

Onboard Mission Replanning for Adaptive Cooperative Multi-Robot Systems

Elim Kwan¹, Rehman Qureshi², Liam Fletcher¹, Colin Laganier¹, Victoria Nockles¹, Richard Walters¹

Abstract—Cooperative autonomous robotic systems have significant potential for executing complex multi-task missions across space, air, ground, and maritime domains. But they commonly operate in remote, dynamic and hazardous environments, requiring rapid in-mission adaptation without relying on fragile or slow communication links to centralized compute. Fast, onboard replanning algorithms are therefore essential to enhance resilience for these systems, but do not yet exist. Reinforcement Learning (RL) shows strong promise for efficiently solving mission planning tasks formulated as Travelling Salesperson Problems (TSPs), but existing methods: 1) are unsuitable for replanning, where agents do not start at a single location; 2) do not allow cooperation between agents; 3) are unable to model tasks with variable durations; or 4) lack practical considerations for onboard deployment. Here we address this gap by defining the Cooperative Mission Replanning Problem as a novel adaptation of multiple TSP, and develop a new encoder/decoder-based RL model to solve it effectively and efficiently. Using a simple example of cooperative drones, we show our replanner consistently (90% of the time) maintains performance within 10% of the state-of-the-art LKH3 heuristic solver, whilst running 85-370 times faster on a Raspberry Pi. This work paves the way for increased resilience in autonomous multi-agent systems.

Index Terms—Path Planning for Multiple Mobile Robots or Agents, Task Planning, Multi-Robot Systems, AI and Machine Learning in Manufacturing and Logistics Systems, Aerial Systems: Perception and Autonomy

I. INTRODUCTION

TEAMS of autonomous robots, including uncrewed aerial, underwater, surface, or ground vehicles (UAVs, UUVs, USVs, UGVs) and multi-satellite constellations, have huge potential for efficiently and safely completing a diverse range of complex missions. In comparison to single-agent operations, such multi-agent systems have increased robustness to failure

Received 2 June 2025; accepted 5 October 2025. Date of publication 30 October 2025; date of current version 12 November 2025. This article was recommended for publication by Associate Editor J. Rojas and Editor J. Kober upon evaluation of the reviewers’ comments. The work of Rehman Qureshi was supported by an international internship on behalf of the Science, Mathematics and Research for Transformation (SMART) Scholarship—for-Service Program within the OUSD/R&E (The Under Secretary of Defense-Research and Engineering), National Defense Education Program (NDEP)/BA-1, Basic Research. This work was supported by the Alan Turing Institute’s Defence and National Security Programme through a partnership with Dstl. Computing resources were provided by Baskerville Tier 2 HPC, funded by EPSRC and UKRI through Advanced Research Computing, University of Birmingham under Grant EP/T022221/1 and Grant EP/W032244/1. (Corresponding authors: Elim Kwan ; Rehman Quresh.)

¹The Alan Turing Institute, British Library, 96 Euston Rd., London NW1 2DB, United Kingdom ekwan@turing.ac.uk

²Department of Aerospace Engineering, Auburn University, Auburn, AL 36849, United States rsg0001@auburn.edu

Code available at: <https://github.com/alan-turing-institute/graph-attention-replanner>

Digital Object Identifier (DOI): 10.1109/LRA.2025.3627093

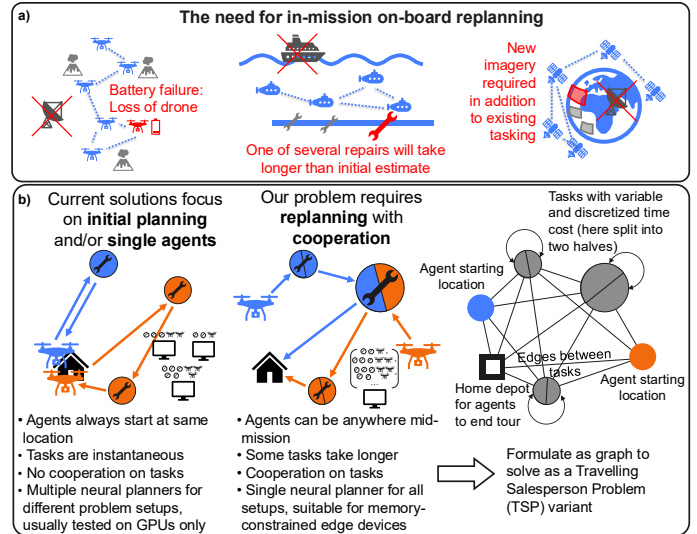


Fig. 1. (a) A cartoon illustrating the need for in-mission, on-board replanning through three examples that all feature major unforeseen changes and limited ground station communication in dynamic, challenging environments. (b) A cartoon illustrating differences between a more traditional initial planning, non-cooperative problem (left) and our cooperative replanning problem (middle). The right-hand diagram shows how this replanning problem can be formulated in a graph format as a variant of the multiple Travelling Salesperson Problem (mTSP).

[1], greater areal coverage [2], and faster mission completion [3]. Furthermore, they can cooperatively tackle complex missions comprised of multiple tasks, where these tasks can be as diverse as capturing images [4], manufacturing/ repairing items or transporting materials [5], or making environmental measurements [6]. However, these systems also commonly operate in environments that are both remote or isolated and also highly dynamic and hazardous, e.g. in nuclear power plants [7], over active volcanoes [6] or in congested low earth orbits in space [4]. Operating in such challenging environments requires these systems to be able to adapt and replan in response to major unexpected changes that impact the mission plan [8] such as the identification of new tasks or the loss or failure of a robotic agent, without relying on fragile or slow communication links to centralized ground stations [9] (Fig. 1a). Therefore, lightweight and fast replanning algorithms that can be deployed on-board cooperative robotic systems are critical for enabling autonomous adaptability during a mission; but these do not yet exist.

