



# Multi-start team orienteering problem for UAS mission re-planning with data-efficient deep reinforcement learning

Dong Ho Lee<sup>1</sup> · Jaemyung Ahn<sup>1</sup>

Accepted: 26 February 2024  
© The Author(s) 2024

## Abstract

In this paper, we study the Multi-Start Team Orienteering Problem (MSTOP), a mission re-planning problem where vehicles are initially located away from the depot and have different amounts of fuel. We consider/assume the goal of multiple vehicles is to travel to maximize the sum of collected profits under resource (e.g., time, fuel) consumption constraints. Such re-planning problems occur in a wide range of intelligent UAS applications where changes in the mission environment force the operation of multiple vehicles to change from the original plan. To solve this problem with deep reinforcement learning (RL), we develop a policy network with self-attention on each partial tour and encoder-decoder attention between the partial tour and the remaining nodes. We propose a modified REINFORCE algorithm where the greedy rollout baseline is replaced by a local mini-batch baseline based on multiple, possibly non-duplicate sample rollouts. By drawing multiple samples per training instance, we can learn faster and obtain a stable policy gradient estimator with significantly fewer instances. The proposed training algorithm outperforms the conventional greedy rollout baseline, even when combined with the maximum entropy objective. The efficiency of our method is further demonstrated in two classical problems – the Traveling Salesman Problem (TSP) and the Capacitated Vehicle Routing Problem (CVRP). The experimental results show that our method enables models to develop more effective heuristics and performs competitively with the state-of-the-art deep reinforcement learning methods.

**Keywords** Deep reinforcement learning · Data-efficient training · Combinatorial optimization · Mission re-planning · Autonomous systems

## 1 Introduction

As the operational technology of Unmanned Aerial Systems (UAS) matures, there is a growing need for fast and accurate high-level decision-making for autonomous mission planning. The ability to adjust evolving mission objectives is essential for addressing the dynamic nature of real-world scenarios, enhancing safety, optimizing resources, and ensuring the success of missions across various industrial, civil, and defense sectors. For example, E-commerce drones for last-mile delivery should be able to optimize routes to

ensure timely service. Surveillance drones monitoring ongoing traffic require adaptive responses for optimal data collection. Several other potential applications include forest fire detection in emergency response, geographical monitoring for scientific research (where objectives may change based on initial findings), surveying and mapping for urban planning, airborne reconnaissance for border control, and search and rescue operations in disaster-stricken areas where UAS assist in locating and aiding survivors [1, 2].

Prior UAS mission studies addressed variants of the vehicle routing problem formulated as the NP-hard combinatorial optimization (CO), such as the Traveling Salesman Problem (TSP) and the Capacitated Vehicle Routing Problem (CVRP). These classical CO problems are primarily concerned with mission preplanning based on the current knowledge of the environment. However, missions in real life involve many unknown and possibly changing factors such as sudden gusts, GPS denial, unexpected threats, terrain uncertainties, fuel leakage, and hardware malfunction.

---

✉ Jaemyung Ahn  
jaemyung.ahn@kaist.ac.kr

Dong Ho Lee  
leedh0124@kaist.ac.kr

<sup>1</sup> Department of Aerospace Engineering, Korea Advanced Institute of Science and Technology (KAIST), Daejeon, South Korea

Once the vehicles have left the base, it is critical to respond to the unexpected environmental changes by managing mission objectives autonomously, thus prompting the need for near-optimal mission re-planning in real-time. Furthermore, visiting all nodes may not be practical considering resource availability. Instead, such applications may require vehicles to visit as many nodes as possible within a maximum duration given on each route. These characteristics of real-life applications give rise to the Multi-Start Team Orienteering Problem (MSTOP), which is a generalization of the Team Orienteering Problem (TOP) with additional degrees of freedom on launch location and available fuel for each vehicle. Many routing problems assume vehicles that identically begin routing from the depot. In contrast, MSTOP models the real-life mission re-planning scenario by launching vehicles located away from the depot, each with a different amount of fuel available.

The MSTOP is formulated in the context of route planning for intelligent UAS and robotic agent systems. Given the nature of higher-level decision making, more efficient route plans for optimal assignments among agents are desirable. For example, a fleet of UAVs suppressing forest fires needs an optimal order of visiting sites to make the most out of their limited volume of extinguishing water. The fleet may also be subject to frequently updating their assigned spots as wildfires can spread unpredictably, which calls for re-planning the routes. Another application is the efficient operation of unmanned delivery drones. If a delivery drone were to visit a number of sites to deliver multiple parcels, the order of sites to be visited can be optimized so that operational revenue is maximized. On top of that, a scheduled delivery site can be modified at the request of the customer, and the drones already in delivery require a new mission plan. In this manner, the MSTOP belongs to a general higher-level planning framework for a wide range of applications in the UAS and robotic systems.

Various traditional approaches have been applied to solve the CO problems so far. For example, exact algorithms are generally based on branch-and-bound or branch-and-cut approaches to obtain optimal solutions. However, finding an optimal solution may take an inordinate amount of time when the problem size grows. Approximate algorithms rapidly produce near-optimal solutions that are often tailored for specific CO problems. Heuristic approaches utilize domain expertise to design hand-crafted strategies for progressively constructing a solution. These approaches may not be straightforwardly applicable to other routing problems.

The deep reinforcement learning (RL) approach has recently emerged as a fast and powerful heuristic solver to find near-optimal solutions to many CO problems. This paper aims to develop a deep RL-based construction framework for solving the MSTOP. We propose a data-efficient training methodology that improves the solution quality


and learning speeds. To demonstrate the effectiveness of our training methodology, we experiment on two classical CO problems: TSP and CVRP. These experiments confirm that our training methodology outperforms the conventional methodology in [3] and is comparable to the state-of-the-art *policy optimization with multiple optima for reinforcement learning* (POMO) [4] while using significantly smaller data. In addition, we identify the asymmetry in the solution representation of MSTOP and use it to improve performance during inference further. With this advanced inference strategy, our model can generate high-quality solutions in a notably short time, bringing us a step closer to real-time mission re-planning.

In summary, our primary contributions are threefold. First, we explore the MSTOP, a routing problem that reflects a real-life mission re-planning scenario, using a data-driven method (deep reinforcement learning). Specifically, we follow the Transformer's encoder-decoder architecture [5]. We use a standard encoder with a multi-head attention mechanism. For the decoder, we adapt the decoding strategy in [6], the current state-of-the-art deep RL solver for a single vehicle TSP, and generalize the strategy to handle multiple vehicle launch locations. Our overall approach adopts the nested inner/outer loop framework similar to [7] that updates the current state after each vehicle returns to the depot to reflect the changes after a partial tour is complete. We name our neural policy network the Deep Dynamic Transformer Model (DDTM).

Second, we propose a data-efficient training approach based on a baseline derived from multiple instances generated by applying linear coordinate transformations to a single instance. These augmented instances are distinct in their raw form since each node in the 2D cartesian plane has been transformed. But, as a graph, these are identical because the lengths between the nodes are preserved. We replace the greedy rollout baseline with a local, mini-batch mean (obtained by rolling out all augmented instances) and combine it with the maximum entropy RL method [8, 9]. Our proposed methodology outperforms the computationally expensive greedy rollout baseline [3] and significantly expedites the learning process.

Finally, we improved the efficiency of the inference phase by using the instance-augmentation tailored for the MSTOP. Unlike TSP and CVRP, solutions to MSTOP are inherently asymmetric since the order of vehicles breaks the symmetry in the solution representations (see Fig. 1). We utilize the asymmetry in MSTOP solutions by permuting all vehicle orders and generating multiple rollouts for each permutation of vehicle order at the inference stage. This method is more efficient than the conventional sampling and instance-augmentation inference (using a single-vehicle order).

The remainder of this paper is as follows. Section II briefly introduces past studies related to our work (e.g., deep



**Fig. 1** Multiple representations for an optimal solution exist in TSP and CVRP. However, for MSTOP, the order of vehicles breaks the symmetry in solution representation

RL approaches for classical CO problems). Section III formulates the MSTOP as the Mixed Integer Linear Programming (MILP) and Markov Decision Process (MDP). Section IV describes our DDTM policy network in detail. Section V describes our proposed REINFORCE baseline and presents inference results on various routing problems. In Section VI, to corroborate the effectiveness of our method, we report an ablation study among several training baselines and present generalization results. Finally, Section VII concludes the paper and discusses future research directions.

## 2 Literature review

The Team Orienteering Problem (TOP) belongs to the broader Vehicle Routing Problem with Profits (VRPP) class. A fleet of vehicles is given, but the vehicles are not required to visit all the nodes or customers. Each node is associated with a prize (profit), denoting its relative attractiveness. The objective is to find a subset of nodes that maximizes the total collected profits while satisfying a limit on the maximum duration of each route [10–12]. Exact algorithms to solve the TOP include approaches based on column generation and constraint branching [13] and branch-and-price algorithm [14]. Taking the TOP as a basis, we devise the MSTOP by extending it with two additional degrees of freedom: launch locations of vehicles and remaining fuel for each vehicle. The MSTOP stands in contrast to traditional CO problems in that the launch locations for each vehicle are distinct. Therefore, the problem state seen by each vehicle is naturally different at each construction step [15]. It is also important to note that the context of multi-start in MSTOP stands in contrast with a number of existing works sharing the same term. For example, Lin et al. [16, 17] uses the term multi-start to refer to a variant of simulated annealing approach they used to solve TOP, while Hapsari et al. [18] deals with multi-objective TOP.

One of the early attempts to apply the deep RL approach to CO in a constructive manner is the study by Bello et al. [19]. They used the pointer network (PtrNet) architecture [20] to encode input sequences and construct the node

sequence in the decoder. Their model was tested on the TSP and the 0–1 knapsack problem (KP) and yielded close-to-optimal results. The PtrNet model is further improved by Khalil et al. [21] and Nazari et al. [22]. Deudon et al. [23] used the pointer network with an attention encoder. Inspired by the Transformer model for machine translation [5], Kool et al. [3] proposed the attention model (AM) based on the transformer architecture to solve various CO problems such as the TSP, VRP, and Orienteering Problem (OP). Cappart et al. [24] combined the RL and constraint programming (CP) to solve the TSP with Time Windows (TSPTW) by learning branching strategies. Additionally, Bono et al. [15] proposed a modified Transformer model to handle the dynamic and stochastic VRPs (DS-VRPs) by using online measurements of the environment to online select the next vehicle via a vehicle-customer intersection module. More recently, Li et al. [25] improved the AM to solve the Heterogeneous Capacitated VRP (HCVRP). Li et al. [26] proposed the attention-dynamic model to solve the covering salesman problem (CSP). Xu et al. [27] designed an attention model with multiple relational attention mechanism that better captures the transition dynamics. Pan and Liu [28] designed a graph-based partially observable MDP (POMDP) that captures the changes in the customer demands to solve a dynamic and uncertain VRP using a deep neural network model with dynamic attention mechanism. Besides attention model, Wang [29] proposed a variational autoencoder-based reinforcement learning methodology using a graph reasoning network for classic vehicle routing problems. In terms of performance, Kwon et al. [4] introduced the POMO method which has demonstrated state-of-the-art results on TSP, CVRP, and KP. During training, the POMO decoder generates multiple heterogeneous trajectories that start at every node to maximize entropy on the first action.

The majority of past studies used policy gradient approaches, which have advantages over supervised learning (SL) [30]. Bello et al. [19] used an actor-critic algorithm to train their model. However, Kool et al. [3] showed that a greedy rollout baseline yields better results than a (learned) critic baseline. Many subsequent works, including [6, 25–27], and [7], used the greedy rollout baseline.

Although the greedy rollout baseline is effective, it requires an additional forward-pass of the model, increasing the computational load on the device. To leverage more data parallelism for efficient learning of training instances, Kool et al. [31, 32] proposed to use a local baseline equal to the average return over  $k$  trajectories sampled without replacement from a single instance using *Stochastic Beam Search*. They reported that this baseline performed on par or slightly better than the computationally expensive greedy rollout and significantly better than the batch baseline. The benefit of sampling without replacement is that the gradient estimators do not lose much final performance while learning from substantially fewer instances (number of training instances is reduced by factor of  $k$ ).

In addition, Kwon et al. [4] used a shared baseline based on all POMO samples, taking the average tour length over  $n$  sample trajectories from a single instance, where  $n$  is the number of nodes. Like multiple-sample baselines in [31], the POMO-shared baseline is local, concentrating on a single instance. As reported in [4], their baseline is very effective since it generates  $n$ , typically larger than  $k$  in [31], non-duplicative sample trajectories for a single instance. However, the POMO requires an additional tensor dimension, and as the graph size  $n$  increases, the tensor size increases by  $n$ -fold. Consequently, while the training time of POMO is comparable to that of REINFORCE with greedy rollout (owing to the parallel generation of trajectories), it requires more GPU memory. Moreover, the POMO training may not be readily applicable on problems such as MSTOP, where we cannot simply use all the nodes as starting points for exploration.

Many strategies for efficient inference were also proposed in prior studies. Bello et al. [19] proposed the “one-shot” greedy inference and sampling strategies. Deudon et al. [23] improved their solution quality by refining it with the 2-Opt heuristic [33]. Kwon et al. [4] suggested  $\times 8$

instance-augmentation to generate multiple trajectories and select the best solution to obtain better results.

### 3 Problem definition

#### 3.1 Mathematical formulation of MSTOP

This section presents the MILP formulation of MSTOP. In particular, this formulation is defined on a graph following [10]. A complete graph  $G$  consists of the set of all nodes ( $N$ ) and the set of arcs ( $A$ ). We summarize key notations in the mathematical formulation of MSTOP in Table 1. Since each vehicle is associated with a unique starting location, we drop the subscript  $k$  in the notation  $v_k$  for simplicity whenever its inclusion is implied.

In the MSTOP, multiple vehicles begin at locations different from the depot. Each vehicle has an available amount of fuel at the start. Given the vehicle set, the MSTOP determines  $K$  routes that maximize the total profits collected over the partial routes while satisfying a maximum duration constraint on each route.

