the advantages of iMTSP in solution quality, computational efficiency, and optimization stability. In future work, we will also consider the constraints of environmental obstacles and inter-agent collision, which will enable iMTSP to conduct multi-agent path planning in real-world scenarios.

REFERENCES

- [1] K. Sundar and S. Rathinam, "Algorithms for routing an unmanned aerial vehicle in the presence of refueling depots," *IEEE Transactions on Automation Science and Engineering*, vol. 11, no. 1, pp. 287–294, 2013.
- [2] Y. Ma, H. Zhang, Y. Zhang, R. Gao, Z. Xu, and J. Yang, "Coordinated optimization algorithm combining GA with cluster for multi-UAVs to multi-tasks task assignment and path planning," in *International Conference on Control and Automation*. IEEE, 2019, pp. 1026–1031.
- [3] R. F. Carpio, J. Maiolini, C. Potena, E. Garone, G. Ulivi, and A. Gasparn, "Mp-stsp: A multi-platform steiner traveling salesman problem formulation for precision agriculture in orchards," in *International Conference on Robotics and Automation*. IEEE, 2021, pp. 2345–2351.
- [4] Z. Ren, S. Rathinam, and H. Choset, "CBSS: A new approach for multiagent combinatorial path finding," *IEEE Transactions on Robotics*, vol. 39, no. 4, pp. 2669–2683, 2023.
- [5] T. Bektas, "The multiple traveling salesman problem: an overview of formulations and solution procedures," *omega*, vol. 34, no. 3, pp. 209–219, 2006.
- [6] E. Gaztañaga, "The mass of our observable universe," *Monthly Notices of the Royal Astronomical Society: Letters*, vol. 521, no. 1, pp. L59–L63, 2023.
- [7] J. Park, C. Kwon, and J. Park, "Learn to solve the min-max multiple traveling salesmen problem with reinforcement learning," in *Proceedings of the 2023 International Conference on Autonomous Agents and Multiagent Systems*, 2023, pp. 878–886.
- [8] A. M. Ham, "Integrated scheduling of m-truck, m-drone, and m-depot constrained by time-window, drop-pickup, and m-visit using constraint programming," *Transportation Research Part C: Emerging Technologies*, vol. 91, pp. 1–14, 2018.
- [9] Z. Ren, S. Rathinam, and H. Choset, "MS: A new exact algorithm for multi-agent simultaneous multi-goal sequencing and path finding," in *International Conference on Robotics and Automation*. IEEE, 2021, pp. 11 560–11 565.
- [10] K. Bérczi, M. Mních, and R. Vincze, "Approximations for many-visits multiple traveling salesman problems," *Omega*, vol. 116, p. 102816, 2023.
- [11] M. A. Al-Omeir and Z. H. Ahmed, "Comparative study of crossover operators for the MTSP," in *2019 International Conference on Computer and Information Sciences*. IEEE, 2019, pp. 1–6.
- [12] R. Necula, M. Breaban, and M. Raschip, "Tackling the bi-criteria facet of multiple traveling salesman problem with ant colony systems," in *27th International Conference on Tools with Artificial Intelligence*. IEEE, 2015, pp. 873–880.
- [13] H. Liang, Y. Ma, Z. Cao, T. Liu, F. Ni, Z. Li, and J. Hao, "SplitNet: a reinforcement learning based sequence splitting method for the MinMax multiple travelling salesman problem," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 37, no. 7, 2023, pp. 8720–8727.
- [14] Y. Hu, Y. Yao, and W. S. Lee, "A reinforcement learning approach for optimizing multiple traveling salesman problems over graphs," *Knowledge-Based Systems*, vol. 204, p. 106244, 2020.
- [15] L. Xin, W. Song, Z. Cao, and J. Zhang, "NeuroLKH: Combining deep learning model with Lin-Kernighan-Helsgaun heuristic for solving the traveling salesman problem," *Advances in Neural Information Processing Systems*, vol. 34, pp. 7472–7483, 2021.
- [16] S. Miki, D. Yamamoto, and H. Ebara, "Applying deep learning and reinforcement learning to traveling salesman problem," in *International Conference on Computing, Electronics and Communications Engineering*. IEEE, 2018, pp. 65–70.
- [17] O. Vinyals, M. Fortunato, and N. Jaitly, "Pointer networks," *Advances in Neural Information Processing Systems*, vol. 28, 2015.
- [18] W. Kool, H. Van Hoof, and M. Welling, "Attention, learn to solve routing problems!" *arXiv preprint arXiv:1803.08475*, 2018.
- [19] M. Nazari, A. Oroojlooy, L. Snyder, and M. Takác, "Reinforcement learning for solving the vehicle routing problem," *Advances in Neural Information Processing Systems*, vol. 31, 2018.
- [20] J. Park, S. Bakhtiyar, and J. Park, "Schedulenet: Learn to solve multi-agent scheduling problems with reinforcement learning," *arXiv preprint arXiv:2106.03051*, 2021.
- [21] E. Khalil, H. Dai, Y. Zhang, B. Dilkina, and L. Song, "Learning combinatorial optimization algorithms over graphs," *Advances in Neural Information Processing Systems*, vol. 30, 2017.