Because mission planning with multiple agents requires finding the optimal tour path for each agent while minimizing the length of the longest tour, previous studies have often formulated and solved it as a variant of the multiple Traveling Salesman Problem (mTSP) with a min-max objective [20]. However, our particular planning problem has three major

TABLE I

COMPARING THE LITERATURE ON USING ATTENTION MECHANISM AND REINFORCEMENT LEARNING FOR TRAVELLING SALESPERSON PROBLEMS

Paper	Problem	Encoder	Decoder	Algo	Key Adaptation			Practical Consideration	
					Vary Start Locs	Vary Time Cost	Collab. Tasking	Generalization	On Edge Devices
[10]	TSP	RNN	Ptr. Net.	REINFORCE	×	×	×	×	×
[11]	TSP	Attention	Ptr. Net.	REINFORCE (Rollout)	×	×	×	×	×
[12]	TSP	RNN + Graph Attention	Ptr. Net.	REINFORCE (Rollout)	×	×	×	×	×
[13]	CVRP	Graph Attention	Attention Model	REINFORCE (Rollout)	×	×	×	×	×
[14]	multi-tasks VRP	Graph Attention	Attention Model	POMO + MoE Layer	×	×	×	✓(Goal)	×
[15]	mTSP with Multiple Depots	Graph Attention	Ptr. Net.	REINFORCE	✓	×	×	×	×
[16]	mTSP with Multiple Depots	MoE Net.	Ptr. Net.	Pretrained	✓	×	×	✓(Agent)	✓
[17]	Split Delivery VRP	Convolutional	Ptr. Net.	REINFORCE	✓	×	✓	×	✓
[18]	VRP (Min-Max)	Graph Attention (Partition/Navigation)	Attention Model	REINFORCE (Rollout)	✓	×	×	✓(Depot)	×
[19]	Heterogeneous CVRP (Min-Max)	Graph Attention	Attention Model (Parallel)	REINFORCE	✓	×	×	✓(Agent)	×
This Study	CMRP	Graph Attention	Attention Model	REINFORCE (Rollout)	✓	✓	✓	✓(Task, Agent, Collab.)	✓

VRP refers to the Vehicle Routing Problem; CVRP refers to the Capacitated VRP; Ptr. Net. refers to the Pointer Network; MoE refers to the Mixture of Experts technique; POMO is a modified REINFORCE algorithm which uses policy optimization with multiple optima.

differences to traditional mTSP (Fig. 1b), and in addition requires practical consideration for edge deployment onboard robotic systems. First, variable start locations for agents must be considered, as each agent may be at any location within the mission domain when replanning is required. Second, nodes with variable time costs are necessary to reflect the common real-world situation where mission time depends not just on how long it takes for an agent to travel to a task, but also how long it takes to complete that task. This models a wide range of applications where tasks have differing execution times [21]–[23]. Third, cooperative tasking is required, such that multiple agents can work together on a single task to complete it faster, even if this joint contribution is asynchronous rather than simultaneous. No previous studies have addressed this important combination of multiple adaptations (see “Key Adaptation” columns of Table I), which we name the Cooperative Mission Replanning Problem (CMRP). Finally, to enable onboard replanning, any model for solving CMRP must be suitably small and fast [24]. These limitations also necessitate a single generalized model that can solve a range of problem sizes, as the number of available agents and remaining tasks are likely to vary during a mission [25] and it is impractical to store one model per problem size on edge devices.

The mTSP class of problems are NP-hard [26] with the total number of possible mission plans scaling as $\frac{(n+m-1)!}{2^{(m-1)!}}$, where n is the number of tasks and m is the number of agents. Reinforcement Learning (RL) techniques, especially those using Attention mechanisms (e.g. GAT; Graph Attention Networks), have proven to be effective at overcoming these scaling issues and efficiently solving TSP [10]–[12], mTSP [27], min-max Vehicle Routing Problems (VRP) [18], [19], or other similar multi-agent problems [13], [14].

These methods are able to scale to larger problems than exact solvers such as CPLEX [28], are faster at execution time and generalize better than heuristic methods [27], [29], [30] such as the Lin-Kernighan heuristic [31], show more consistent and generalizable performance than metaheuristic methods [32], [33], and can be deployed on edge devices as lightweight models [34].

But despite its real-world importance, no work to date has addressed the replanning problem as outlined above. Several studies have made one of the required adaptations, e.g. multiple depots, where agents start and finish at different locations [15], [16], [18] or cooperative tasking [17], but none

have considered more than one together, and no previous studies have considered nodes with variable time costs (Table I), which is a fundamental requirement. Similarly, most prior research has focused on high-performance computing systems with GPUs, with limited investigation of deployment on ground robots [16] or CPUs [17]. And whilst [14], [17]–[19], developed a single neural solver for different problem variants, numbers of agents or depots, they did not consider varying numbers of tasks. Our work fills this critical gap, developing a cooperative replanning solver that is suitable for edge deployment and real-world on-board application.

- We are the first to formulate the Cooperative Mission Replanning Problem (CMRP); a novel and important real-world variant of the mTSP adapted to include variable agent start locations, nodes with variable time cost, and cooperative tasking.
- We propose a new lightweight mission replanner based on the GAT and Attention Model that can effectively and efficiently solve the CMRP on edge hardware.
- Using an example of UAVs completing tasks in a 2D domain, our approach produces solutions that are consistently within 10% of those generated by the state-of-the-art LKH3 heuristic solver, but runs 85-370 times faster on a Raspberry Pi. Our single generalized model has high performance across a range of problem sizes.
- Our framework also valuably supports future extension to additional important capabilities (e.g. multi-objective, probabilistic, anticipatory replanning).

Our new formulation and approach enables fast and robust in-mission, on-device replanning for the first time, for cooperative teams of robots. This has the potential to improve the resilience of diverse real-world multi-agent robotic systems.

II. METHODOLOGY

We formally define the CMRP in Section II-A and detail how it differs from the traditional mTSP. We then introduce our GAT-based solution to the CMRP in Section II-B.

A. The Cooperative Mission Replanning Problem (CMRP)

As detailed in the introduction, CMRP varies from the traditional mTSP in three key ways: 1) flexible start locations for robotic agents; 2) non-zero and variable time costs associated with each task; and 3) cooperative tasking. Additionally, there

