

Scalable Multi-Robot Informative Path Planning for Target Mapping via Deep Reinforcement Learning

Apoorva Vashisth , Manav Kulshrestha, Damon Conover, and Aniket Bera 

Abstract—Autonomous robots are widely utilized for mapping and exploration tasks due to their cost-effectiveness. Multi-robot systems offer scalability and efficiency, especially in terms of the number of robots deployed in more complex environments. These tasks belong to the set of Multi-Robot Informative Path Planning (MRIPP) problems. In this letter, we propose a deep reinforcement learning approach for the MRIPP problem. We aim to maximize the number of discovered stationary targets in an unknown 3D environment while operating under resource constraints (such as path length). Here, each robot aims to maximize discovered targets, avoid unknown static obstacles, and prevent inter-robot collisions while operating under communication and resource constraints. We utilize the centralized training and decentralized execution paradigm to train a single policy neural network. A key aspect of our approach is our coordination graph that prioritizes visiting regions not yet explored by other robots. Our learned policy can be copied onto any number of robots for deployment in more complex environments not seen during training. Our approach outperforms state-of-the-art approaches by at least 26.2% in terms of the number of discovered targets while requiring a planning time of less than 2 sec per step. We present results for more complex environments with up to 64 robots and compare success rates against baseline planners.

Index Terms—Motion and path planning, reinforcement learning, multi-robot systems.

I. INTRODUCTION

AUTONOMOUS robotic systems are used in several tasks, such as search and rescue missions [1], environment mapping [2], and orchard monitoring [3]. Multi-robot systems are gaining popularity in these domains due to their increased efficiency, compared to single-robot systems [4] and manual approaches [5]. Key challenges for deploying multi-robot systems in these tasks include planning efficient paths for all robots to optimize the task objective, avoiding inter-robot and robot-obstacle

collisions, scaling to larger multi-robot systems deployed in more complex environments, and considering communication and resource constraints.

In this work, we aim to develop a deep reinforcement learning-based, scalable, multi-robot path planning approach for discovering stationary targets in an unknown 3D environment. Here, each robot is constrained to a resource budget (e.g., battery capacity or mission time). Our considered problem setting belongs to the family of multi-robot informative path planning (MRIPP) problems. Our 3D environment contains unknown static obstacles. Our multi-robot system consists of autonomous aerial vehicles (AAVs), where each AAV is equipped with two range-constrained modules - a unidirectional RGB-D sensor and a communication module. The challenges considered in this work include - the ability to scale to a larger number of robots deployed in more complex environments, consideration of regions explored by other robots while planning, and avoiding inter-robot and robot-obstacle collisions as the robots operate under communication constraints. Applications of our work include search and rescue missions, reconnaissance for military applications, mapping fruits in an orchard for precision agriculture, and discovering targets of interest in urban environments.

Several approaches have been proposed for the MRIPP problem. Classical approaches [6], [7], [8], [9], [10] extend the single-robot planners for multi-robot systems via sequential allocation by planning path for each robot one after another, in a specific sequence. Centralized approaches [2], [11], [12], [13], [14], [15], [16], [17], [18], [19] plan over the joint action space of all robots. However, these planners assume availability of global communication and hence are not applicable in our problem setting with a limited communication range. Moreover, as centralized planners plan in the joint action space of all robots, they do not scale well with increasing number of robots. Recently proposed decentralized planners [20], [21], [22], [23], [24], [25], [26], [27] decouple the action space - each robot in the multi-robot system plans its own action. However, these approaches do not consider limited communication range [20], [22], [23], [25], [26], inter-robot collision avoidance [22], or presence of unknown static obstacles in the environment [25], [27]. To address these issues, we propose a novel decentralized approach based on deep reinforcement learning. Our approach considers the regions previously explored by other robots during planning, and constrains the planning to each robot's local region. This aids our method in not only avoiding collisions with newly discovered static obstacles, but also in preventing inter-robot collisions. As our approach is decentralized, it can

Received 2 June 2025; accepted 16 December 2025. Date of publication 19 January 2026; date of current version 3 February 2026. This work was supported by the DEVCOM Army Research Laboratory under cooperative agreement W911NF2520170. This article was recommended for publication by Associate Editor A. M. Metelli and Editor J. Kober upon evaluation of the reviewers' comments. (Corresponding author: Apoorva Vashisth.)

Apoorva Vashisth, Manav Kulshrestha, and Aniket Bera are with the Department of Computer Science, Purdue University, West Lafayette, IN 47907 USA (e-mail: vashista@purdue.edu; mkulshre@purdue.edu; aniketbera@purdue.edu).

Damon Conover is with DEVCOM Army Research Laboratory, Adelphi, MD 20783 USA (e-mail: damon.m.conover.civ@army.mil).

Our code and trained model are available at - https://github.com/AccGen99/mafl_ipp.

This article has supplementary downloadable material available at <https://doi.org/10.1109/LRA.2026.3655204>, provided by the authors.