In the MILP formulation below,  $x_{ijk}$  denotes a binary variable, which equals one if arc  $(i, j)$  in  $A$  is traversed by vehicle  $k$  (in  $K$ ), and zero otherwise. Also, binary variable  $y_{ik}$  equals one if node  $i$  (in  $X$ ) is visited by vehicle  $k$  (in  $K$ ) and otherwise zero.  $t_{ij}$  is measured as the Euclidean distance between the two nodes, and the subscript  $v$  denotes a vehicle’s launching node. The MILP formulation for the MSTOP is as follows:

(MILP Formulation for MSTOP)

$$\max \sum_{i \in X \setminus \{0\}} p_i \sum_{k=1}^K y_{ik}, \quad (1)$$

subject to

**Table 1** Notation table for MSTOP

Notation	Description
$G = (N, A)$	Complete graph
$N = X \cup V$	Set of all nodes
$X = \{0, 1, \dots, n\}$	Customer (1 ~ $n$ ) and depot (0) nodes
$V = \{v_k\}_{k=1}^K$	Initial locations of $K$ vehicles
$A = A_n \cup A_k$	Set of arcs or edges
$A_n = \{(i, j)   i, j \in X, i \neq j\}$	Set of arcs among customer/depot nodes
$A_k = \{(k, j)   k \in V, j \in X\}$	Set of arcs among vehicle locations and remaining nodes
$T_{max}$	Maximum route duration for all vehicles
$t_{ij}$	Travelling length associated with arc $(i, j)$
$f_k$	Available fuel amount at the start for vehicle $k$ ( $\in K$ )
$p_i$	Scalar prize associated with visiting node $i$ ( $\in X$ )

$$\sum_{i \in X \setminus \{0\}} x_{i0k} + x_{v0k} = 1 \quad k = 1, \dots, K, \quad (2)$$

$$\sum_{j \in X, j < i} x_{ijk} + \sum_{j \in X, i < j} x_{jik} + x_{vik} = 2y_{ik} \quad \forall i \in X \setminus \{0\}, k = 1, \dots, K, \quad (3)$$

$$\sum_{j \in X} x_{vjk} = y_{vk} \quad k = 1, \dots, K, \quad (4)$$

$$\sum_{k=1}^K y_{vk} = K, \quad (5)$$

$$\sum_{k=1}^K y_{0k} = K, \quad (6)$$

$$\sum_{k=1}^K y_{ik} \leq 1 \quad \forall i \in X \setminus \{0\}, \quad (7)$$

$$\sum_{(i,j) \in A, j < i} t_{ij} x_{ijk} + f_k \leq T_{\max} \quad k = 1, \dots, K, \quad (8)$$

$$y_{ik} \in \{0, 1\} \quad \forall i \in X \cup \{v\}, k = 1, \dots, K, \quad (9)$$

$$x_{ijk} \in \{0, 1\} \quad \forall (i,j) \in A, j < i, i \in X \setminus \{0\} \cup \{v\}, k = 1, \dots, K. \quad (10)$$

Equation (1) expresses the objective of the problem, which is maximizing the collected profit from routes. Equations (2)-(10) present the constraints of the problem. Equation (2) ensures that all routes end at the depot. Equation (3) guarantees that an arc enters a node and leaves from that node. Equations (4)-(5) ensure that a route begins at the initial vehicle location. Equation (6) constrains the number of total routes ( $K$ ). Equation (7) imposes a constraint that each node is visited at most once. Equation (8) limits the maximum duration or length for each route. Lastly, Eqs. (9)-(10) define the decision variables.

Note that the local constraints of the formulation do not guarantee that all nodes in a route are properly connected without subtours. To generate a feasible set of routes, we add the subtour elimination constraints. However, given the nature of routing problems, adding such constraints before the optimization can significantly increase the model size for large-scale problems. As a result, we add the subtour elimination constraints in a *lazy fashion* [34]. This way, we can remove solutions with subtours during the optimization.

### 3.2 MDP formulation of MSTOP

This section introduces the MDP formulation of the MSTOP. To apply reinforcement learning to MSTOP, we model the

problem as a sequential decision-making process, where an agent performs a sequence of actions (i.e., decides which node to visit) through interactions with the surrounding environment (i.e., observing changes in the state) to maximize the cumulative reward.

In our MDP setting, a vehicle is first assigned at random. The agent selects nodes to visit starting from the initial position of the assigned vehicle. Once a partial route is constructed, the agent chooses the next vehicle starting at a different location. The complete solution is constructed by concatenating the individual partial routes. We model the MSTOP as an MDP defined by a 4-dimensional tuple  $\langle S, A, P, R \rangle$ , where  $S$  denotes the state space,  $A$  the action space,  $P$  the state transition model, and  $R$  the reward model.

**State space (S)** Each state at time step  $t$  is defined as a tuple  $s_t (= \langle X_t, V_t \rangle)$ . The first component of the tuple,  $X_t$ , denotes the set of all nodes ( $= \{x_i^t\}$ ), and the second component,  $V_t$ , expresses the states of all vehicles ( $= \{v_k^t\}$ ). Here,  $x_i^t (= (r_i, p_i^t))$  contains the information of a node where  $r_i (= (x_i, y_i))$  is the location and  $p_i^t$  is the prize assigned to the node. Also,  $v_k^t = (\rho_k^t, f_k^t, O_k^t)$  denotes the vehicle information where  $\rho_k^t = (x_k, y_k)$  represents the vehicle location,  $f_k^t$  is the vehicle's available/remaining fuel amount, and  $O_k^t$  is the total prizes collected until step  $t$ . We denote the terminal time as  $T$  at which all vehicles arrive at the depot.

**Action space (A)** The permissible set of actions in our MDP is the choice of the next node to visit by considering the vehicle's current partial route and the amount of fuel. We denote each action at time step  $t$  ( $a_t \in A$ ) as  $x_j^t$  and view the action as an addition of a node to the partial route. The construction of partial route satisfies the maximum travel duration constraint for each vehicle by action masking policy, i.e. masking the nodes that cannot be visited.

**State transition model ( $P: S \times A \rightarrow S$ )** The state transition model describes how the current state ( $s_t$ ) transitions to the next state ( $s_{t+1}$ ) when an action ( $a_t$ ) is taken. We adopt deterministic transition dynamics, i.e., a vehicle moves to the chosen node with the probability of 1. Given the current vehicle  $k$  and chosen action  $a_t = (x_j^t)$  (i.e., the vehicle visits node  $j$ ), we update the elements of  $\{x_i^t\}$  and  $\{v_k^t\}$  at step  $t$  as follows.

$$p_i^{t+1} = \begin{cases} 0 & i = j \\ p_i^t & i \neq j \end{cases}, \quad (11)$$

$$\rho_k^{t+1} = r_j, \quad (12)$$

$$f_k^{t+1} = f_k^t - t_{ij}, \tag{13}$$

$$O_k^{t+1} = O_k^t + p_j^t. \tag{14}$$

Equation (11) sets the prize associated with node  $j$  as 0 when visited, and Eq. (12) updates the current location of vehicle  $k$ . Equation (13) updates the available amount of fuel by subtracting  $t_{ij}$  (distance between nodes  $i$  and  $j$ ) from it. Equation (14) updates the total prize by adding the prize value obtained at node  $j$  ( $p_j$ ).

**Reward model ( $R: S \times A \rightarrow \mathbb{R}$ )** We model the cumulative reward as the sum of total prizes collected from all partial routes. To be specific, the reward is defined as  $\mathcal{R} = \sum_{k=1}^K O_k^T$ . Termination time  $T$ , determined by the number of actions executed until the completion of all partial routes, defines the trajectory length.

## 4 Proposed model and solution procedure

### 4.1 Proposed framework

Figure 2 explains a framework proposed to solve the MSTOP, which contains inner and outer loops. The inner loop begins at the vehicle’s initial location and generates a partial route that terminates at the depot. Each partial route is a permutation of numbers ending with 0, as shown in Fig. 3. When the inner loop is finished, the outer loop updates the graph instance.

This procedure contrasts the models in [3], where the encoder is executed only once initially ( $t=0$ ). In classical CO problems, when a vehicle returns to the depot, the graph instance changes only slightly because the next vehicle starts at the same depot. However, constructing a partial route in an MSTOP modifies the graph instance. Not only does the next vehicle face a different set of nodes

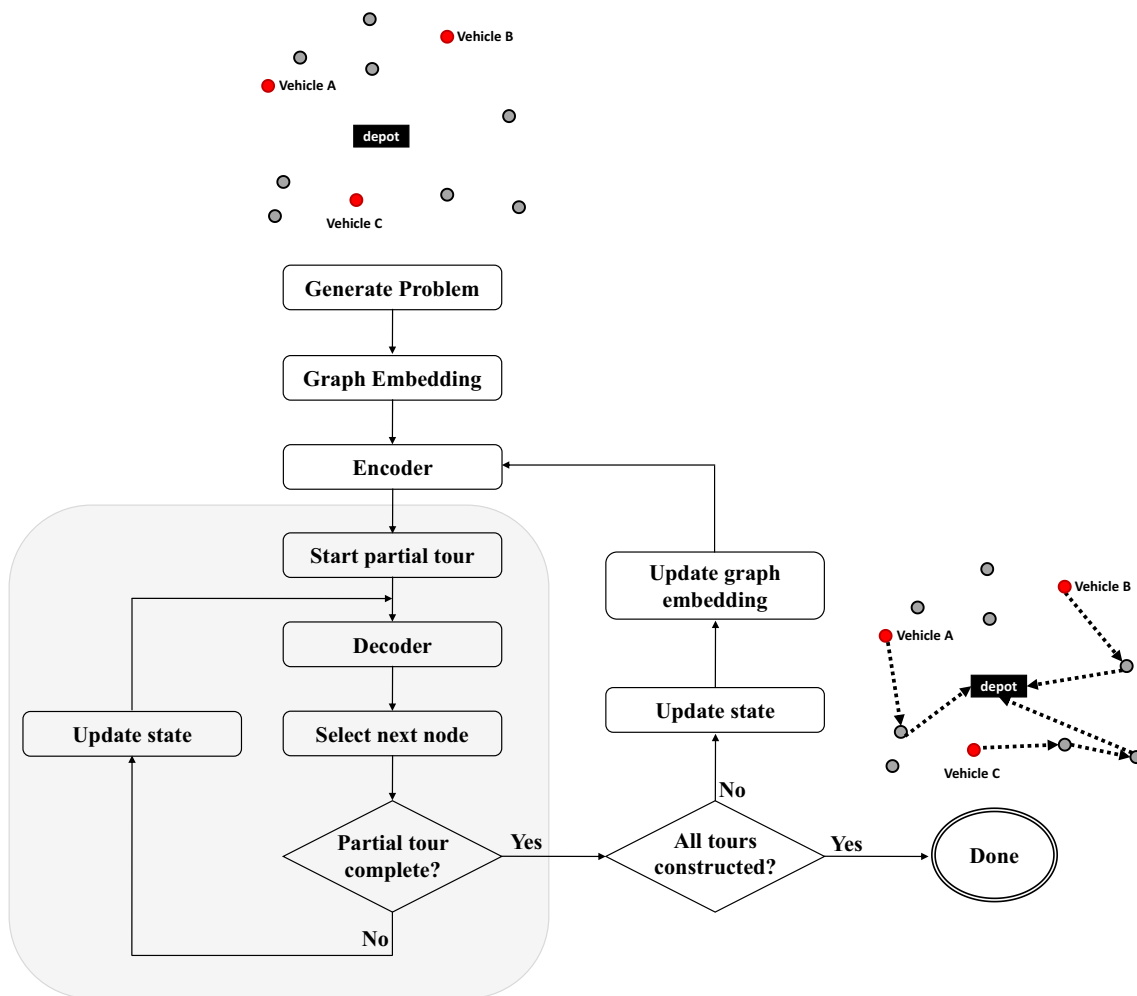
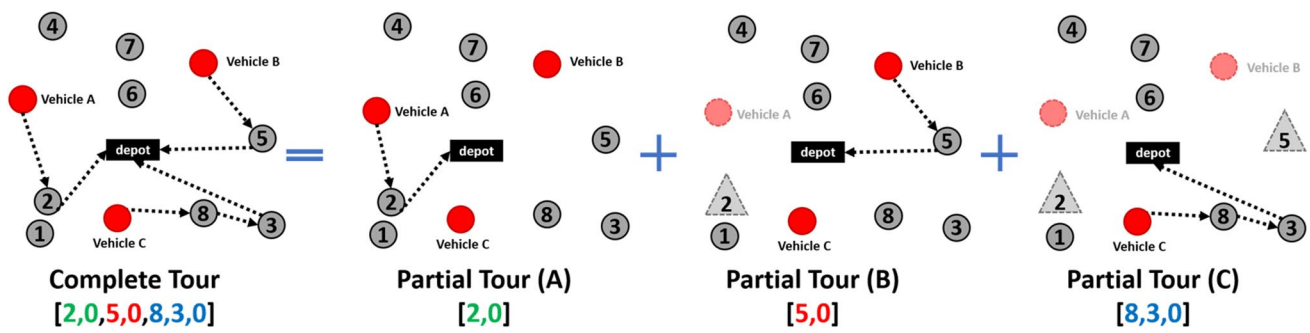


Fig. 2 Diagram explaining the proposed framework



**Fig. 3** Complete MSTOP solution obtained by combining individual routes – each route is constructed by a single vehicle. Opaque nodes indicate either (i) visited nodes (triangular) or (ii) vehicles that have arrived at the depot

**Table 2** Summary of key terminologies for DDTM

Terminology	Description
$(x_0, y_0)$	Depot location
$(x_i, y_i, p_i)$	Node location and prize for mission $i$
$(\hat{x}_k, \hat{y}_k, f_k)$	Location and fuel for vehicle $k$
$h^{(l)}$	Embedded feature data at $l^{\text{th}}$ encoder layer
$n_{enc}, n_{dec}$	Number of encoder/decoder layers
$h^{(n_{enc})}$	Node-vehicle embedding
$\bar{h}^{(n_{enc})}$	Graph embedding
$t_{dec}$	Current decoding step
$PE_{t_{dec}}$	Positioning encoding at $t_{dec}$
$p_t^{dec}$	Probability distribution over all candidate nodes

(i.e., without visited nodes), but it also starts at a different location.