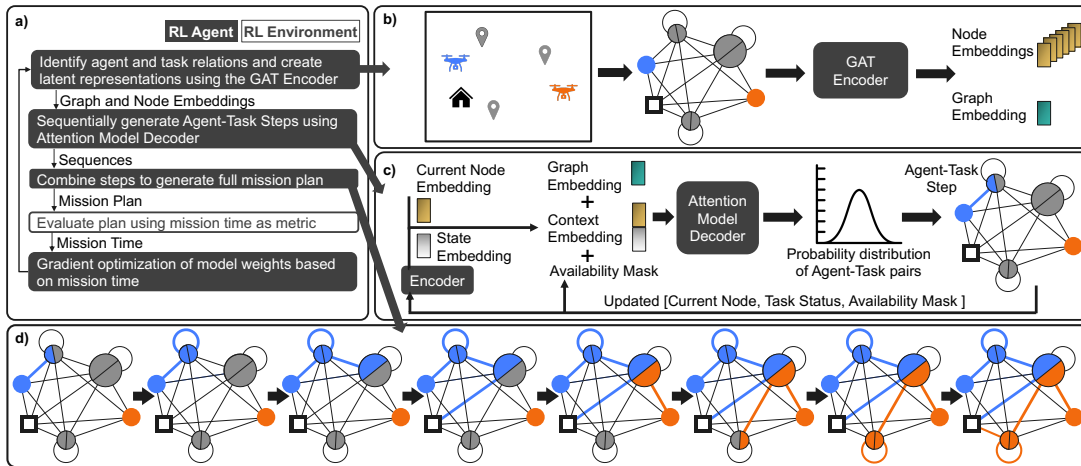


Fig. 2. (a) The RL training process. (b) The transformation of the mission planning data into higher-level node embeddings and a graph embedding using a GAT encoder. (c) The Attention Model decoder generates a probability distribution for available agent-task combinations based on the graph embedding, the current node embedding, the current state of the mission (in embedding space) and the availability mask. The agent-task step with the highest probability is then selected as the next step. (d) The sequential generation of all steps of the mission plan.

is a practical consideration of model generalization, which will be discussed later in Section III-B.

Traditional mTSP with m salespeople or agents and n cities or tasks can be defined as a graph $G = (V, E)$ where the vertices $V = \{v_{depot}, v_1, \dots, v_n\}$ include all n task locations as well as v_{depot} , which serves as the single start and end location for all agents. The edge set E is defined as $E = \{(v_i, v_j) : v_i, v_j \in V, i \neq j\}$ and represents all possible valid paths between tasks. For all edges, $T = (d_{ij})$ represents the travel time associated with traversing from task v_i to v_j , assuming a constant travel speed of 1 m/s.

First, we incorporate flexible start locations into this formulation by adding m start locations to the vertex set, similar to the multiple depot problems [15], [16]. Our new vertex set is $V' = \{v_{depot}, v_{start1}, \dots, v_{startm}, v_1, \dots, v_n\}$, where v_{startk} is the start location of the k -th agent, and v_i is the location of the i -th task. Moreover, unlike formulations in the literature where agents complete their tours at their respective starting locations, our v_{depot} serves as a common destination for all agents, simulating a shared landing space for drones. Two constraints are set to ensure that agents always begin at their designated start locations: (1) $d_{depot i} = 0$; (2) $d_{i startj} = \infty$; where $v_i \in \{v_1, \dots, v_n\}$ and $v_{startj} \in \{v_{start1}, \dots, v_{startm}\}$. Hence, the T cost matrix is no longer symmetric and therefore changes the problem to an asymmetric mTSP.

Second, we incorporate variable time costs for each task so that each vertex v_i has an associated non-negative time cost t_i , with the exception of the depot and the start locations, where $t_i = 0$ for $v_i \in \{v_{depot}, v_{start1}, \dots, v_{startm}\}$. This is similar to the nodes' demand in Capacitated VRP (CVRP) [13] with the distinction that t_i is incorporated into the edge costs rather than being a separate vector. Our new cost matrix is therefore calculated as $T' = d_{ij} + t_j$, and again this modification introduces more asymmetry to this matrix, as time costs are no longer determined just by the distance between two vertices, but also by the direction of travel.

Third, each task is discretized into δ sub-tasks similar to the work in Split Delivery VRP [17] to enable cooperative tasking,

such that the total number of sub-tasks $n' = n \times \delta$. We refer to δ as the task discretization level. Sub-tasks replace tasks in the vertex set. They have the same spatial position as the original task but contain a corresponding fraction of the original task completion time. The time cost for a sub-task associated with task i is $t'_i = t_i / \delta$.

These combined changes from V, T, n, t to V', T', n', t' respectively define our problem as a variant of mTSP with m agents and $n' + m + 1$ vertices (including the home depot; see Fig. 1). Hence, the number of unique solutions (i.e. possible mission plans) is:

$$\frac{(n' + m - 1)!}{(m - 1)!} \quad (1)$$

B. The Graph Attention Replanner (GATR)

Our proposed GAT-based mission replanner (GATR) builds upon the encoder-decoder architecture from [13]. We adapt the mTSP environment from the RL4CO Library [27] to incorporate different agents' start locations and time costs associated with nodes (the first and second changes to standard mTSP described in Table I). Discretization of tasks to allow splitting of tasks between cooperative agents (the third change) does not require any additional adaptation, and can simply be accommodated during problem setup by replacing a single task of a certain time cost with several shorter collocated tasks. Finally, we enhance the data ingestion pipeline to enable training of a single generalized model for multiple problem sizes.

The overall training process is shown in Fig. 2a. Using the environment described above, we trained our model using the REINFORCE algorithm with a greedy rollout baseline [13] (see Appendix). The reward is defined as the negative of the maximum (largest) mission time among the agents as in [18].

1) *Encoder*: The primary role of the GAT encoder is to transform an input graph into latent representations, as illustrated in the first step of the flowchart in Fig. 2a. In addition to the standard mTSP graph features (depot locations

and task locations), we also include the time cost of vertices and agents' starting locations as part of the input to the encoder. This is similar to the approach used for specifying the 'demand' for deliveries in the CVRP [13], and for incorporating multiple depots in [15]. The GAT encoder then transforms these features into a set of node embeddings and a graph embedding that are passed to the decoder using the procedure described in the Appendix.

2) *Decoder*: Our decoder is based on the Attention Model, which was selected over the traditional Pointer Network due to its empirically superior performance [12]. The decoder forms the mission plan sequentially, as illustrated in Fig. 2c-d. At each time step, the decoder selects the next sub-task to visit for an agent based on the graph embedding, the current node embedding, the current state of the mission (in embedding space) and the availability mask. If the depot is selected, the algorithm proceeds to the subsequent agent. Note that an agent has the option to stop completing tasks and return to the depot at any point, provided there is still at least one other agent available to complete remaining tasks, similar to the assumption made in [16]. The state of the environment is summarized by the following features: the number of remaining agents, the accumulated mission time for the current agent, the maximum mission time across all agents, and the distance from the depot. These features are mapped into an embedding space using affine linear transformations and concatenated with the embedding for the current agent vertex (i.e. the agent's current location) to form the context embeddings. Unlike in [13], we do not include the embedding of the first vertex traversed by the agent, as the agent does not return to this location in the replanning problem. Cross-attention with an availability mask is applied to the context embeddings and vertex embeddings to compute the attention scores. In the replanning problem, the mask prevents selection of: previously visited sub-tasks; start locations when the current node is not the depot; and the depot if the current agent is the last one available, and unvisited sub-tasks remain. The next sub-task is then selected based on these scores, as detailed in the Appendix. This process is repeated until all sub-tasks are assigned to the agents. The final output is a set of routes, one for each agent, covering all sub-tasks, as illustrated in Fig. 2d. We note that in future work, the GATR architecture could be easily extended, e.g. to other encoders tailored for min-max objectives (e.g. [18], which has separate embeddings for task partitioning and navigation), or decoders that output actions for agents in parallel to encourage collaborative behavior [19].