- [22] Y. Wu, W. Song, Z. Cao, J. Zhang, and A. Lim, "Learning improvement heuristics for solving routing problems," *IEEE Transactions on Neural Networks and Learning Systems*, vol. 33, no. 9, pp. 5057–5069, 2021.
- [23] R. J. Williams, "Simple statistical gradient-following algorithms for connectionist reinforcement learning," *Machine Learning*, vol. 8, pp. 229–256, 1992.
- [24] T. Fu, S. Su, Y. Lu, and C. Wang, "iSLAM: Imperative SLAM," *IEEE Robotics and Automation Letters (RA-L)*, 2024.
- [25] F. Yang, C. Wang, C. Cadena, and M. Hutter, "iPlanner: Imperative path planning," in *Robotics: Science and Systems (RSS)*, 2023.
- [26] Z. Zhan, D. Gao, Y.-J. Lin, Y. Xia, and C. Wang, "iMatching: Imperative correspondence learning," *arXiv preprint arXiv:2312.02141*, 2023.
- [27] M. Figurnov, S. Mohamed, and A. Mnih, "Implicit reparameterization gradients," *Advances in Neural Information Processing Systems*, vol. 31, 2018.
- [28] B. L. Nelson, "Control variate remedies," *Operations Research*, vol. 38, no. 6, pp. 974–992, 1990.
- [29] L. Perron and F. Didier, "ORTools routing options," Google, 2023. [Online]. Available: https://developers.google.com/optimization/routing/routing_options
- [30] O. Cheikhrouhou and I. Khoufi, "A comprehensive survey on the multiple traveling salesman problem: Applications, approaches and taxonomy," *Computer Science Review*, vol. 40, p. 100369, 2021.
- [31] C. Wei, Z. Ji, and B. Cai, "Particle swarm optimization for cooperative multi-robot task allocation: a multi-objective approach," *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 2530–2537, 2020.
- [32] P. Kitjacharoenchai, M. Ventresca, M. Moshref-Javadi, S. Lee, J. M. Tanchoco, and P. A. Brunese, "Multiple traveling salesman problem with drones: Mathematical model and heuristic approach," *Computers and Industrial Engineering*, vol. 129, pp. 14–30, 2019.
- [33] C. C. Murray and R. Raj, "The multiple flying sidekicks traveling salesman problem: Parcel delivery with multiple drones," *Transportation Research Part C: Emerging Technologies*, vol. 110, pp. 368–398, 2020.
- [34] D. E. Rumelhart, G. E. Hinton, R. J. Williams *et al.*, "Learning internal representations by error propagation," 1985.
- [35] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, "Attention is all you need," *Advances in Neural Information Processing Systems*, vol. 30, 2017.
- [36] E. Greensmith, P. L. Bartlett, and J. Baxter, "Variance reduction techniques for gradient estimates in reinforcement learning," *Journal of Machine Learning Research*, vol. 5, no. 9, 2004.
- [37] W. Grathwohl, D. Choi, Y. Wu, G. Roeder, and D. Duvenaud, "Back-propagation through the void: Optimizing control variates for black-box gradient estimation," *arXiv preprint arXiv:1711.00123*, 2017.
- [38] S. S. Gu, T. Lillicrap, R. E. Turner, Z. Ghahramani, B. Schölkopf, and S. Levine, "Interpolated policy gradient: Merging on-policy and off-policy gradient estimation for deep reinforcement learning," *Advances in Neural Information Processing Systems*, vol. 30, 2017.
- [39] H. Liu, R. Socher, and C. Xiong, "Taming MAML: Efficient unbiased meta-reinforcement learning," in *International Conference on Machine Learning*. PMLR, 2019, pp. 4061–4071.
- [40] S. P. Karimireddy, S. Kale, M. Mohri, S. Reddi, S. Stich, and A. T. Suresh, "Scaffold: Stochastic controlled averaging for federated learning," in *International Conference on Machine Learning*. PMLR, 2020, pp. 5132–5143.
- [41] J. Baker, P. Fearnhead, E. B. Fox, and C. Nemeth, "Control variates for stochastic gradient MCMC," *Statistics and Computing*, vol. 29, pp. 599–615, 2019.
- [42] T. Geffner and J. Domke, "Using large ensembles of control variates for variational inference," *Advances in Neural Information Processing Systems*, vol. 31, 2018.
- [43] W. Guo, S. Wang, P. Ding, Y. Wang, and M. I. Jordan, "Multi-source causal inference using control variates," *arXiv preprint arXiv:2103.16689*, 2021.
- [44] J. Gilmer, S. S. Schoenholz, P. F. Riley, O. Vinyals, and G. E. Dahl, "Neural message passing for quantum chemistry," in *International Conference on Machine Learning*. PMLR, 2017, pp. 1263–1272.
- [45] S. Meyn, *Control techniques for complex networks*. Cambridge University Press, 2008.
- [46] H. Gao, X. Zhou, X. Xu, Y. Lan, and Y. Xiao, "AMARL: An attention-based multiagent reinforcement learning approach to the min-max multiple traveling salesmen problem," *IEEE Transactions on Neural Networks and Learning Systems*, 2023.