The rationale behind this sequential construction framework, which addresses one vehicle at a time, is grounded in empirical observations that simultaneous consideration of the next node for each vehicle can impede training convergence due to additional freedom in decision-making. During early training epochs, this additional complexity presents challenges for the model to “learn” to generate routes.

In the solution procedure, the encoder plays a pivotal role in transforming the raw features of the graph instance, encompassing mission node and vehicle data, into a hidden representation known as node-vehicle and graph embeddings. These embeddings, as computed in Eqs. (23) and (24), capture essential information about the spatial relationships and characteristics of the nodes and vehicles within the graph. The major interplay between the encoder and decoder occurs when the output of the encoder, comprising the node-vehicle embeddings and graph embedding, is sent as input to the decoder. Subsequently, the decoder leverages this information to

extract relevant features, generating a probability distribution over non-visited (candidate) nodes that guides the selection of the next node in the route. This iterative process continues until the depot is chosen (i.e., completing an individual vehicle route). Following each partial route, the graph is updated before advancing to the next vehicle. Table 2 outlines key terminologies used in this section that describe the structure of DDTM.

### 4.2 Encoder-decoder architecture of DDTM

Figure 4 presents the encoder-decoder architecture of DDTM used for MSTOP. Figure 5 illustrates the encoder structure (for a single encoding layer). The encoder embeds the MSTOP features using separate parameters for the additional vehicle features – vehicle location and available fuel. We denote the embedded feature data as  $h^{(l)}$ , where  $l$  is the encoder layer. The embedded data as a whole represents the graph instance, and each element in  $h^{(l)}$  is a mapping corresponding to each feature.

A good feature mapping needs to consider the feature’s context within the graph. For example, the node representation should contain sufficient information to be selected among its neighbors and to determine its position in the output sequence. To understand how one feature is related to another from a broader perspective, we apply multi-head self-attention, which generates enhanced feature embeddings. The self-attention mechanism enables the encoder to effectively weigh and consider the significance of different features of the input graph. The encoding steps are formally expressed as follows.

$$h_0^{(l)} = [x_0, y_0] W_0^{init}, \tag{15}$$

$$h_i^{(l)} = [x_i, y_i, p_i] W_{node,i}^{init} \text{ for } i \in \{1, \dots, N\}, \tag{16}$$

$$\hat{h}_k^{(l)} = [\hat{x}_k, \hat{y}_k, f_k] W_{veh,k}^{init} \text{ for } k \in \{1, \dots, K\}, \tag{17}$$

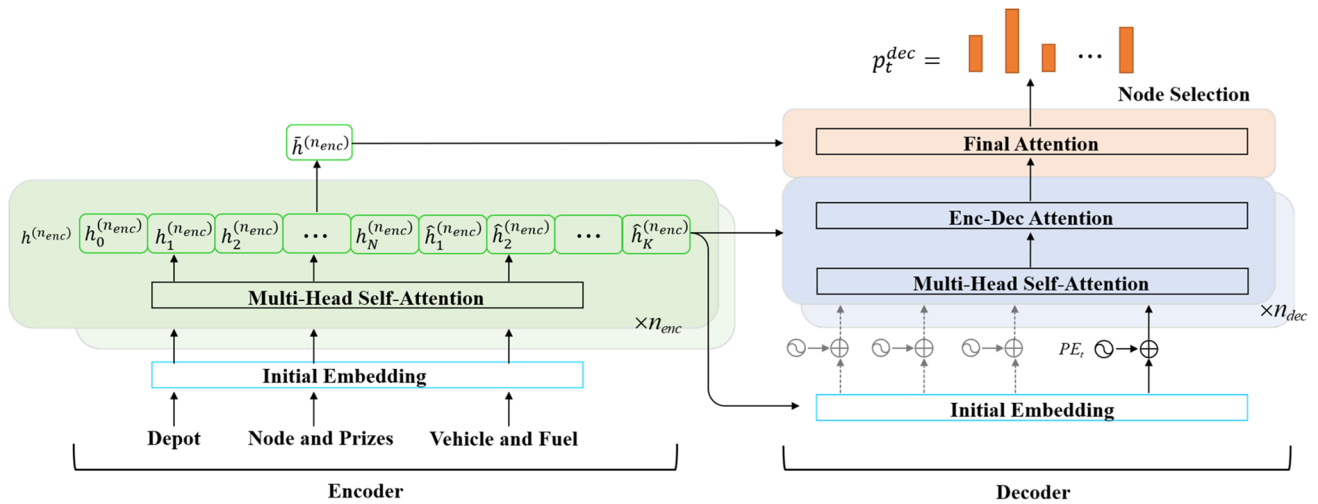


Fig. 4 Encoder-Decoder architecture of DDTM

$$h^{(l)} = [h_0^{(l)}, h_1^{(l)}, \dots, h_N^{(l)}, \hat{h}_0^{(l)}, \dots, \hat{h}_K^{(l)}], \tag{18}$$

$$Q_l = h^{(l)} W_l^Q, K_l = h^{(l)} W_l^K, V_l = h^{(l)} W_l^V, \tag{19}$$

$$Z_l^h = attention(Q_l, K_l, V_l) = \text{Softmax}\left(\frac{Q_l K_l^T}{\sqrt{d_k}}\right) V_l, \tag{20}$$

where  $d_k = d/H$  with  $d (= 128)$  is a hyperparameter and  $H (= 8)$  is the number of heads. To compute multi-head attention, we concatenate the attention outputs of each head ( $Z_l^h$ ) as

$$MHA(h^{(l)}) = [Z_l^1, Z_l^2, \dots, Z_l^H] W_l^{out}. \tag{21}$$

The next embedded feature,  $h^{(l+1)}$ , is obtained by passing  $h^{(l)}$  through a feed-forward layer with batch normalization, residual connection, and ReLU activation as follows,

$$\tilde{h}^{(l)} = BN(h^{(l)} + MHA(h^{(l)})), \tag{22}$$

$$h^{(l+1)} = FF(\tilde{h}^{(l)}) = BN(W_1^{ff} \text{ReLU}(W_0^{ff} \tilde{h}^{(l)}) + \tilde{h}^{(l)}), \tag{23}$$

where  $W_0^{ff} \in \mathbb{R}^{d \times d_h}$  and  $W_1^{ff} \in \mathbb{R}^{d_h \times d}$  are trainable parameters with  $d_h (= 512)$ . After  $n_{enc}$  encoding layers, the final output of the encoder is the node-vehicle embedding ( $h^{(n_{enc})}$ ) and the graph embedding ( $\hat{h}^{(n_{enc})}$ ) defined as

$$\hat{h}^{(n_{enc})} = \begin{cases} \frac{1}{N+K+1} \left( \sum_{i=0}^{N+1} h_i^{(n_{enc})} + \sum_{k=1}^K \hat{h}_k^{(n_{enc})} \right) & \text{if } t = 0 \\ \frac{1}{N'+K'+1} \left( \sum_{i=0}^{N+1} h_i^{(n_{enc})} + \sum_{k=1}^K \hat{h}_k^{(n_{enc})} \right) & \text{if } t > 0 \end{cases} \tag{24}$$

where  $N' (= N - N_{visited})$  is the remaining number of nodes and  $K'$  is the remaining number of vehicles. After a partial route is constructed ( $t > 0$ ), the graph instance seen by the next vehicle differs from that seen by the previous ones. We update the graph instance by computing  $h^{(n_{enc})}$  and  $\hat{h}^{(n_{enc})}$  using Eqs. (15)–(24), and mask the visited nodes using the outer product as,

$$\mathcal{M}_{att} = \mathcal{M} \otimes 1^T + 1 \otimes \mathcal{M}^T - \mathcal{M} \otimes \mathcal{M}^T \in \mathbb{R}^{(N+K) \times (N+K)}, \tag{25}$$

$$Z_l = attention(Q_l, K_l, V_l) = \text{Softmax}\left(\frac{Q_l K_l^T}{\sqrt{d_k}} \odot \mathcal{M}_{att}\right) V_l, \tag{26}$$

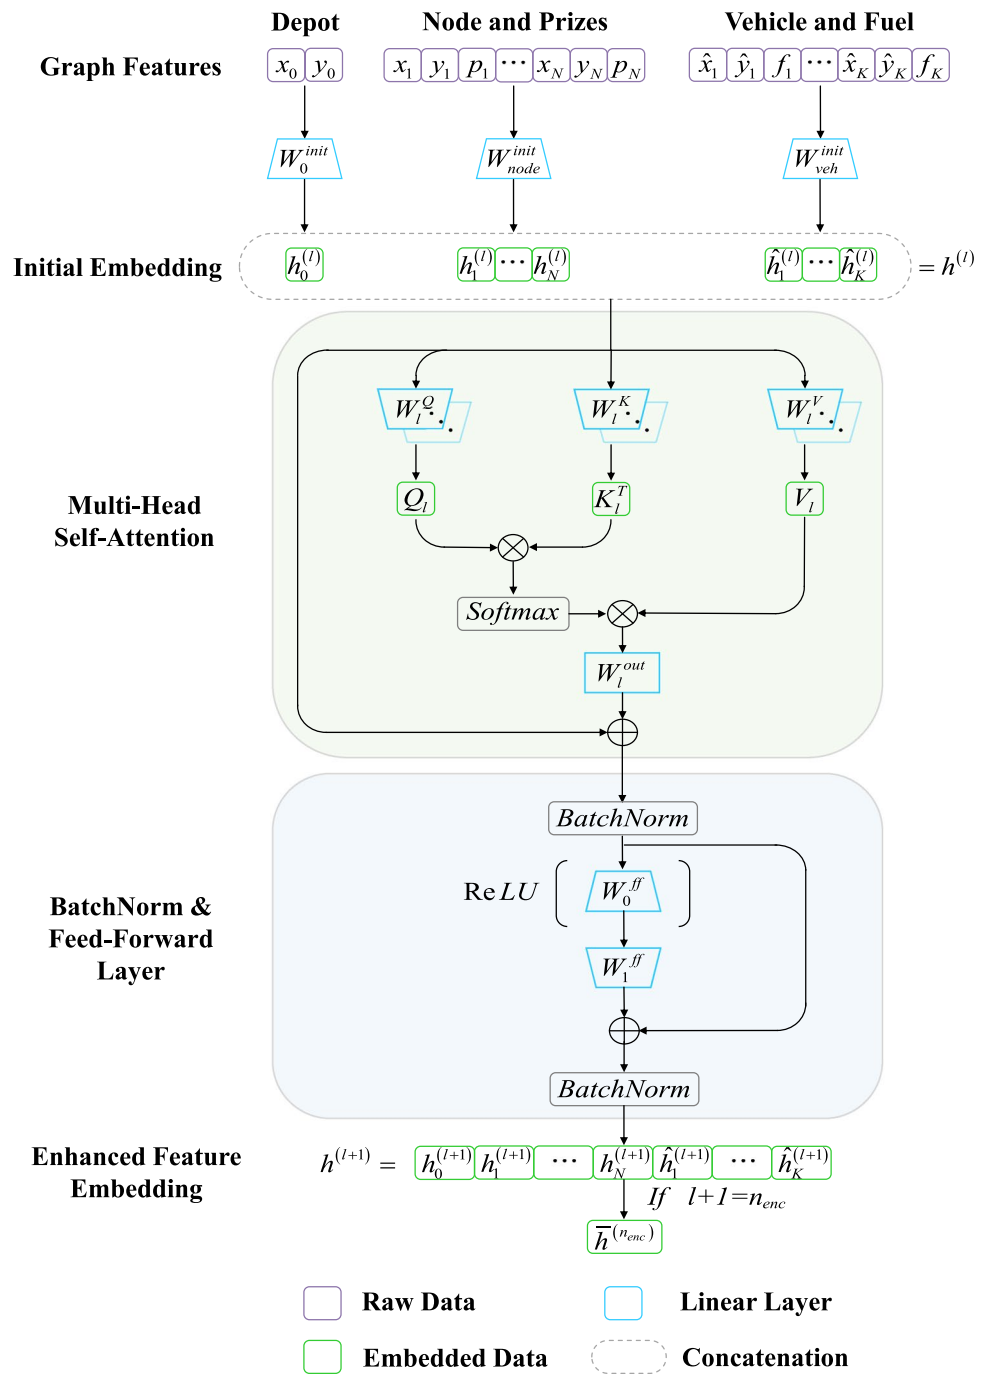
where  $\mathcal{M} \in \mathbb{R}^{(N+K) \times 1}$  is a column mask vector that masks visited nodes and vehicles at the depot,  $1 \in \mathbb{R}^{(N+K) \times 1}$  is a column vector of ones, and  $\odot$  is the Hadamard product for matrices.

Given the node-vehicle and graph embeddings by the encoder, the decoder produces probability distributions ( $p_t^{dec}$ ) for all candidate nodes and selects the next node. Candidate nodes are those not visited by any vehicle at the start of decoding. Our decoding strategy consists of three steps based on [6] as follows:

**Step 1:** We begin by computing the multi-head self-attention between the current node and the nodes in the current partial route. By examining the history of visited nodes for the current node, we obtain the contextual information up to the current decoding time,  $t_{dec}$ . We first extract the current node embedding ( $\tilde{h}_{t_{dec}}$ ) from the node-vehicle embeddings ( $h^{(n_{enc})}$ ), then concatenate it with the current amount of fuel ( $f_k^t$ ). We set  $t_{dec}$  as zero at the start of the decoding for each partial route and increment it by one



Fig. 5 Encoder structure



per each node selection within the inner loop. Since the decoding starts at the vehicle’s initial location, we select the current node embedding as  $\tilde{h}_0 = \hat{h}_k^{(n_{enc})}$  and update it as  $\tilde{h}_{t_{dec}} = h_a^{(n_{enc})}$ , where  $a$  ( $:= a_{t_{dec}-1} \in \{1, \dots, N\}$ ) is the node selected in the previous step. Since the partial route begins at the vehicle’s location and ends at the depot, the order of nodes in the partial route matters. This characteristic requires the addition of positional encoding [5] (which describes the position of a node within the graph instance so that each node can have a unique representa-

tion) to the linearly projected pair to generate  $h_{t_{dec}}^{(l)} \in \mathbb{R}^{1 \times d}$  as follows,

$$h_{t_{dec}}^{(l)} = [\tilde{h}_{t_{dec}}, f_k^{t_{dec}}] W_o^{proj} + PE_{t_{dec}}, \quad (27)$$

where  $PE_{t_{dec}}$  is a  $d$ -dimensional row vector. Each element of the vector is defined as

$$PE_{t_{dec},i} = \begin{cases} \sin(t_{dec}/10000^{2i/d}) & \text{if } i \text{ is even} \\ \cos(t_{dec}/10000^{2i/d}) & \text{if } i \text{ is odd} \end{cases}, \quad (28)$$

where  $i \in \{0, 1, \dots, d-1\}$  is the position along the  $d$  dimension.