III. RESULTS

We first demonstrate GATR's effectiveness for the CMRP on small problem sizes where it is possible to evaluate our replan against all possible solutions (Section III-A). Then we show how a single generalized GATR model is able to solve problems with varying numbers of tasks, agents, and task discretization levels, and how our model scales to larger problem sizes (Section III-B).

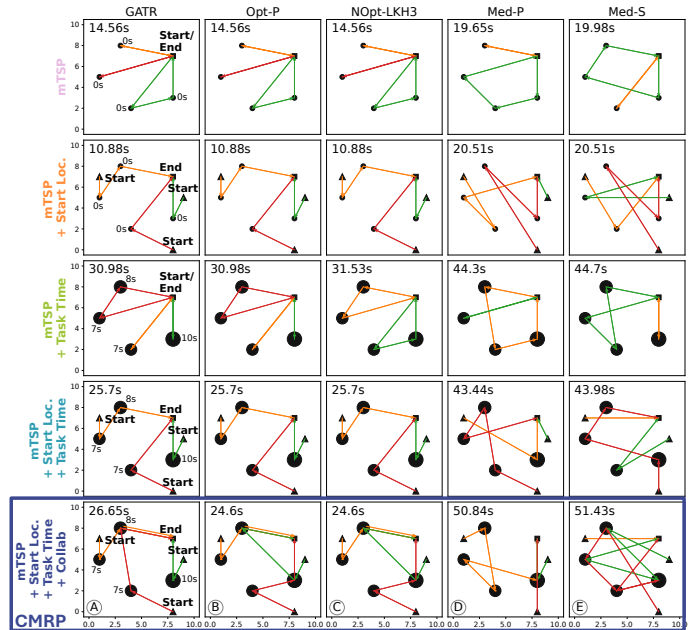


Fig. 3. Comparison of how GATR and other benchmarks (columns) solve an example three-agent-four-task problem for our five increasingly complex problem types (rows). The tours of the three agents are shown in red, amber, and green, the tasks are shown in black circles (with size corresponding to the time cost of each task for the last three rows), and the black square shows the home depot. In second, fourth and fifth rows, where agent start locations differ from the depot, they are marked with a black triangle.

A. Analysis for small problem sizes

Our experimental setup simulates a small indoor flight arena for aerial drones, with a 10 m × 10 m 2D domain. Task, depot and agent start locations are randomly distributed, and task time costs are randomly set between 1-10 s. We assume agents move at a fixed speed of 1 m/s. In order to systematically assess GATR's capabilities, we incrementally incorporate the changes to mTSP from Section II-A, defining five problem types: mTSP; mTSP with multiple agent start locations; mTSP with varying task time costs; mTSP with both multiple start locations and varying task time costs; and finally our CMRP problem which also features agent cooperation (tasks discretized into multiple sub-tasks). Initially, we focus on a three agent, four task problem, and train five models (one for each problem type), using three different seeds for each. We use task discretization level $\delta = 2$ for collaborative tasking, which means that our CMRP has 8 sub-tasks (4 tasks × 2) whilst all other problem types have 4 sub-tasks (only 1 sub-task per task). We generate 100 mTSP evaluation instances also across three seeds (300 instances). We use mission time (time to complete all tasks and return to homebase) as our primary evaluation metric.

We evaluate mission times against four benchmarks: the optimal and median population solutions (Opt-P and Med-P respectively), calculated via brute-force enumeration of the full population of possible solutions (e.g., 1,814,400 for 3 agents and 8 sub-tasks, from Equation 1, evaluation of which takes ~1 minute for a single scenario on the GPU); the median of a sample of the population - 101 randomly-generated solutions (Med-S); and a near-optimal solution generated by the LKH3 solver [31] (NOpt-LKH3). Opt-P and NOpt-LKH3 serve as




Fig. 4. (a) Histograms showing the distribution of mission times for our method GATR (top row) and four benchmarks for the test set of 300 scenarios for the three-agent-four-task CMRP. The circled letters represent the five different solutions for the single example scenario shown in the bottom row of Fig. 3. (b) A cumulative frequency histogram showing the percentage of GATR solutions (y-axis) with normalized mission times less than a given value (x-axis). The solid lines represent comparison against the Opt-P benchmark ($\hat{m}t_{Opt-P}$) for the five progressively complex problem types illustrated in Fig. 3, building up from mTSP (pink) to CMRP, our full problem (dark blue). The dashed dark blue line represents comparison against NOpt-LKH3 benchmark ($\hat{m}t_{NOpt-LKH3}$) for CMRP. (c) A graph illustrating the proportion of solutions with $\hat{m}t$ below the 0.01 (1%) threshold for different problem types and numbers of sub-tasks. The triangles represent the four-task problems and match the colors for the same set of experiments shown in (b).

the upper bounds on expected performance, whilst the two median benchmarks represent a lower bound as they simulate the expected performance of a randomly generated plan. We note that although it is feasible to obtain enumeration solutions for these small problem sizes, this is not possible for the larger problems we investigate in Section III-B, for which NOpt-LKH3 and Med-S are essential for evaluation instead.

Qualitative inspection of example solutions (e.g. single example scenario in Fig. 3) shows that GATR exhibits sensible and desired behaviors across all five problem types: effectively distributing tasks between agents, selecting tasks close to agent locations, accounting for variable task time cost when assigning agents, and collaboration on tasks.

This is supported by analysis of the full test set of 300 scenarios; GATR solutions in general have mission times similar to those from Opt-P and NOpt-LKH3 solutions (Fig. 4a) and much smaller than those from Med-P and Med-S solutions. In addition, the close alignment between Med-S and Med-P distributions in Fig. 4a provides confidence in approximating the median of the full population with a sample-based estimate for evaluation of larger problem sizes.