Figure 6 illustrates the decoding **Step 1**. There are  $t_{dec}$  visited nodes in the current partial route. We first compute the self-attention between  $h_{t_{dec}}^{(l)}$  and  $[h_0^{(l)}, h_1^{(l)}, \dots, h_{t_{dec}-1}^{(l)}] \in \mathbb{R}^{t_{dec} \times d}$ . **Step 1** is mathematically described as follows (where  $d_k = d/H$ ).

$$\overset{\circ}{Q}_l = h_{t_{dec}}^{(l)} W_{l,sa}^Q \in \mathbb{R}^{1 \times d_k}, W_{l,sa}^Q \in \mathbb{R}^{d \times d_k} \quad (29)$$


$$K_l = [h_0^{(l)}, h_1^{(l)}, \dots, h_{t_{dec}-1}^{(l)}] W_{l,sa}^K \in \mathbb{R}^{t_{dec} \times d_k}, W_{l,sa}^K \in \mathbb{R}^{d \times d_k}, \quad (30)$$

$$V_l = [h_0^{(l)}, h_1^{(l)}, \dots, h_{t_{dec}-1}^{(l)}] W_{l,sa}^V \in \mathbb{R}^{t_{dec} \times d_k}, W_{l,sa}^V \in \mathbb{R}^{d \times d_k}, \quad (31)$$

$$\overset{\circ}{Z}_l^h = \text{attention}(Q_l, K_l, V_l) = \text{Softmax}\left(\frac{Q_l K_l^T}{\sqrt{d_k}}\right) V_l \in \mathbb{R}^{1 \times d_k}, \quad (32)$$

$$\overset{\circ}{Z}_{l,att}^h = \text{attention}(Q_{l,att}, K_{l,att}, V_{l,att}) = \text{Softmax}\left(\frac{Q_{l,att} K_{l,att}^T}{\sqrt{d_k}} \odot \mathcal{M}^T\right) V_{l,att} \in \mathbb{R}^{1 \times d_k}, \quad (38)$$

$$h_{t_{dec}}^{(l)} \leftarrow \text{MHA}(\cdot) \Big|_{att} = [Z_{l,att}^{\circ 1}, Z_{l,att}^{\circ 2}, \dots, Z_{l,att}^{\circ H}] W_{l,att}^{out} \in \mathbb{R}^{1 \times d}, W_{l,att}^{out} \in \mathbb{R}^{d \times d}. \quad (39)$$



**Fig. 6** Step 1 of the decoding procedure. The orange contour indicates the partial route at time step  $t_{dec}$

$$h_{t_{dec}}^{(l)} \leftarrow \text{MHA}(\cdot) \Big|_{sa} = [Z_{l,sa}^{\circ 1}, Z_{l,sa}^{\circ 2}, \dots, Z_{l,sa}^{\circ H}] W_{l,sa}^{out} \in \mathbb{R}^{1 \times d}, W_{l,sa}^{out} \in \mathbb{R}^{d \times d}. \quad (33)$$

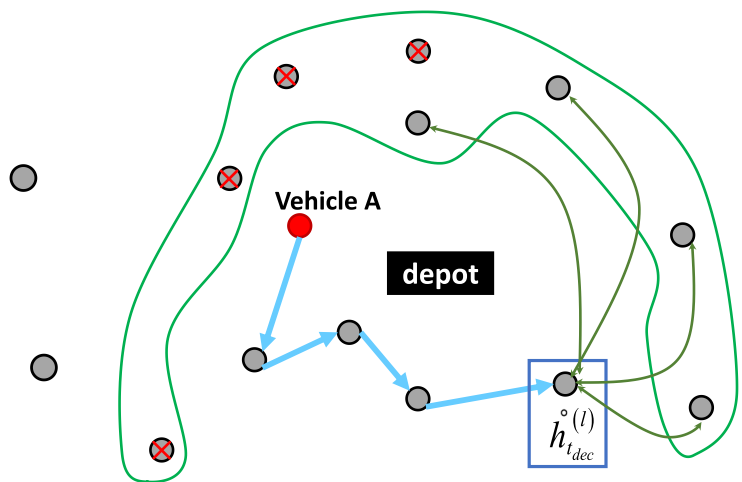
**Step 2:** This step queries the next node to visit among all candidate nodes. The step uses the encoder-decoder attention between the self-attention of a partial route (output of **Step 1**; denoted as  $h_{t_{dec}}^{(l)}$  for coherence) and context node embeddings ( $H_{node} \in \mathbb{R}^{(N+2) \times d}$ ; node-vehicle embeddings with current vehicle embedding only (Eq. (34)). We mask the nodes that cannot be visited from the current location. Figure 7 illustrates the encoder-decoder attention in **Step 2** of the decoding procedure. The following equations express **Step 2**.

$$H_{node} = [h_0^{(n_{enc})}, h_1^{(n_{enc})}, \dots, h_N^{(n_{enc})}, \hat{h}_k^{(n_{enc})}] \in \mathbb{R}^{(N+2) \times d}, \quad (34)$$

$$Q_{l,att} = h_{t_{dec}}^{(l)} W_{l,att}^Q \in \mathbb{R}^{1 \times d_k}, W_{l,att}^Q \in \mathbb{R}^{d \times d_k}, \quad (35)$$

$$K_{l,att} = H_{node} W_{l,att}^K \in \mathbb{R}^{(N+2) \times d_k}, W_{l,att}^K \in \mathbb{R}^{d \times d_k}, \quad (36)$$

$$V_{l,att} = H_{node} W_{l,att}^V \in \mathbb{R}^{(N+2) \times d_k}, W_{l,att}^V \in \mathbb{R}^{d \times d_k}, \quad (37)$$



**Fig. 7** Step 2 of the decoding procedure. The blue box denotes the current node, the green contour represents the set of candidate nodes, and the red cross indicates masked nodes

**Step 3:** Step 1 and Step 2 form a single decoding layer. After  $n_{dec}$  decoding layers, the resultant output  $h_{t_{dec}}^{o(l)}$  is sent to the final attention layer, where we compute a single-head attention to get probability distribution across all candidate nodes. The decoder receives a graph embedding ( $\tilde{h}^{(n_{enc})}$ ) from the encoder, and its linear projection is added to  $h_{t_{dec}}^{o(l)}$ . The query is constructed from the sum. The key is obtained by a linear projection of  $\tilde{H}_{node} \in \mathbb{R}^{(N+1) \times d}$ , which is the context node embedding in Eq. (34) without current vehicle embedding ( $\hat{h}_k^{(n_{enc})}$ ). The decoding step 3 is described as the following equations and illustrated in Fig. 8.


$$\tilde{H}_{node} = [h_0^{(n_{enc})}, h_1^{(n_{enc})}, \dots, h_N^{(n_{enc})}] \in \mathbb{R}^{(N+1) \times d}, \quad (40)$$

$$Q_{f.att} = h_{t_{dec}}^{o(l)} W_{f.att}^Q \in \mathbb{R}^{1 \times d}, \quad W_{f.att}^Q \in \mathbb{R}^{d \times d}, \quad (41)$$

$$K_{f.att} = \tilde{H}_{node} W_{f.att}^K \in \mathbb{R}^{(N+1) \times d}, \quad W_{f.att}^K \in \mathbb{R}^{d \times d}, \quad (42)$$

$$p_t^{dec} = \text{Softmax} \left( C \cdot \text{Tanh} \left( \frac{Q_{f.att} K_{f.att}^T}{\sqrt{d}} \odot \mathcal{M}^T \right) \right) \in \mathbb{R}^{1 \times (N+1)}. \quad (43)$$

The value of  $C$  in Eq. (43) is selected as 10. Consequently, the next node  $a \in \{0, 1, \dots, N\}$  is sampled from the output probability distribution  $p_t^{dec}$  (following a categorical distribution or greedy fashion), and  $t$  and  $t_{dec}$  are incremented by one.



**Fig. 8** Step 3 of the decoding procedure. The purple boxes above candidate nodes and depot indicate the selection probability

## 5 Data-efficient training with proposed REINFORCE baseline

This section presents our proposed training methodology that improves learning efficiency. In terms of data efficiency, our methodology requires fewer (raw) training instances at every epoch compared to the conventional method. Since the training instances are generated on the fly, an epoch in our proposed methodology takes shorter time to generate the training data and transfer them over to the GPU. Moreover, in terms of sample efficiency, our method reaches an equivalent performance (validation score) within fewer training epochs or with fewer training instances in comparison with other methods.

### 5.1 Preliminary

Policy-gradient methods learn the policy directly and explicitly through gradient-based optimization. We define the model’s policy as a parametrized function  $\pi_\theta(a|s)$ , where  $\theta$  denotes the trainable parameters of the model. The function is stochastic in that it defines a probability distribution of actions ( $a$ ) at a given state ( $s$ ). The goal of policy optimization is to maximize the expected cumulative return (sum of rewards,  $R(\tau)$ ) of the trajectory ( $\tau = (s_0, a_0, s_1, a_1, \dots, s_T)$ ) whose actions are chosen by the policy defined as

$$J(\theta) = \mathbb{E}_{\tau \sim \pi_\theta} [R(\tau)] = \mathbb{E}_{\tau \sim \pi_\theta} \left[ \sum_{t=0}^T r(s_t, a_t) \right]. \quad (44)$$

The objective of the policy optimization problem expressed in Eq. (44) uses the expectation over all possible trajectories. For a given stochastic policy ( $\pi_\theta$ ), the trajectory probability ( $P(\tau; \pi_\theta) := P(\tau; \theta)$ ) represents the probability of generating a trajectory following the policy. The trajectory probability is factorized as

$$P(\tau; \theta) = \prod_{t=0}^T \pi_\theta(a_t | s_t) p(s_{t+1} | s_t, a_t), \quad (45)$$

where  $p(s_{t+1} | s_t, a_t)$  is the state-transition probability of the MDP defined in Section III. Williams [35] proposed a viable estimator of the policy gradient using Monte-Carlo sampling by assuming that  $R(\tau)$  is independent of  $\theta$ :

$$\nabla_\theta J(\theta) = \mathbb{E}_{\tau \sim \pi_\theta} [R(\tau) \nabla_\theta \log P(\tau; \theta)]. \quad (46)$$

In practice, the unbiased REINFORCE gradient estimator presented in Eq. (46) suffers from a high variance of the returns  $R(\tau_i)$  and is sample inefficient since it requires many sample episodes to converge. We can overcome these issues by including a baseline ( $b(s)$ ), an action-independent function, in the policy gradient estimation. Consequently, an unbiased estimate of the gradient with reduced variance is expressed as

$$\nabla_{\theta} J(\theta) = \mathbb{E}_{\tau \sim \pi_{\theta}} [(R(\tau) - b) \nabla_{\theta} \log P(\tau; \theta)]. \tag{47}$$

### 5.2 Choice of REINFORCE baseline $b(s)$

An example of the baseline is the average return over sample trajectories ( $b = \mathbb{E}_{\tau \sim \pi_{\theta}} [R(\tau)] \approx \frac{1}{N} \sum_{i=1}^N R(\tau_i)$ ), where  $N$  is the number of samples in a mini-batch. Although the mini-batch baseline can effectively reduce variance in Gradient-Bandit algorithms [36], Kool et al. [31] showed that it performs significantly worse than other state-of-the-art baselines.

Prior studies suggest that designing an effective yet computationally tractable REINFORCE baseline is crucial in training the policy network. In this work, we propose to use the average return of sample trajectories generated by instance augmentation from a single instance as the baseline, referred to as the *instance-augmentation baseline*. Our baseline is a potential alternative to the existing baselines with improved training speed and reduced variance. The proposed baseline is motivated by observations of other baselines in prior works. In general, a local baseline performs significantly better than a batch baseline. In particular, a local baseline based on multiple samples without replacement is expected to perform better because non-duplicate samples are guaranteed [31, 32]. This observation can be extended to POMO [4], whose local batch mean is based on  $N$  non-duplicate sample trajectories from a single instance, despite an increased tensor size. Since each POMO trajectory begins at a unique node, these samples are also guaranteed to be non-identical. These REINFORCE baselines are more data-efficient than the greedy rollout because they require fewer training instances (reduced by some factor).

It would be effective if a baseline as equally *data-efficient* as the multiple-sample baselines and even *computationally lighter* than the POMO shared baseline is used. The proposed baseline meets these requirements by utilizing the instance augmentation, which was first suggested in [4] for effective inference.

Table 3 lists the coordinate transformations applied to all features (nodes, depots, and vehicle locations) to generate additional instances for a given training instance (a total of 8 instances). While each of these instances is distinct, the optimal tour would be identical since these transformations preserve the lengths between nodes. We then rollout sample trajectories of each of these “counterfactuals.” The policy model would perceive these as distinct instances, only to arrive at similar solutions as it generates multiple rollouts in parallel. The model

**Table 3** Unit square transformations

$f(x, y)$	
$(x, y)$	$(y, x)$
$(x, 1 - y)$	$(y, 1 - x)$
$(1 - x, y)$	$(1 - y, x)$
$(1 - x, 1 - y)$	$(1 - y, 1 - x)$

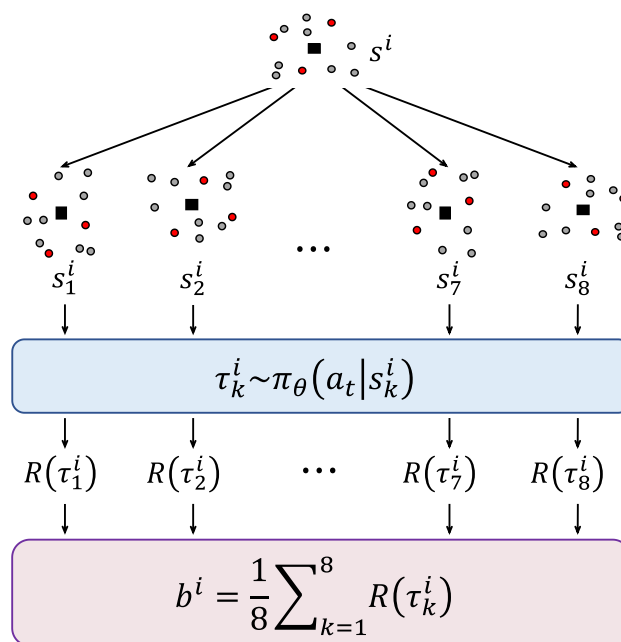
inherently learns to find improved solutions for a given instance based on the local batch mean. The policy model also learns more effective heuristics because the baseline offers a more focused view on a single instance through diverse perspectives. Figure 9 is an illustration of how our local baseline works. We believe that the proposed baseline combines the strengths of multiple-sample baselines and the POMO shared baseline.

#### Comparison with multiple samples with/without replacements

Our baseline does not strictly generate non-duplicate samples. However, it is highly less likely to generate many duplicate samples, especially in the early stages of training, when the policy network  $\pi_{\theta}$  has not yet “learned” much. So, our baseline promotes more “exploration” in the initial learning phase. To see this, we note that each augmented instance is associated with a distinct input embedding in the encoder output ( $h^{(n_{enc})}$ ). Let  $i$  denote the original instance, and let  $k$  and  $j$  denote the augmented instances derived from  $i$ . For  $k (\neq j)$  and  $s_k^i \neq s_j^i$  in raw form,  $h_k^{(n_{enc})i} \neq h_j^{(n_{enc})i}$  in the latent space. Since a trajectory is sampled based on  $h^{(n_{enc})}$ , it is likely that  $\tau_k^i$  is different from  $\tau_j^i$ . Indeed, as training proceeds,  $\pi_{\theta}$  may generate duplicate samples since it learns which action produces high-return trajectories in a more general setting. However, this limitation could be mitigated in large-size problems for which longer trajectories are likely to be unique.

#### Comparison with greedy rollout baseline

In greedy rollout baseline, a solution is generated by running the policy greedily, i.e., at each construction step, the node with the highest probability (where the probability distribution is obtained



**Fig. 9** Proposed REINFORCE baseline

from an earlier version of the model) is visited. This deterministic solution trajectory serves as a baseline in the REINFORCE algorithm. While effective, the greedy rollout baseline incurs an additional forward pass of the earlier model version, which increases computation by 50%. Apart from this, we also empirically found that the greedy rollout baseline entails slightly noisy learning. The current model's (best) performance may not be replicated or generalized to another problem set. This finding is more apparent towards the later stages of training, especially when the model finds it difficult to surpass its greedy self, and there is a noticeable lack of baseline policy updates. At this point, the model does not learn much from the competition with its greedy self.

**Comparison with POMO shared baseline** Compared to the POMO baseline, our approach is more computationally efficient since it uses a fixed local batch size that does not increase with the number of nodes.

### 5.3 Combining with maximum entropy objective

Training the policy model with entropy can smooth out the optimization landscape, speeding up the learning process. In some environments, it yields a better final policy [9]. It also turns out to be robust to internal algorithmic disturbances and external environmental disturbances like dynamics and reward function [8]. We note that robustness to external disturbances is an important factor determining the generalization

capability (i.e., performance on graphs of various sizes). This work combines the maximum entropy RL with our instance-augmentation baseline and shows improved training and inference performance for various problem instances.

We implement the maximum entropy RL as follows. The objective aims to maximize the expected cumulative return augmented by a conditional action entropy as

$$J_{MaxEnt}(\theta) = \sum_t \mathbb{E}_{(s_t, a_t) \sim \rho_\theta^\pi} [r(s_t, a_t) + \alpha \mathbb{H}(\pi_\theta(\cdot | s_t))] \quad (48)$$

where  $\mathbb{H}(\pi_\theta(\cdot | s_t)) = \mathbb{E}_{a_t \sim \pi_\theta} [-\log \pi_\theta(a_t | s_t)] = -\sum_{a_t} [\pi_\theta(a_t | s_t) \log \pi_\theta(a_t | s_t)]$  denotes the Shannon entropy of conditional distribution over actions along the trajectory,  $\rho_\theta^\pi(s_t, a_t)$  is the state-action marginal of trajectory distribution induced by  $\pi_\theta$  and  $\alpha$  is the entropy weight or temperature. The maximum entropy objective function presented in Eq. (48) results in a slightly different gradient [9] (trajectory view):

$$\nabla_\theta J(\theta) = \mathbb{E}_{\tau \sim \pi_\theta} [R(\tau) \nabla_\theta \log P(\tau; \theta) + \alpha \sum_t \nabla_\theta \mathbb{H}(\pi_\theta(\cdot | s_t))] \quad (49)$$

Although Sultana et al. [37] used the entropy maximization term to train the policy with a greedy rollout baseline, we note that its application has not been used with other baselines. By integrating the objective function with entropy and using our instance-augmentation baseline, our policy model learns a more stochastic policy that is applicable in a generalized setting. Algorithm 1 presents our proposed REINFORCE algorithm. The Adam optimizer [38] with a constant learning rate of 0.0001 is used to train the policy model parameters.

**Algorithm 1** Proposed REINFORCE Algorithm (Instance-augmentation baseline with maximum entropy objective)

---

```

1  Training: training set  $S$ , augmentation factor  $K$ , entropy weight  $\alpha$ ,
    $E$  epochs,  $T$  steps per epoch,  $B$  instances per batch
2  Initialize policy network  $\pi_\theta(a | s)$  parameter  $\theta$ 
3  for epoch = 1, ...,  $E$  do
4  |    $s^i \leftarrow \text{SAMPLE}(S) \quad \forall i \in \{1, \dots, B\}$ 
5  |    $\{s_1^i, s_2^i, \dots, s_K^i\} \leftarrow \text{AUGMENT}(s^i)$ 
6  |    $\tau_k^i \leftarrow \text{ROLLOUT}(s_k^i, \pi_\theta) \quad \forall i \in \{1, \dots, B\}, \forall k \in \{1, \dots, K\}$ 
7  |    $b^i \leftarrow \frac{1}{K} \sum_{k=1}^K R(\tau_k^i) \quad \forall i \in \{1, \dots, B\}$ 
8  |    $H(\pi_\theta(\cdot | s_{k,i}^i)) \leftarrow -\sum_{a_t} \pi_\theta(a_t | s_{k,i}^i) \log \pi_\theta(a_t | s_{k,i}^i) \quad \forall i \in \{1, \dots, B\}, \forall k \in \{1, \dots, K\}$ 
9  |    $\nabla_\theta J \leftarrow \frac{1}{BK} \sum_{i=1}^B \sum_{k=1}^K (R(\tau_k^i) - b^i) \nabla_\theta \log P(\tau; \theta) + \alpha \sum_{t=0}^T \nabla_\theta H(\pi_\theta(\cdot | s_{k,i}^i))$ 
10 |    $\theta \leftarrow \text{Adam}(\theta, \nabla_\theta J)$ 
11 end for

```

---

**Table 4** MSTOP problem instances of various sizes

$n$	$N$	$T_{\max}$
10	2	1.5
20	2	2.0
50	3	3.0
70	3	3.0

**Table 5** Permutations of vehicle order. Bold denotes the first vehicle to start routing. The DDTM sequentially begins routing according to the given vehicle order

$N$	Permutations
2	AB/BA
3	ABC/ACB/BCA/ <b>BAC</b> /CAB/CBA

## 6 Experiments and discussion

To establish the effectiveness of our proposed REINFORCE algorithm, we conducted a comprehensive study, first comparing our instance-augmentation baseline with the greedy rollout baseline and subsequently comparing our instance-augmentation baseline with maximum entropy objective against the greedy rollout baseline with entropy. The strengths of our method are substantiated across various problem sets, encompassing TSP, CVRP, and MSTOP, demonstrating consistent improvement in training (in terms of both solution quality and training time) even with increasing problem sizes.

### 6.1 Problem setup and hyperparameters

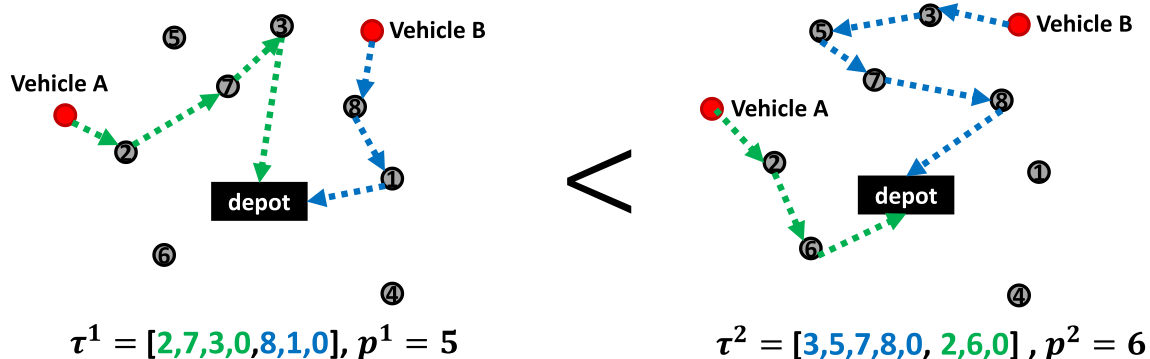
This section describes the controlled experiments to solve the MSTOP using the DDTM. To observe the benefits of our instance-augmentation baseline (over greedy rollout), we conduct an ablation study on classical TSP and CVRP using the

original AM. To this end, we consider three problem/policy pairs – MSTOP/DDTM, TSP/AM, and CVRP/AM. The graph sizes ( $n$ ) of 10, 20, 50, and 70 are set for the MSTOP (Table 4) and we consider these cases with 2 and 3 vehicles. The decision to focus on scenarios involving 2 and 3 vehicles is rooted in our motivation to providing insights into optimizing the efficiency of limited resources in situations where deploying a larger fleet is impractical. For TSP and CVRP, we consider the instances with sizes of 50 and 100. For TSP and CVRP, we consider the instances with sizes of 50 and 100. Furthermore, to check how our proposed training algorithm improves the generalization performance, we test the performance of each AM on problem instances of various sizes.

**Training DDTM to solve MSTOP** We follow the basic problem setup in [3] for the Orienteering Problem (OP), i.e., the coordinates of all customer and depot nodes are randomly sampled within a normalized  $[0,1] \times [0,1]$  world. The prizes of nodes are either initialized as one (constant) or sampled from a uniform distribution between 0 and 1.

Table 4 describes the experimental details, including the graph size ( $n$ ), the number of vehicles ( $N$ ), and the maximum length constraint for each route ( $T_{\max}$ ). Additionally, each vehicle in MSTOP starts at a random location within the same  $[0,1] \times [0,1]$  world and is given a variable remaining tour length (or equivalently fuel amount) with the distance between the current vehicle location and the depot as the lower bound. This setting ensures that the sum of the remaining tour length and the partial tour constructed henceforth is bounded above by  $T_{\max}$ . For all MSTOP cases, the DDTM is initialized with  $n_{\text{enc}} = 4$  and  $n_{\text{dec}} = 2$ , which we found to be an acceptable trade-off between computational load and the quality of learned policy.

For numerical experiments, we train 1,280,000 instances per epoch. Considering the GPU memory constraints, we train 1250 batches of 1024 instances ( $n = 10, 20$ ) for 200 epochs, train 2500 batches of 512 instances ( $n = 50$ ) for 100

**Fig. 10** LEFT: Routing begins with Vehicle A. RIGHT: Routing begins with Vehicle B (optimal tour found)

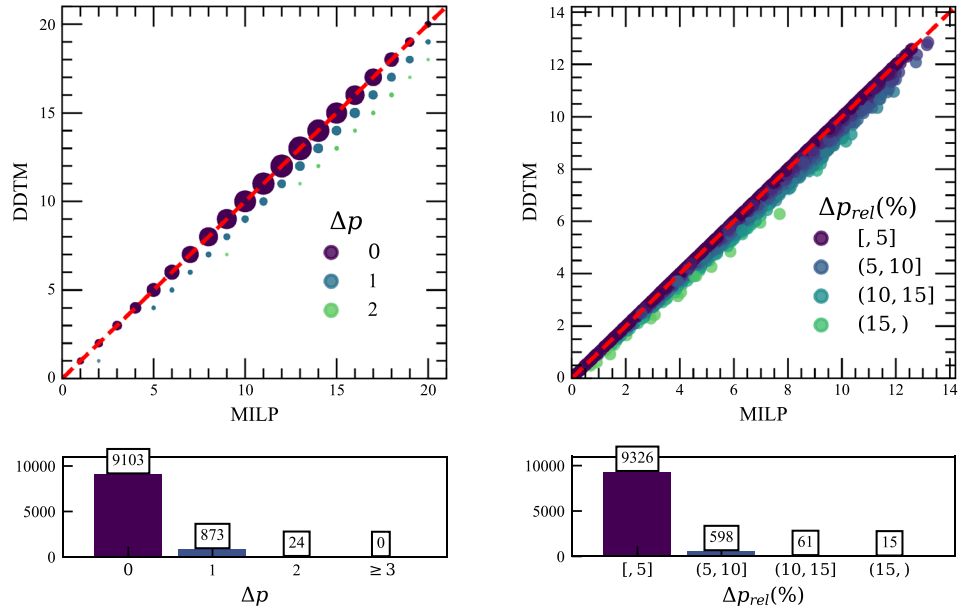
**Table 6** Experimental results on MSTOP (constant prizes; bold: best result)