We also analyse normalized mission time ($\hat{m}t_{benchmark}$), where GATR’s mission times (mt_{GATR}) are compared against those from benchmark solutions:

$$\hat{m}t_{benchmark} = \frac{mt_{GATR} - mt_{benchmark}}{mt_{Med-P} - mt_{benchmark}} \quad (2)$$

where $mt_{benchmark}$ can refer to either mission time from the optimal benchmark (mt_{Opt-P}) or the near-optimal benchmark ($mt_{NOpt-LKH3}$) as desired. Here, $\hat{m}t_{benchmark} = 0$ indicates our model produces a solution with the same mission time as the chosen benchmark (the optimal solution or a near-optimal solution from the LKH3 solver); and $\hat{m}t_{benchmark} = 1$ indicates our model’s solution has the same mission time as an average randomly generated mission plan. As expected, optimality of the solution decreases as the problem complexity increases from mTSP (pink line in Fig. 4b) to the full CMRP (dark blue

line). However, Fig. 4b shows that 89.3% of GATR solutions still achieve $\hat{m}t_{Opt-P} < 0.1$ (normalized mission times within 10% of optimal), indicating consistently high performance despite the increased problem complexity over mTSP, and 12% solutions actually outperform our LKH3 heuristic benchmark (evidenced by the dashed line to the left of 0.00 in Fig. 4b). We also evaluate how performance scales with both problem complexity (Fig. 4b-c) and size (Fig. 4c). Of the two major modifications to mTSP (pink lines in Fig. 4b-c), the introduction of multiple start locations (orange lines) causes a larger negative impact on performance than incorporating variable task times (green lines). We suggest this is because the addition of new start location nodes causes a large increase in the number of possible solutions (as per equation 1), whereas varying task time only affects solution cost, not the number of solutions. In terms of scaling with problem size, we run additional experiments with varying number of sub-tasks, taking the test set from 300 to 4,500 runs per method. We see an expected decrease in performance as the number of sub-tasks increases (Fig. 4c), but we note this trend is not strong - indeed it is only clearly visible if we analyse the proportion of highest-tier solutions ($\hat{m}t_{Opt-P} < 0.01$ in Fig. 4c).

B. Generalization and scaling to larger problem sizes

A key implementation challenge for neural solvers is their inability to generalize to different problem setups (see penultimate column, Table I). But generalization is critical for replanning during a mission as the number of agents, tasks and other parameters will change throughout. This issue can be addressed by training a model on a diverse range of problem setups, but most existing implementations are constrained by fixed input graph sizes, including RL4CO that we adopt here, which supports varying numbers of agents but not tasks. To overcome this limitation, we calculate the maximum number of vertices required in the graph, i.e., $V_{max} = n_{max} \times \delta_{max} + m_{max} + 1$. Based on the scenario, we

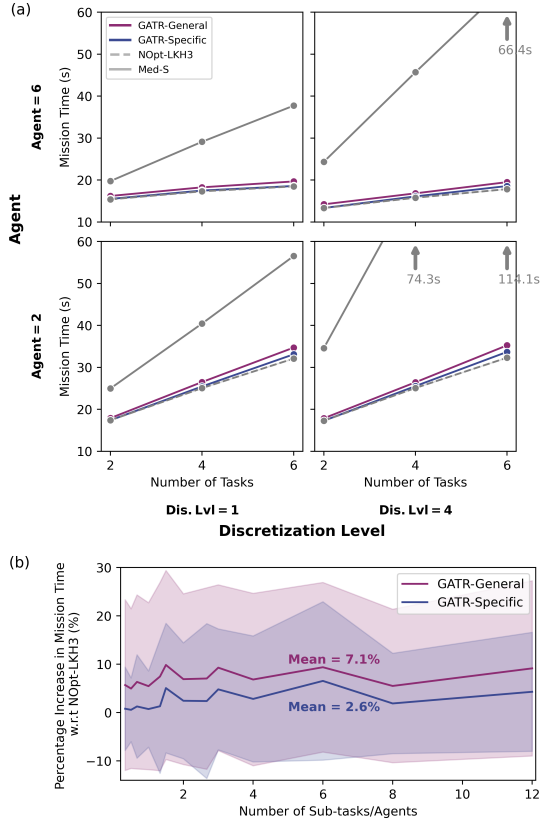


Fig. 5. (a) Comparison of average mission times (mean across 30,000 test scenarios) achieved by GATR-General (purple lines) and GATR-Specific (dark blue) across varying numbers of tasks, agents and task discretization levels; both closely track NOpt-LKH3 performance (dashed grey lines) and are much smaller than randomly generated missions (represented by Med-S, solid grey lines). (b) Comparison of percentage increase in mission time of GATR-General (purple) and GATR-Specific (dark blue) relative to LKH3 as a function of increasing problem size (number of sub-tasks per agent). Colored lines and shaded regions show mean and standard deviations respectively, with average across whole dataset indicated by colored text.

create a graph with $V = n \times \delta + m + 1$ vertices and pad it with $V_{max} - V$ vertices. These extra nodes are masked during the cross-attention procedures and will not be selected as the next action. The fixed-size matrices with custom paddings allow us to parallelize the data generation process while maintaining data diversity, hence enabling generalization. To demonstrate model generalization, we trained a single model (GATR-General) to solve problems with 1-6 agents, 1-6 tasks, and 1-4 task discretization levels. We tested GATR-General across 27 problem sizes (all combinations of 2/4/6 agents, 2/4/6 tasks and 1/2/4 discretization levels) and compared it against 27 specifically-trained models (GATR-Specific), one for each problem size. All models were also trained across three seeds. For fair comparison, we maintained the same network architecture to ensure similar level of RAM utilization. We set up test sets of 10,000 scenarios for each problem size with three seeds (30,000 total) as in [13], and evaluated GATR-General and GATR-Specific against NOpt-LKH3 and Med-S only; brute-force enumeration is not possible for large problems. Even assuming our best-case runtime for evaluating a single solution (3×10^{-5} s) from our simpler experiments (see previous Section), where it took ~ 1 minute to obtain an optimal mission

plan via enumeration, the largest problem size considered here (6 agents, 24 sub-tasks) would instead require 3×10^{-5} s $\times \frac{(24+6-1)!}{(6-1)!}$ possible solutions = 7×10^{16} years. We also compared inference runtimes between GATR-General and NOpt-LKH3 on both a high-performance GPU (single-node 40GB A100 GPU) and on edge hardware (Raspberry Pi 5 8GB), both using the same configurations (software dependencies detailed in the Appendix). We trained all models on the GPU, then transferred the resulting checkpoint files to the Pi for inference.