Method	N=2 Vehicles			N=3 Vehicles								
	n=10			n=20			n=50			n=70		
	Obj.	Gap	Time	Obj.	Gap	Time	Obj.	Gap	Time	Obj.	Gap	Time
Gurobi greedy	5.35	0.00%	(8 m)	11.84	0.00%	(2 h)	–	–	–	–	–	–
	4.82	9.92%	(<1 s)	10.06	15.01%	(<1 s)	34.94	14.81%	(<1 s)	43.58	16.98%	(<1 s)
	5.20	2.76%	(<1 s)	11.20	5.43%	(1 s)	38.72	5.59%	(7 s)	49.08	6.50%	(8 s)
sampling	5.21	2.61%	(<1 s)	11.23	5.13%	(1 s)	39.10	4.66%	(7 s)	49.57	5.55%	(8 s)
	5.32	0.50%	(2 m)	11.66	1.53%	(4 m)	39.28	4.21%	(11 m)	49.21	6.26%	(13 m)
	5.30	1.02%	(8 m)	11.60	1.99%	(16 m)	40.73	0.68%	(1 h)	52.07	0.79%	(1.5 h)
augment	5.30	0.96%	(8 m)	11.62	1.84%	(16 m)	40.84	0.42%	(1 h)	52.34	0.28%	(1.5 h)
	5.29	1.15%	(1 s)	11.50	2.92%	(2 s)	40.06	2.31%	(1 m)	50.91	3.01%	(2 m)
	5.30	1.02%	(1 s)	11.52	2.68%	(2 s)	40.28	1.79%	(1 m)	51.31	2.25%	(2 m)
	5.34	0.24%	(4 s)	11.73	0.90%	(11 s)	40.91	0.25%	(3 m)	52.22	0.52%	(5 m)
	<b>5.34</b>	<b>0.19%</b>	<b>(4 s)</b>	<b>11.75</b>	<b>0.78%</b>	<b>(11 s)</b>	<b>41.01</b>	<b>0.00%</b>	<b>(3 m)</b>	<b>52.49</b>	<b>0.00%</b>	<b>(5 m)</b>

**Table 7** Experimental results on MSTOP (uniformly distributed prizes; bold: best result)

Method	N=2 Vehicles			N=3 Vehicles								
	n=10			n=20			n=50			n=70		
	Obj.	Gap	Time	Obj.	Gap	Time	Obj.	Gap	Time	Obj.	Gap	Time
Gurobi greedy	2.88	0.00%	(9 m)	6.49	0.00%	(2 h)	–	–	–	–	–	–
	2.56	11.18%	(<1 s)	5.40	16.70%	(<1 s)	17.95	16.30%	(<1 s)	22.71	18.81%	(<1 s)
	2.77	3.80%	(<1 s)	6.06	6.57%	(1 s)	20.37	5.05%	(6 s)	26.22	6.27%	(9 s)
sampling	2.78	3.62%	(<1 s)	6.11	5.86%	(1 s)	20.51	4.36%	(6 s)	26.41	5.59%	(9 s)
	2.85	1.12%	(2 m)	6.26	3.44%	(3 m)	20.16	6.00%	(10 m)	25.73	8.02%	(12 m)
	2.85	1.31%	(8 m)	6.33	2.47%	(15 m)	21.36	0.40%	(1 h)	27.82	0.56%	(1.5 h)
augment	2.85	1.31%	(8 m)	6.35	2.07%	(15 m)	21.38	0.31%	(1 h)	27.91	0.26%	(1.5 h)
	2.83	1.97%	(1 s)	6.23	3.96%	(2 s)	21.01	2.02%	(1 m)	27.19	2.82%	(1.5 m)
	2.83	1.78%	(1 s)	6.26	3.50%	(2 s)	21.07	1.75%	(1 m)	27.37	2.18%	(1.5 m)
	2.87	0.62%	(4 s)	6.39	1.52%	(11 s)	21.42	0.16%	(3 m)	27.86	0.43%	(5 m)
	<b>2.87</b>	<b>0.56%</b>	<b>(4 s)</b>	<b>6.40</b>	<b>1.34%</b>	<b>(11 s)</b>	<b>21.45</b>	<b>0.00%</b>	<b>(3 m)</b>	<b>27.98</b>	<b>0.00%</b>	<b>(5 m)</b>

**Fig. 11** DDTM solution quality (optimality gap) on MSTOP20 instances



$$(\Delta p = P_{MILP} - P_{DDTM})$$

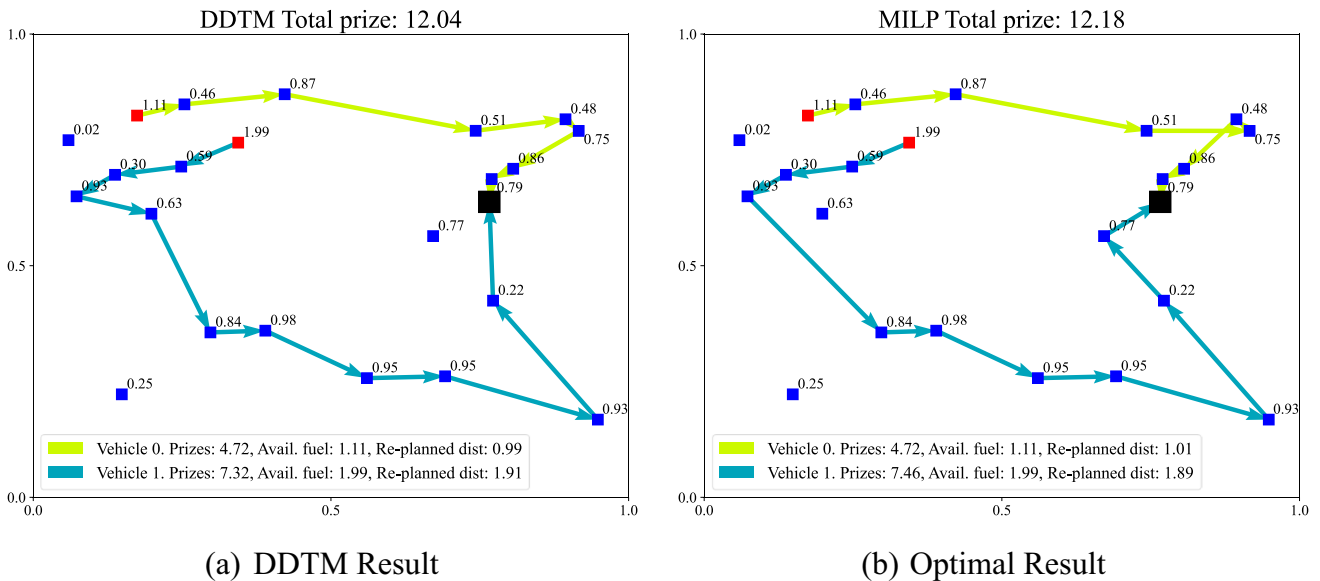
(a) Constant prizes

$$\left( \Delta p_{rel} = \frac{P_{MILP} - P_{DDTM}}{P_{MILP}} \times 100\% \right)$$

(b) Uniformly distributed prizes


epochs, and train 3333 batches of 384 instances ( $n = 70$ ) for 100 epochs. The instance-augmentation baseline uses a batch size reduced by 8, i.e. 128 for  $n = 10$  and  $n = 20$ , 64 for  $n = 50$ , and 48 for  $n = 70$ , so that the total number of training instances is the same. These training instances

are generated randomly on the fly at every epoch to prevent overfitting. After each epoch, we roll out the current model (with greedy decoding) on a held-out validation set of size 10,000 and plot the learning curve to observe the training process.



**Fig. 12** Example solution of MSTOP20 (uniformly distributed prizes). Numerical values next to blue nodes represent node prizes






**Fig. 13** Learning curves for MSTOP20 with uniformly distributed prizes. Dark curves are smoothed results, lighter curves are raw results

**Training AM to solve TSP/CVRP** We adopt the problem setup prescribed in [3]. We used the same hyperparameters for training AM policy network for a fair comparison (except for the application of ‘warmup’).

**Entropy weight** To ensure the benefits of maximum entropy realized in our methodology, we need to use a suitable value for  $\alpha$ . A very large  $\alpha$  value can make the problem close to the maximum entropy problem, whose policy is purely random. On the contrary, if  $\alpha$  is small, premature convergence may occur due to inadequate exploration. The  $\alpha$  value used for training is 0.01 for both MSTOP and TSP/CVRP. We observed that this value works well on MSTOP20 (uniformly distributed prizes) and TSP50.

## 6.2 Inference result

This section presents the performance of DDTM on 10,000 random MSTOP instances. To validate our proposed methodology,



**Fig. 14** Learning curves on TSP50 using the vanilla AM. Dark curves are smoothed results, lighter curves are raw results

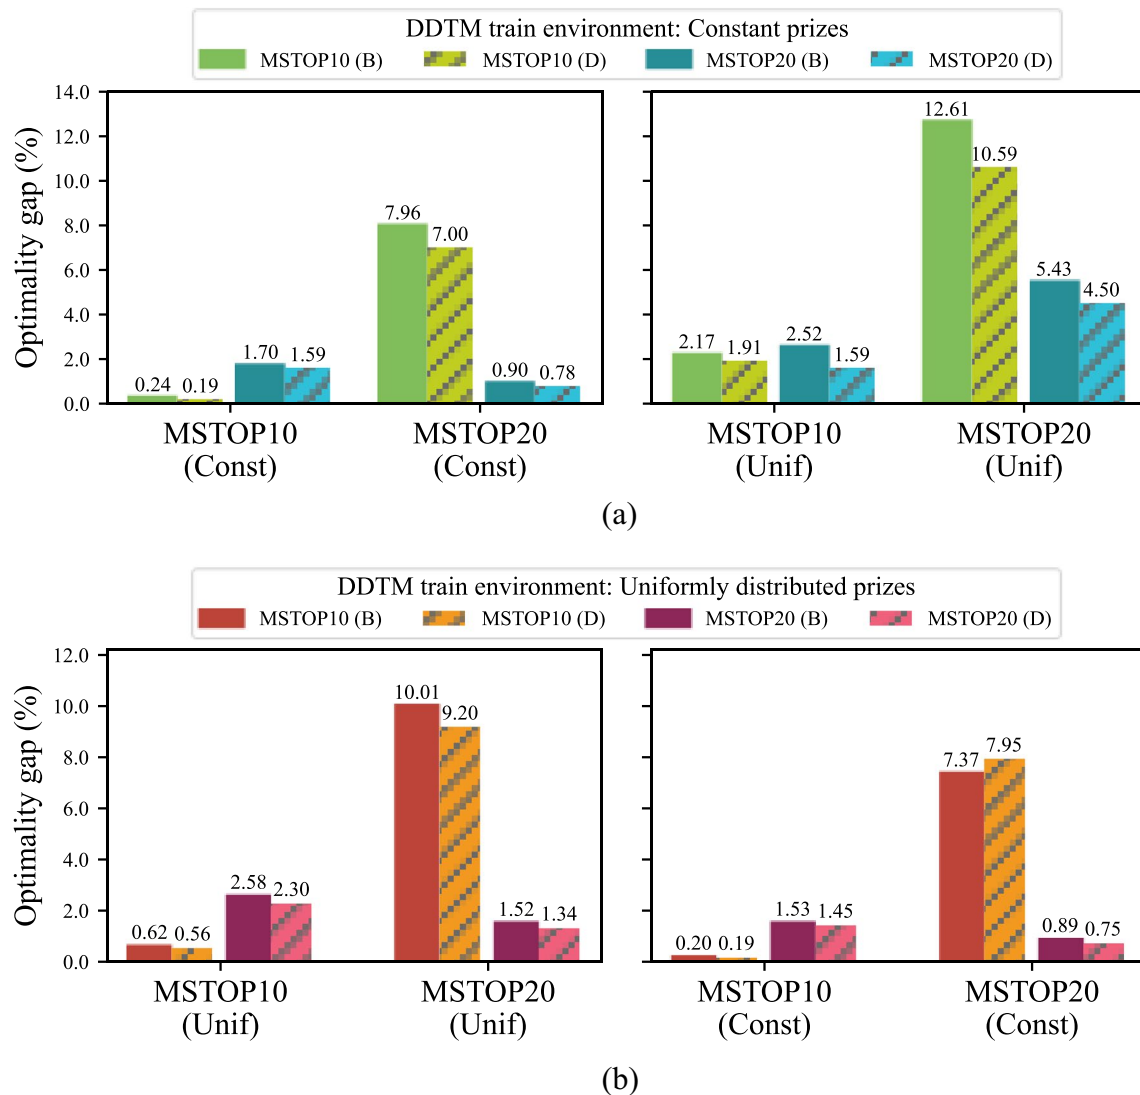
we assess the performance of 1) DDTM trained with our proposed baseline and maximum entropy objective, and 2) DDTM trained with greedy rollout baseline and maximum entropy objective. The following section presents a comprehensive ablation study for various REINFORCE training baselines.

We use three decoding strategies. The *greedy* strategy rolls out a single greedy trajectory for each instance. The *sampling* strategy generates 1280 trajectories (per instance) and selects the best one. Finally, the *instance augmentation* strategy draws multiple greedy trajectories for each instance and selects the best result. To effectively handle inherent asymmetry in the MSTOP solutions, we permute the order of starting vehicles (see Table 5). Then, we generate a single greedy trajectory for each vehicle order and choose the best out of  $N!$  trajectories. To expand the search space, for each permutation, we further rollout eight trajectories about each problem instance (by solving its augmented instances) and select the best out of  $8 \cdot N!$  trajectories. As illustrated in Fig. 10, this increases the chance of finding near-optimal solutions.

To the best of our knowledge, we could not find any algorithms specifically for MSTOP. For  $n$  values of 10 and 20, we compare the results with the optimal solutions obtained using the MILP formulation introduced in Section III (implemented with Gurobi [34]). We also implement the heuristic by Tsiligirides for OP introduced in [39] with slight modification and compare the results. The MILP solution is used as the reference to compute the optimality gap. For larger instances ( $n = 50$  and  $n = 70$ ), it takes prohibitively long to solve the MILP to optimality. Therefore, the best out of the solutions obtained by various methodologies is used as a reference to compute the optimality gap.