TABLE II
RUNTIME OF GATR-GENERAL AND NOPT-LKH3

Machine	Method	4 Agents				6 Agents			
		6 Sub-tasks		24 Sub-tasks		6 Sub-tasks		24 Sub-tasks	
		Abs.	Rel.	Abs.	Rel.	Abs.	Rel.	Abs.	Rel.
Pi	GATR-General	0.0078	(1.18%)	0.0070	(0.33%)	0.0072	(0.67%)	0.0074	(0.27%)
	NOpt-LKH3	0.7157		2.2074		1.1301		2.7467	
GPU	GATR-General	0.0624	(11.11%)	0.0633	(3.70%)	0.0726	(7.97%)	0.0709	(3.38%)
	NOpt-LKH3	0.6138		1.7741		0.9642		2.159	

Absolute (Abs.) times given in seconds, with relative (Rel.) times as % w.r.t. NOpt-LKH3. Standard Errors typically $< 3 \times 10^{-4}$ s for GATR-General and $3 - 7 \times 10^{-3}$ s for NOpt-LKH3.

Overall, both GATR-General and GATR-Specific closely follow LKH3's performance across a range of problem sizes (Fig. 5a). GATR-General and GATR-Specific achieve mission times that are only 7.1% and 2.6% larger than LKH3 respectively, and this gap does not show strong signs of widening with increasing workload per agent (Fig. 5b, purple and dark blue lines). Furthermore, GATR-General's runtime is just 0.27–1.18% of LKH3's runtime on edge hardware (e.g. 0.007 s vs. 2.7 s at lower end, see first two rows in Table II); this fast speed is essential for real-time on-board replanning, especially when many individual replans may be required across multiple agents, e.g. for consensus-based planning (see Discussion). Table II (compare first two rows) shows GATR-General maintains consistent runtimes across varying numbers of sub-tasks from 6 to 24 (similar to [15]), whereas LKH3's runtime scales poorly, increasing by roughly a factor of four. Moreover, execution on the Raspberry Pi is faster than on the GPU because there is limited room for parallelization when replans are only required in isolation.

IV. DISCUSSION

Our GATR framework enables effective and efficient replanning for cooperative multi-agent robotic systems, but also has the major advantage of enabling future extension to incorporate additional key capabilities; multi-objective, probabilistic consensus, and anticipatory replanning.

1) Managing competing aims - multi-objective planning:

In this study, our GATR method is designed to minimize overall mission time only, which implies that there are no practical constraints on maximum overall mission time or on fuel usage for each robotic agent, and that all tasks are expected to be completed. However, these assumptions are not valid for many real-world scenarios, where the ability to consider and explore trade-offs between mission completion time, task completion rate (percentage of tasks completed), and agent attrition rate (percentage of agents lost, e.g. due to insufficient fuel for return to homebase) is important. Existing heuristic or meta-heuristic methods are not designed to handle such multi-objective planning, but our encoder-decoder GATR

framework can easily support adaptation to incorporate these considerations in future work, primarily through modifying our model’s reward function to include terms for task completion rate and agent attrition rate, but also through modification of the model’s state space (e.g. adding elements to represent agent fuel levels, sensor functionality, actuator health), asymmetric edge weights and node costs (e.g. to represent wind or currents, which influence agent traversal and task completion times). These adaptations would also allow replanning for response to mission environment changes (e.g. reduction in mission time window due to incoming storm or imposition of a no-fly zone, change in travel times due to increased wind [6]) or more granular changes in system health (e.g. reduction of agent speed [35]), in addition to responding to the changes already demonstrated in this study; i.e. in tasking (changes in location, number, time-cost of tasks) and high-level system health (loss or addition of agents).

2) *Managing current uncertainty - probabilistic consensus planning:* To develop our GATR method in this study, we have assumed for simplicity that all robotic agents have the same perfect (i.e. zero uncertainty) knowledge of the environment state. Under these idealized conditions, all agents will generate the same replan for any given scenario; we therefore do not explicitly consider which agent generates any replan. But in real-world dynamic environments with limited communication, agents’ observations of the environment will be subject to uncertainty, and these observations and uncertainties will differ between agents [36]. Both of these challenges can be addressed through future incorporation of probabilistic Bayesian methods into GATR, which is straightforward for RL-based models [37] such as ours, but not for meta-heuristic or heuristic planning methods. The lightweight and fast nature of GATR (runtime order of 0.07 s) makes it feasible to use such a probabilistic approach for consensus replanning, by generating replans from multiple robotic agents with differing beliefs. This consensus replanning could be achieved via a proposer-responder framework, in which any agent within a subset can temporarily assume the role of leader [38], propose a plan, and allow other agents to vote based on their own beliefs about the global state, merging multiple beliefs to reduce input uncertainties and achieving replanning in a semi-decentralized fashion. The combination of diverse beliefs and independent planning fosters ensemble learning, where collective intelligence can be used to enhance the overall planning quality [39]. Our GATR framework can therefore serve as a fundamental building block for development of a probabilistic-guided, reliable and resilient semi-decentralized replanner.

3) *Managing future uncertainty - anticipatory planning:* Our GATR method provides the ability to replan onboard, which is essential for adapting to unforeseen mission changes [9]. But if we have knowledge of the type of changes and their likelihood, then anticipatory planning can enable our replans to mitigate the impact of further changes. The effectiveness of GATR could therefore be further enhanced by integrating anticipatory planning techniques. This could be achieved through further leveraging Bayesian methods to construct beliefs not just about neighbouring agents’ positions and task progress

[37], but also about the likelihood of future changes that might necessitate a replan, e.g. to system health (failure of agents, loss of fuel), tasking (addition of new tasks, changes in tasks), or to the environment (reduction of overall mission time). For instance, there is a significant likelihood of path deviations when wind speeds are high, or sensors may degrade more rapidly in hazardous environments, preventing agents from completing assigned tasks [35]. Replanning based on probabilistic beliefs of further changes, rather than assuming deterministic execution of the replan, would enable the system to regulate minor deviations before they escalate into agent failures or mission-wide breakdowns [40]. This proactive error regulation would reduce the frequency of replanning and the magnitude of required deviations from the initial plan or from any replans, leading to more efficient overall mission execution.