Tables 6 and 7 summarize the experimental results for comparison. We report the average of total prizes over 10,000 test MSTOP instances. Using the greedy strategy, the DDTM finds near-optimal solutions with optimality gaps of around 4 – 5%. The optimality gap values for DDTM solutions obtained using the sampling strategy are 1 – 2%. In almost all strategies, the DDTM outperforms the heuristic by Tsiligirides. The DDTM performs best with the  $\times 8N!$  instance augmentation strategy, which finds high-quality solutions much faster than the sampling technique, demonstrating its superiority.

Figure 11 presents the quality (optimality gap) of the solutions obtained using the DDTM trained under the proposed methodology for 10,000 test MSTOP20 instances. The optimality gap of the DDTM solutions is 0% in more than 90% of constant-prize instances. Also, in over 90% of instances with uniformly-distributed prizes, the optimality gap is smaller than 5%. Figure 12 show example solutions of MSTOP20 for different prize distributions. The DDTM inference solutions with “ $\times 8N!$ -augmentation-strategy” are plotted on the left. The corresponding MILP solutions are presented on the right for comparison.



**Fig. 15** Generalization performance of DDTMs trained and tested between MSTOP10 and MSTOP20 environments. Models trained under (a) Constant prizes and (b) Uniformly distributed prizes. Optimality gaps reported as the performance measure

### 6.3 Ablation study

The ablation study analyses the contribution of our proposed training methodology (instance-augmentation baseline with maximum entropy objective) to training policy network models. Specifically, we compare the learning curves using different baselines on the DDTM (for solving MSTOP) and the original AM (for TSP and CVRP). Each learning curve is obtained by evaluating the model on a held-out validation set of 10,000 random instances. The following learning curves are plotted for four different training strategies: greedy rollout baseline (A), greedy rollout baseline with maximum entropy objective (B), instance-augmentation baseline (C), and instance-augmentation baseline with maximum entropy objective (D).

**DDTM & training baselines (MSTOP)** Fig. 13 shows the learning curves of the four training methods – (A) to (D) – on MSTOP20 with uniformly distributed prizes. It can be observed that our proposed baseline (C) helps

**Table 8** Comparison of training time for different training strategies (per epoch, in *min: sec*); training performed on a single 3090Ti GPU

	(A)	(B)	(C)	(D)
TSP50	9:21	8:18	6:28	6:19
TSP100	17:43	19:00	12:30	14:13
CVRP50	13:26	12:08	8:38	8:23
CVRP100	24:07	25:08	16:00	17:41

**Table 9** Test results of vanilla AM trained with different methods

Method	TSP50		TSP100		CVRP50		CVRP100		
	Obj.	Gap	Obj.	Gap	Obj.	Gap	Obj.	Gap	
Concorde*/LKH3*	5.70*	–	7.76*	–	10.38*	–	15.65*	–	
greedy	AM (A)	5.81	1.89%	8.10	4.43%	11.01	6.07%	16.66	6.47%
	AM (B)	5.79	1.65%	8.11	4.50%	10.97	5.70%	16.64	6.32%
	AM (C)	5.80	1.78%	8.08	4.14%	10.96	5.59%	16.65	6.36%
	AM (D)	5.78	1.38%	8.04	3.59%	10.93	5.30%	16.50	5.45%
sampling	AM (A)	5.73	0.48%	7.91	1.97%	10.64	2.47%	16.14	3.13%
	AM (B)	5.71	0.21%	7.93	2.15%	10.57	1.87%	16.15	3.22%
	AM (C)	5.73	0.49%	7.90	1.74%	10.62	2.31%	16.12	2.98%
	AM (D)	<b>5.71</b>	<b>0.15%</b>	7.89	1.69%	<b>10.56</b>	<b>1.75%</b>	<b>16.07</b>	<b>2.69%</b>
× 8 augment	AM (A)	5.72	0.40%	7.93	2.24%	10.69	3.03%	16.30	4.16%
	AM (B)	5.72	0.32%	7.94	2.30%	10.67	2.81%	16.28	4.03%
	AM (C)	5.72	0.37%	7.92	2.01%	10.68	2.86%	16.28	4.05%
	AM (D)	5.71	0.23%	<b>7.89</b>	<b>1.68%</b>	10.66	2.67%	16.19	3.42%

\*: Values reported in [3]

the model learn better policy than both the greedy rollout baseline (A) and its combination with the maximum entropy objective (B). As an added benefit, the instance-augmentation baseline substantially speeds up learning by generating fewer training data. With the maximum entropy objective (D), the proposed methodology significantly outperforms the rest of the methodologies and achieves high validation scores in fewer training epochs, demonstrating the sample efficiency of the proposed training methodology.

**AM & training baselines (TSP, CVRP)** We believe that the proposed methodology is a general technique that can be used instead of the conventional greedy rollout baseline. To validate this, we perform additional experiments on the vanilla AM network using the original code to solve TSP and CVRP. For a fair comparison, we plot the learning curves on the same validation set (with seed 1234) and also report the inference results on the same test set (with seed 4321) used in [3].

Figure 14 shows the learning curves for the original AM with different baselines for TSP50. The instance-augmentation baseline (C) performs comparatively better than the greedy rollout baseline (A) and slightly worse than the greedy rollout baseline with maximum entropy objective (B). However, the proposed methodology (D) substantially improves the quality of the learned policy. Moreover, using the instance-augmentation baseline – (C) and (D) – instead of greedy rollout baseline – (A) and (B) – significantly reduces the per-epoch training time by over 30% (see Table 8). Our proposed method is thus effective in expediting per-epoch training time while simultaneously keeping competitive performance. This indeed supports the

claim that our proposed method strikes a favorable balance between training speed and overall performance.

Table 9 summarizes the inference test results on TSP and CVRP. Our proposed methodology (D) outperforms the other training methods across all decoding strategies in all cases. In particular, the proposed approach is comparable to the state-of-the-art POMO method in terms of the optimality gap. The best performance for TSP50 obtained by the proposed approach (optimality gap: 0.15%, sampling) is better than that by the POMO inference without augmentation (0.24% [4]). Similarly, in CVRP50 instances, the best result obtained by the proposed method (1.75%; sampling) outperforms the POMO inference with a single trajectory (3.52% [4]). Even on large instances ( $n = 100$ ), the proposed methodology (D) shows improvement over all decoding strategies.

## 6.4 Generalization result

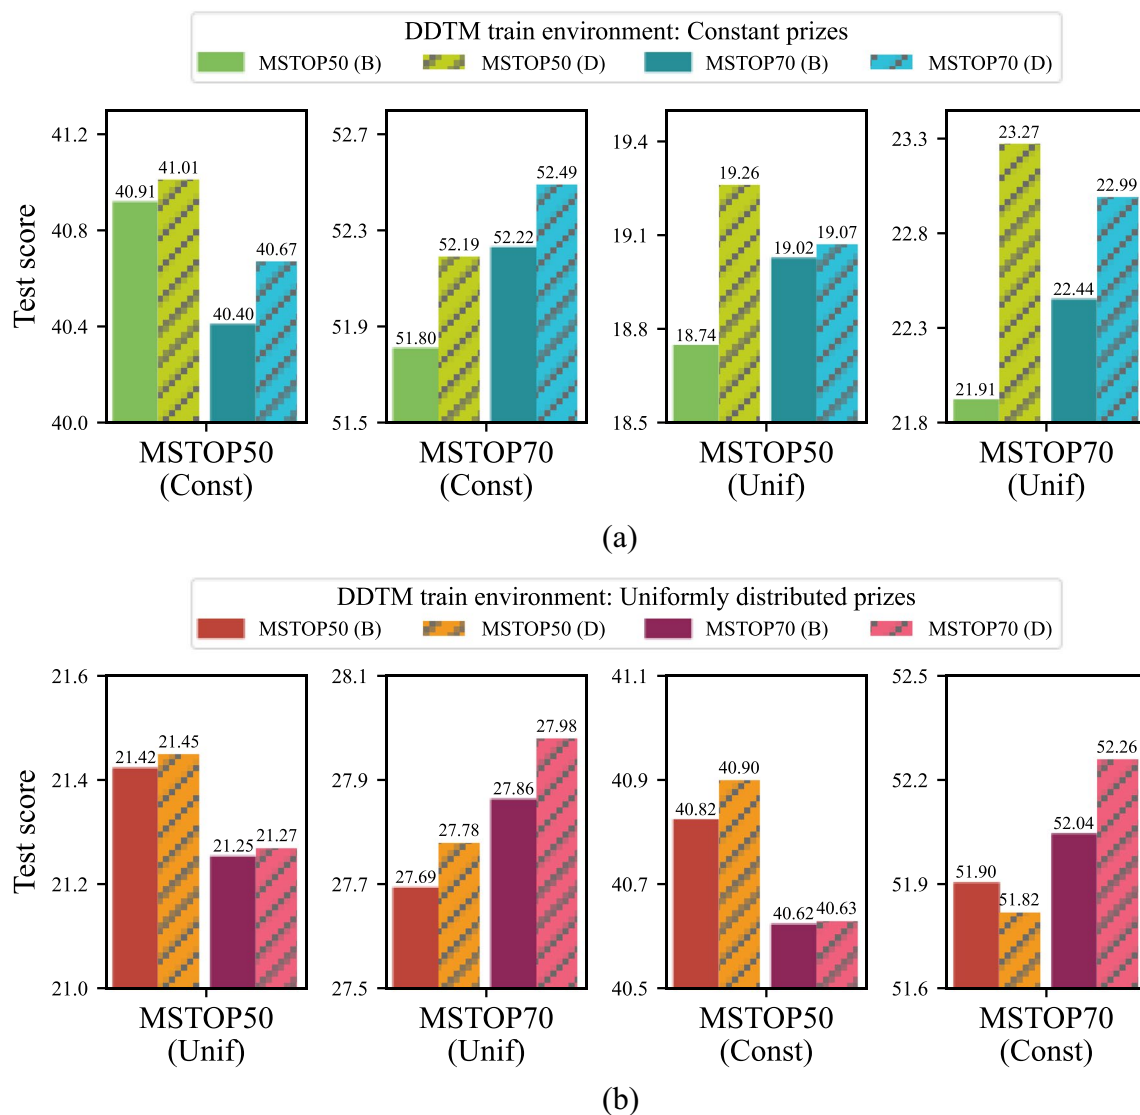
This section discusses the generalization capability of our training methodology. Kool et al. [3] demonstrated that the AM and greedy rollout baseline can be generalized to problems with different graph sizes, although the error increases as the graph size increases. Since training with the maximum entropy objective is known to improve the model's robustness, we conduct a comparative study on generalization performance between greedy rollout with maximum entropy objective (B) and our proposed methodology (D) to see how our proposed methodology reduces generalization error. Note that the generalization results are reported according to the instance-augmentation decoding strategy on the same test datasets as in the previous sections.

Figure 15 illustrates the generalization performance of DDTMs trained on MSTOP10 and MSTOP20 environments

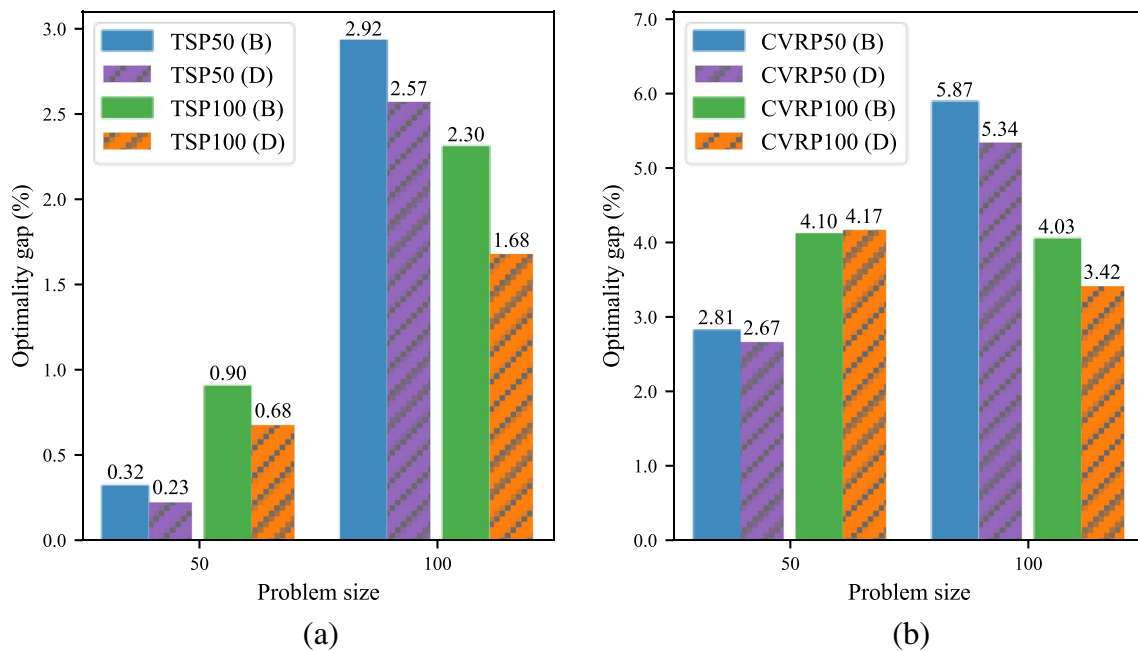
for  $N=2$  vehicles where the horizontal axis represents the test environment (i.e. prize distribution and graph size) and the vertical axis refers to the optimality gap. Part (a) reports the performance of DDTM trained under constant prizes whereas part (b) corresponds to that of DDTM trained under uniformly distributed prizes. We observe that the models naturally perform best when tested under the same conditions as the training environment. However, optimality gaps tend to increase when tested on different graph sizes. In general, the proposed methodology (D) shows better generalization than the conventional method (B) in terms of reduced optimality gaps for changing graph sizes. Moreover, we also observe that models trained under uniformly distributed prizes generalize better than the counterparts trained under constant prizes when tested on environments with different

prize distributions. This is not surprising since uniformly distributed prizes can be seen as a generalized version of constant prizes, and the problems with constant prizes are generally considered easier to solve. One exception is the case of DDTM trained under MSTOP10 with uniformly distributed prizes being tested on MSTOP20 with constant prizes, where the model trained using the proposed methodology (D) performs worse than the conventional approach (B). The reason behind this result might be attributed to using entropy weight  $\alpha$  tuned for MSTOP20 (with uniformly distributed prizes) problems.