4) *Conclusion:* In this study we have defined the Cooperative Mission Replanning Problem (CMRP), which represents a key real-world challenge for developing resilient, autonomous multi-robot systems, and have developed GATR, a new model to solve this problem effectively and efficiently. This work enables on-board replanning for cooperative multi-robot systems for the first time, improving their adaptability and resilience in dynamic, uncertain environments.

APPENDIX

Graph Attention Model Overview: Following [41], we use query/key dimensionality $d_k = \frac{\text{Hidden Dimensions}}{\text{Attention Heads}} = \frac{128}{8} = 16$.

Encoder: The features of the input graph’s vertices (coordinates and time cost) are represented as $\mathbf{x}_i = [x_i, y_i, t_i]$. These are transformed into embeddings: $h_i^{(0)} = \text{Linear}(\mathbf{x}_i)$. The embeddings $h_i^{(0)}$ are processed through N attention layers to generate vertex embeddings $h_i^{(N)}$, and the graph embedding is computed as $\bar{h}^{(N)} = \frac{1}{n} \sum_{i=1}^n h_i^{(N)}$.

Decoder: The state embedding is defined as $\mathbf{s} = [m_{\text{remaining}}, t_{\text{cur}}, t_{\text{max}}, \text{dist}]$ and transformed into $h_{\text{state}} = \text{Linear}(\mathbf{s})$. The context embedding combines the graph embedding with the previous vertex embedding: $h_{(c)}^{(N)} = [\bar{h}^{(N)}, \text{Linear}([h_{\text{state}}, h_{\text{pre. vertex}}^{(N)}])]$. The probability of selecting a vertex is computed as: $p_i = \frac{e^{u_i}}{\sum_j e^{u_j}}$, where $u_j = -\infty$

if $j \in \text{Mask}$; otherwise $u_j = 10 \cdot \tanh\left(\frac{q' \cdot k'_j}{\sqrt{d_k}}\right)$, where queries q' and keys k' are products of learnable weights and embeddings.

Training: We train with a batch size of 5000 and a learning rate of 1×10^{-4} for up to 50 epochs with early stopping. Using 1,000,000 training and 100,000 validation samples, model training on an NVIDIA A100 GPU takes ~ 1 hour. Major software dependencies include: Python 3.10.16, RLCO 0.5.0.dev, TorchRL 0.6.0, PyTorch-Lightning 2.3.8.