Figure 16 presents the generalization result for DDTM trained on MSTOP50 and MSTOP70 environments for  $N=3$  vehicles where the vertical axis represents the test score. Similar to Fig. 15, the proposed methodology (D) generally



**Fig. 16** Generalization performance of DDTMs trained and tested between MSTOP50 and MSTOP70 environments. Models trained under (a) Constant prizes and (b) Uniformly distributed prizes. Test scores reported as performance measure



**Fig. 17** Generalization results for (a) TSP and (b) CVRP

performs better for both changing graph sizes and prize distributions, as evidenced by larger test scores. The degree of improvement is more apparent for large-scale problems, demonstrating that the proposed methodology generalizes well with scalability on graph size.

Figure 17 presents the generalization performance for TSP and CVRP versus the graph size. For both TSP and CVRP, the proposed methodology (D) shows better generalization performance (reduced optimality gaps) except for the CVRP100 model on graph size  $n=50$ , which is likely a result of using entropy weight  $\alpha$  that is tuned for TSP50. From the various tests on different routing problems, it can be observed that our proposed methodology generally results in an improved generalization performance compared to the existing conventional method.

## 7 Conclusion

The Multi-Start Team Orienteering Problem (MSTOP) is introduced to address the routing problems arising in dynamic environments. An attention-based policy network model referred to as the Deep Dynamic Transformer Model (DDTM) is proposed to solve the MSTOP. The proposed learning procedure modifies the REINFORCE algorithm by introducing a new baseline with instance-augmentation and combining it with the maximum entropy objective, improving its learning efficiency and inference capability. A set of numerical experiments comparing the performance of the proposed procedure with existing

methodologies demonstrates its effectiveness. For a suitable value of entropy weight, the instance-augmented baseline outperforms the conventional greedy rollout baseline both in terms of inference performance, generalization performance and training speed. The test result indicates that the proposed approach performs comparably to the current state-of-the-art POMO baseline while requiring less computational resources. The procedure is further applied to classical TSP and CVRP, showing the potential to be a general technique for solving various routing problems. It would be interesting to apply the proposed methodology to other asymmetric CO problems, such as the Multi-Depot VRP and Multi-Depot MSTOP, where the order of vehicles break the symmetry in solution representations. Applying the proposed approach to missions involving the cooperation between agents would be also a meaningful extension of this study [37]. Another promising subject for future study is to handle the instance-augmentation inference for problems with many vehicles. We can tackle these large problems by breaking them into smaller, more manageable subproblems. By doing so, we can utilize our model (that is trained for 2 or 3 vehicles) to iteratively solve portions of the larger problem. Subsequently, we can then concatenate the individual solutions to generate a comprehensive solution for the entire fleet. While this iterative approach may not yield optimal solutions, it may produce near-optimal solutions rapidly, as our model solves in the order of 10 ms. We also acknowledge that the current implementation of DDTM architecture is heavy, resulting in a longer training time compared to the original AM. One possible resolution would

be to “compress” the model [38, 39] for efficient training and inference.

**Funding** Open Access funding enabled and organized by KAIST. This work was prepared at the Korea Advanced Institute of Science and Technology, Department of Aerospace Engineering, under a research grant from the National Research Foundation of Korea (2020R1A2C1005037).

## Declarations

**Conflicts of interest** The authors have no competing interests to disclose.

**Code** Our DDTM implementation and training methodology code based on the instance-augmentation baseline with maximum entropy is publicly available at <https://github.com/leedh0124/Deep-Dynamic-Transformer-Model-for-Multi-Start-Team-Orienteering-Problem>.

**Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article’s Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article’s Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

## References

- Coutinho WP, Battarra M, Fliege J (2018) The unmanned aerial vehicle routing and trajectory optimisation problem, a taxonomic review. *Comput Ind Eng* 120:116–28. <https://doi.org/10.1016/j.cie.2018.04.037>
- Rojas Vilorio D, Solano-Charris EL, Muñoz-Villamizar A, Montoya-Torres JR (2021) Unmanned aerial vehicles/drones in vehicle routing problems: a literature review. *Int Trans Oper Res* 28:1626–57. <https://doi.org/10.1111/itor.12783>
- Kool W, Hoof HV, Welling M (2019) Attention, Learn to Solve Routing Problems! In: 2019 International Conference on Learning Representations (ICLR). <https://doi.org/10.48550/arXiv.1803.08475>
- Kwon Y-D, Choo J, Kim B, Yoon I, Gwon Y, Min S (2020) Pomo: Policy optimization with multiple optima for reinforcement learning. In: *Advances in Neural Information Processing Systems (NeurIPS)*, 21188–98. <https://doi.org/10.48550/arXiv.2010.16011>
- Vaswani A, Shazeer N, Parmar N, Uszkoreit J, Jones L, Gomez AN, Kaiser Ł, Polosukhin I (2017) Attention is all you need. In: *Proceedings of the 31st International Conference on Neural Information Processing Systems*, 6000–10. Curran Associates Inc, Long Beach, California, USA. <https://dl.acm.org/doi/10.5555/3295222.3295349>
- Bresson X, Laurent T (2021) The transformer network for the traveling salesman problem. In: *ArXiv*. <https://doi.org/10.48550/arXiv.2103.03012>
- Peng B, Wang J, Zhang Z (2020) A deep reinforcement learning algorithm using dynamic attention model for vehicle routing problems. In: <https://doi.org/10.48550/arXiv.2002.03282>
- Eysenbach B, Levine S (2021) Maximum entropy rl (provably) solves some robust rl problems. In: <https://doi.org/10.48550/arXiv.2103.06257>
- Ahmed Z, Le Roux N, Norouzi M, Schuurmans D (2019) Understanding the impact of entropy on policy optimization. In: *International conference on machine learning*, 151–60. PMLR. <https://doi.org/10.48550/arXiv.1811.11214>
- Archetti C, Grazia Speranza M, Vigo D (n.d.) Chapter 10: Vehicle Routing Problems with Profits. In: *Vehicle Routing (MOS-SIAM Series on Optimization)*. <https://epubs.siam.org/doi/abs/10.1137/1.9781611973594.ch10>
- Archetti C, Bianchessi N, Speranza MG (2013) Optimal solutions for routing problems with profits. *Discret Appl Math* 161:547–57. <https://doi.org/10.1016/j.dam.2011.12.021>
- Vansteenwegen P, Souffriau W, Van Oudheusden D (2011) The orienteering problem: a survey. *Eur J Oper Res* 209:1–10. <https://doi.org/10.1016/j.ejor.2010.03.045>
- Butt SE, Ryan DM (1999) An optimal solution procedure for the multiple tour maximum collection problem using column generation. *Comput Oper Res* 26:427–41. [https://doi.org/10.1016/S0305-0548\(98\)00071-9](https://doi.org/10.1016/S0305-0548(98)00071-9)
- Boussier S, Feillet D, Gendreau M (2007) An exact algorithm for team orienteering problems. *4OR* 5:211–30. <https://doi.org/10.1007/s10288-006-0009-1>
- Bono G, Dibangoye JS, Simonin O, Matignon L, Pereyron F (2021) Solving multi-agent routing problems using deep attention mechanisms. *IEEE Trans Intell Transp Syst* 22:7804–13. <https://doi.org/10.1109/TITS.2020.3009289>
- Lin S-W (2013) Solving the team orienteering problem using effective multi-start simulated annealing. *Appl Soft Comput* 13:1064–73. <https://doi.org/10.1016/j.asoc.2012.09.022>
- Lin S-W, Yu VF (2017) Solving the team orienteering problem with time windows and mandatory visits by multi-start simulated annealing. *Comput Ind Eng* 114:195–205. <https://doi.org/10.1016/j.cie.2017.10.020>
- Hapsari I, Surjandari I, Komarudin K (2019) Solving multi-objective team orienteering problem with time windows using adjustment iterated local search. *J Ind Eng Int* 15:679–93. <https://doi.org/10.1007/s40092-019-0315-9>
- Bello I, Pham H, Le QV, Norouzi M, Bengio S (2017) Neural Combinatorial Optimization with Reinforcement Learning. In: *2017 International Conference on Learning Representations (ICLR)*. <https://doi.org/10.48550/arXiv.1611.09940>
- Vinyals O, Fortunato M, Jaitly N (2015) Pointer networks. In: *Advances in Neural Information Processing Systems (NeurIPS)*. <https://doi.org/10.48550/arXiv.1506.03134>
- Khalil E, Dai H, Zhang Y, Dilkina B, Song L (2017) Learning combinatorial optimization algorithms over graphs. In: *Advances in Neural Information Processing Systems (NeurIPS)*. <https://doi.org/10.48550/arXiv.1704.01665>
- Nazari M, Oroojlooy A, Snyder L, Takác M (2018) Reinforcement learning for solving the vehicle routing problem. In: *Advances in Neural Information Processing Systems (NeurIPS)*. <https://doi.org/10.48550/arXiv.1802.04240>
- Deudon M, Cournut P, Lacoste A, Adulyasak Y, Rousseau LM (2018) Learning heuristics for the tsp by policy gradient. In: *International conference on the integration of constraint programming, artificial intelligence, and operations research*, 170–81. Springer. [https://doi.org/10.1007/978-3-319-93031-2\\_12](https://doi.org/10.1007/978-3-319-93031-2_12)
- Cappart Q, Moisan T, Rousseau L-M, Prémont-Schwarz I, Cire A (2020) Combining reinforcement learning and constraint programming for combinatorial optimization. In: *ArXiv*. <https://doi.org/10.48550/arXiv.2006.01610>
- Li J, Ma Y, Gao R, Cao Z, Lim A, Song W, Zhang J (2021) Deep reinforcement learning for solving the heterogeneous capacitated

- vehicle routing problem. *IEEE Trans Cybern.* <https://doi.org/10.48550/arXiv.2110.02629>
26. Li K, Zhang T, Wang R, Wang Y, Han Y, Wang L (2021) Deep reinforcement learning for combinatorial optimization: covering salesman problems. *IEEE Trans Cybern.* <https://doi.org/10.48550/arXiv.2102.05875>
  27. Xu Y, Fang M, Chen L, Gangyan X, Yali D, Zhang C (2021) Reinforcement learning with multiple relational attention for solving vehicle routing problems. *IEEE Trans Cybern.* <https://doi.org/10.1109/TCYB.2021.3089179>
  28. Pan W, Liu SQ (2023) Deep reinforcement learning for the dynamic and uncertain vehicle routing problem. *Appl Intell* 53:405–22. <https://doi.org/10.1007/s10489-022-03456-w>
  29. Wang Q (2022) VARL: a variational autoencoder-based reinforcement learning Framework for vehicle routing problems. *Appl Intell* 52:8910–23. <https://doi.org/10.1007/s10489-021-02920-3>
  30. Joshi CK, Laurent T, Bresson X (2019) On learning paradigms for the travelling salesman problem. In *ArXiv*. <https://doi.org/10.48550/arXiv.1910.07210>
  31. Kool W, van Hoof H, Welling M (2019) Buy 4 reinforce samples, get a baseline for free! In: *ICLR 2019 Deep Reinforcement Learning meets Structured Prediction Workshop*. <https://openreview.net/forum?id=r1lgTGL5DE>. Accessed 23 Jun 2022
  32. Kool W, van Hoof H, Welling M (2019) Stochastic beams and where to find them: The gumbel-top-k trick for sampling sequences without replacement. In: *International Conference on Machine Learning (ICML)*, 3499–508. PMLR. <https://doi.org/10.48550/arXiv.1903.06059>
  33. Croes GA (1958) A method for solving traveling-salesman problems. *Oper Res* 6:791–812. <https://www.jstor.org/stable/167074>. Accessed 23 Jun 2022
  34. Gurobi Optimization, LLC (2018) Gurobi optimizer reference manual. <https://www.gurobi.com>
  35. Williams RJ (1992) Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Mach Learn* 8:229–56. <https://doi.org/10.1007/BF00992696>
  36. Sutton RS, Barto AG (2018) Reinforcement learning: An introduction. Cambridge, MIT press
  37. Sultana N, Chan J, Sarwar T, Qin AK (2021) Learning to Optimise Routing Problems using Policy Optimisation. In: *2021 International Joint Conference on Neural Networks (IJCNN)*, 1–8. IEEE. <https://doi.org/10.1109/IJCNN52387.2021.9534010>
  38. Kingma DP, Ba J (2014) Adam: A method for stochastic optimization. In. <https://doi.org/10.48550/arXiv.1412.6980>
  39. Tsiligirides T (1984) Heuristic methods applied to orienteering. *J Oper Res Soc* 35:797–809. <https://www.jstor.org/stable/2582629>. Accessed 23 Jun 2022

**Publisher's Note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

**Dong Ho Lee** is a research scientist within the Department of Aerospace Engineering at KAIST. He received both his bachelor's degree and master's degree in Aerospace Engineering from KAIST in 2017 and 2022, respectively. His primary focus lies at the intersection of optimization, machine learning and control theory, especially in the areas of interactive autonomous systems.

**Jaemyung Ahn** is a professor of aerospace engineering at the Korea Advanced Institute of Science and Technology (KAIST; Daejeon, South Korea). His research interests include design and optimization of aerospace missions/systems, complexity analysis, and flight mechanics. From 2008 to 2010, he worked for Bain & Company as a management consultant. He also worked for the Korea Aerospace Research Institute from 1999 to 2004 as a system engineer. He received his B.S. and M.S. degrees from Seoul National University in 1997 and 1999, and Ph.D. degree in aeronautics and astronautics from MIT in 2008.