REFERENCES

- [1] A. F. T. Winfield and J. Nembrini, "Safety in Numbers: Fault-Tolerance in Robot Swarms," *Int. J. Model. Identif. Control.*, vol. 1, pp. 30–37, 2006.
- [2] W. Zheng-Jie and L. Wei, "A Solution to Cooperative Area Coverage Surveillance for a Swarm of MAVs," *Int. J. Adv. Robotic Syst.*, vol. 10, no. 12, pp. 398–405, 2013.
- [3] D. Thakur, Y. Tao, R. Li, A. Zhou, A. Kushleyev, and V. Kumar, *Swarm of Inexpensive Heterogeneous Micro Aerial Vehicles*. Springer International Publishing, 2021.
- [4] J. Wu, X. Wang, Z. Shi, F. Zhao, Y. Ma, and Z. Jin, "A Large-Scale Mission Planning Method for Agile Earth Observation Satellite," in *Proc. 35th Chin. Control Decis. Conf*, 2023, pp. 5012–5017.
- [5] Y. Khosiawan and I. Nielsen, "A system of UAV application in indoor environment," *Prod. Manuf. Res.*, vol. 4, no. 1, pp. 2–22, 2016.
- [6] A. Román et al., "UAV as a Tool for Hazard Assessment: The 2021 eruption of Cumbre Vieja volcano," *Sci. Total Environ.*, 2022.
- [7] S. D. J. McArthur et al., "Multi-Agent Systems for Power Engineering Applications—Part I: Concepts, Approaches, and Technical Challenges," *IEEE Trans. Power Syst.*, vol. 22, no. 4, pp. 1743–1752, 2007.
- [8] Y. Rizk, M. Awad, and E. W. Tunstel, "Cooperative Heterogeneous Multi-Robot Systems: A Survey," *ACM Comput. Surv.*, vol. 52, no. 2, pp. 29:1–29:31, 2019.
- [9] T. H. Chung, M. R. Clement, M. A. Day, K. D. Jones, D. Davis, and M. Jones, "Live-Fly, Large-Scale Field Experimentation for Large Numbers of Fixed-Wing UAVs," in *2016 IEEE International Conference on Robotics and Automation (ICRA)*, 2016, pp. 1255–1262.
- [10] I. Bello, H. Pham, Q. V. Le, M. Norouzi, and S. Bengio, "Neural combinatorial optimization with reinforcement learning," in *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24–26, 2017, Workshop Track Proceedings*. OpenReview.net, 2017.
- [11] M. Deudon, P. Cournut, A. Lacoste, Y. Adulyasak, and L. Rousseau, "Learning Heuristics for the TSP by Policy Gradient," in *Integration of Constraint Programming, Artificial Intelligence, and Operations Research - 15th International Conference, CPAIOR 2018, Delft, The Netherlands, June 26–29, 2018, Proceedings*, ser. Lecture Notes in Computer Science, W. J. van Hoeve, Ed., vol. 10848. Springer, 2018, pp. 170–181.
- [12] A. Stohy et al., "Hybrid Pointer Networks for Traveling Salesman Problems Optimization," *PLOS ONE*, vol. 16, no. 12, pp. 1–17, 12 2021.
- [13] W. Kool, H. van Hoof, and M. Welling, "Attention, Learn to Solve Routing Problems!" in *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6–9, 2019*. OpenReview.net, 2019, pp. 7166–7190.
- [14] J. Zhou et al., "MVMoE: Multi-Task Vehicle Routing Solver with Mixture-of-Experts," in *Forty-first International Conference on Machine Learning, ICML 2024, Vienna, Austria, July 21–27, 2024*, no. 2557. JMLR.org, 2024, pp. 61 804–61 824.
- [15] W. Shi and C. Yu, "Multi-Agent Task Allocation with Multiple Depots Using Graph Attention Pointer Network," *Electronics*, vol. 12, no. 16, Jan. 2023.
- [16] S. Ma, J. Ruan, Y. Du, R. Bucknall, and Y. Liu, "An End-to-End Deep Reinforcement Learning Based Modular Task Allocation Framework for Autonomous Mobile Systems," *IEEE Trans. Autom. Sci. Eng.*, vol. 22, pp. 1519–1533, 2025.
- [17] M. Nazari, A. Oroojlooy, L. V. Snyder, and M. Takác, "Reinforcement Learning for Solving the Vehicle Routing Problem," in *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems, NeurIPS 2018, S. Bengio, H. M. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, Eds.*, 2018, pp. 9861–9871.
- [18] Z. Zheng, S. Yao, Z. Wang, T. Xialiang, M. Yuan, and K. Tang, "DPN: Decoupling partition and navigation for neural solvers of min-max vehicle routing problems," in *Proceedings of the 41st International Conference on Machine Learning*, ser. Proceedings of Machine Learning Research, R. Salakhutdinov et al., Eds., vol. 235. PMLR, 21–27 Jul 2024, pp. 61 559–61 592.
- [19] F. Berto et al., "PARCO: Parallel AutoRegressive Models for Multi-Agent Combinatorial Optimization," in *Advances in Neural Information Processing Systems*, 2025.
- [20] P. M. França, M. Gendreau, G. Laporte, and F. M. Müller, "The emphm-Traveling Salesman Problem with Minmax Objective," *Transp. Sci.*, vol. 29, no. 3, pp. 267–275, 1995.
- [21] A. Aggarwal, F. Ho, and S. Nakadai, "Extended Time Dependent Vehicle Routing Problem for Joint Task Allocation and Path Planning in Shared Space," in *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Oct. 2022, pp. 12 037–12 044, iSSN: 2153-0866.
- [22] S. Asadzadeh, W. J. de Oliveira, and C. R. de Souza Filho, "UAV-based remote sensing for the petroleum industry and environmental monitoring: State-of-the-art and perspectives," *J. Petroleum Sci. Eng.*, vol. 208, p. 109633, 2022.
- [23] M. Thammawichai, S. P. Baliyarasimhuni, E. C. Kerrigan, and J. B. Sousa, "Optimizing Communication and Computation for Multi-UAV Information Gathering Applications," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 54, no. 2, pp. 601–615, 2018.
- [24] E. Y. L. Kwan and J. Nunez-Yanez, "Entropy-Driven Adaptive Filtering for High-Accuracy and Resource-Efficient FPGA-Based Neural Network Systems," *Electronics*, vol. 9, no. 11, 2020.
- [25] H. Khorasani, H. Wang, H.-K. Tang, and C. Gupta, "K-nearest Multi-agent Deep Reinforcement Learning for Collaborative Tasks with a Variable Number of Agents," in *2021 IEEE International Conference on Big Data (Big Data)*, 2021, pp. 3883–3889.
- [26] T. Bektas, "The multiple traveling salesman problem: An overview of formulations and solution procedures," *Omega*, vol. 34, pp. 209–219, 06 2006.
- [27] F. Berto et al., "R14co: An extensive reinforcement learning for combinatorial optimization benchmark," in *Proceedings of the 31st ACM SIGKDD Conference on Knowledge Discovery and Data Mining V.2, ser. KDD '25*. New York, NY, USA: Association for Computing Machinery, 2025, p. 5278–5289.
- [28] I. I. Cplex, "V12. 1: Users Manual for CPLEX," *Int. Bus. Mach. Corp.*, vol. 46, no. 53, pp. 1–192, 2009.
- [29] N. Mazyavkina, S. Sviridov, S. Ivanov, and E. Burnaev, "Reinforcement learning for combinatorial optimization: A survey," *Comput. Oper. Res.*, vol. 134, p. 105400, 2021.
- [30] Y. Bengio, A. Lodi, and A. Prouvost, "Machine learning for combinatorial optimization: A methodological tour d'horizon," *Eur. J. Oper. Res.*, vol. 290, no. 2, pp. 405–421, 2021.
- [31] K. Helsgaun, "An Extension of the Lin-Kernighan-Helsgaun TSP Solver for Constrained Traveling Salesman and Vehicle Routing Problems: Technical Report," Dec. 2017.
- [32] J. Turner, Q. Meng, G. Schaefer, and A. Soltoggio, "Fast Consensus for Fully Distributed Multi-Agent Task Allocation," in *Proceedings of the 33rd Annual ACM Symposium on Applied Computing*, ser. SAC '18. New York, NY, USA: Association for Computing Machinery, Apr. 2018, pp. 832–839.
- [33] Q. Li, S.-Y. Liu, and X.-S. Yang, "Influence of initialization on the performance of metaheuristic optimizers," *Appl. Soft Comput.*, vol. 91, p. 106193, Jun. 2020.
- [34] Y. Dai, D. Xu, K. Zhang, Y. Lu, S. Maharjan, and Y. Zhang, "Deep Reinforcement Learning for Edge Computing and Resource Allocation in 5G Beyond," in *2019 IEEE 19th International Conference on Communication Technology (ICCT)*, 2019, pp. 866–870.
- [35] N. K. Ure, G. Chowdhary, J. P. How, M. A. Vavrina, and J. Vian, "Health Aware Planning under uncertainty for UAV missions with heterogeneous teams," in *2013 European Control Conference (ECC)*, 2013, pp. 3312–3319.
- [36] M. Bettini and A. Prorok, *Heterogeneous Teams*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2020, pp. 1–8.
- [37] X. Zhou, W. Wang, T. Wang, Y. Lei, and F. Zhong, "Bayesian Reinforcement Learning for Multi-Robot Decentralized Patrolling in Uncertain Environments," *IEEE Trans. Veh. Technol.*, vol. 68, no. 12, pp. 11 691–11 703, 2019.
- [38] D. Liu, L. Dou, R. Zhang, X. Zhang, and Q. Zong, "Multi-Agent Reinforcement Learning-Based Coordinated Dynamic Task Allocation for Heterogeneous UAVs," *IEEE Trans. Veh. Technol.*, vol. 72, no. 4, pp. 4372–4383, 2023.
- [39] J. T. Ebert, M. Gauci, F. Mallmann-Trenn, and R. Nagpal, "Bayes Bots: Collective Bayesian Decision-Making in Decentralized Robot Swarms," in *2020 IEEE International Conference on Robotics and Automation (ICRA)*, 2020, pp. 7186–7192.
- [40] S. Shao, H. Li, Y. Zhao, and X. Wu, "A New Method for Multi-UAV Cooperative Mission Planning Under Fault," *IEEE Access*, vol. 11, pp. 52 653–52 667, 2023.
- [41] A. Vaswani et al., "Attention is All you Need," in *Advances in Neural Information Processing Systems 30*, I. Guyon et al., Eds., 2017, pp. 5998–6008.