Digital Object Identifier 10.1109/LRA.2026.3655204

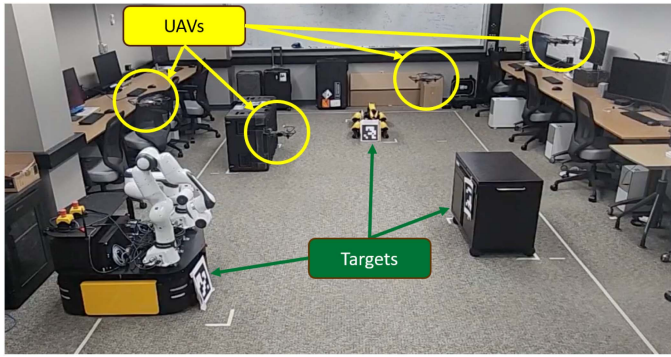


Fig. 1. We implement our approach on Ryze Tello drones in a real-world monitoring scenario. We use Aruco tags as the targets to be discovered. Our approach successfully plans collision-free paths online for maximizing the number of discovered targets while under mission-time constraints. Here, 4 Tello drones and 3 Aruco tags are visible.

be implemented on larger multi-robot systems deployed within more complex environments unseen during training.

To summarize, we develop a scalable, decentralized, and efficient deep-reinforcement learning-based solution to the MRIPP problem. Our deep reinforcement learning based policy can be implemented on a multi-robot system consisting of a large number of robots in environments not seen during training. The core aspect of our approach is the coordination graph that models the regions of the environment previously explored by other robots. Fig. 1 illustrates our approach implemented on a multi-robot system consisting of UAVs in a real-world monitoring scenario. We present the following four contributions:

- Our coordination graph models the regions explored by other robots, enabling the policy to plan actions visiting unexplored regions of the environment.
- Our proposed reward function, aligned with the MRIPP objective, encourages inter-robot communication.
- Our method enables more efficient discovery of targets (66%) compared to state-of-the-art methods (52%) when deployed in previously unseen environments.
- Our learned policy can be deployed over higher number of robots without requiring re-training.

We assess the effectiveness of our approach in an urban monitoring application in a simulator and also perform real-world experiments with Ryze Tello drones.

II. RELATED WORK

Classical approaches [6], [7], [8], [9], [10] to the MRIPP problem attempt to extend the single-robot methods to multi-robot planning via sequential allocation [6]. Here, one planner plans the path for each robot one after another in an arbitrary order. These approaches decompose the environment into clusters and plan each robot's path over the clusters [7], use a random sequence order of the robots for sequential allocation [10], or utilize a deep reinforcement learning approach for generating the planning order of the robots [9]. However, these approaches assume infinite communication range and do not consider inter-robot collisions, leading to inapplicability in our problem setting.

Centralized methods [2], [11], [12], [13], [14], [15], [16], [17], [18], [19] introduce cooperative behavior by planning in the joint action-space of all robots. Some approaches plan paths that minimize the final uncertainty of a given uncertainty map [11], [12], [13]. Other methods decompose the environment into disjoint clusters and allocate one robot per cluster [2], [14], [15], [18]. Other approaches rely on a consensus filter [17], on deep reinforcement learning [16], or attempt to resolve the inter-robot collisions in pre-computed paths [19]. However, these approaches require availability of infinite communication range. Moreover, as they plan in the joint action space of the robots, these methods are not scalable to a large number of robots. Decentralized planners [20], [21], [22], [23], [24], [25], [26], [27], [28], [29] provide scalable solutions by allowing each robot to independently plan their next action. These approaches propose a decentralized variant of Monte-Carlo tree search (MCTS) [24], or utilize consensus filters for encouraging cooperation among the robots [26], [27], [28]. Recently, deep reinforcement learning based decentralized planners have been developed that are not only computationally efficient at deployment but also have the capability of generalizing to similar environments not seen during training. These approaches utilize parameter sharing to encourage cooperation among robots [20], employ imitation learning [29], utilize the centralized training and decentralized execution paradigm [21], [25], employ Q-Learning to learn collision avoidance behavior [23], or attempt to utilize attention mechanism for modeling the paths of other robots [22]. However, these approaches operate in 2D environment with known obstacles [29], do not consider inter-robot collisions [22], [23], assume an obstacle-free environment [25], [27], or consider static communication connectivity [28]. A key difference of our approach with the prior works is that each robot models the regions explored by other robots within communication range and plans only within its local region. This encourages visiting unexplored regions, prevents collisions with discovered obstacles by planning in known local regions, and avoids inter-robot collisions due to planning outside of the collision range of other robots. Our results demonstrate that our approach outperforms the state-of-the-art learning and non-learning methods in our considered problem setting and is scalable to larger number of robots deployed in more complex environments not seen during training.

III. BACKGROUND

A. Problem Setting

In this work, we aim to maximize the number of discovered targets in a 3D environment containing undiscovered static obstacles. Our multi-robot system consists of $N \in \mathbb{Z}^+$ robots constrained to a total resource budget $B \in \mathbb{R}^+$. We model the budget as the sum of the maximum cost of the paths executed by each robot. Hence, each robot $i \in 1, \dots, N$ receives a budget of $B^i = B/N$. We model the robots as UAVs equipped with a unidirectional range sensor (e.g., RGB-D camera) and a communication module which is single-hop and range-constrained, i.e., robots communicate only when they are within distance $\rho \in \mathbb{R}^+$. To discover the targets, we assume presence of a noiseless

communication Gaussian process over the set of candidate actions aid our policy network in considering the regions explored by other robots while planning the next action.

B. MRIPP Objective

We model the path followed by robot i as a sequence of consecutively executed actions $\psi^i_{0:T} = (\mathbf{a}_0^i, \mathbf{a}_1^i, \dots, \mathbf{a}_T^i)$ where \mathbf{a}_0^i is the initial pose and \mathbf{a}_T^i is the action executed upon depletion of the budget $B^i = B/N$, causing the termination of its mission. In general, the MRIPP problem searches the space of all possible paths $\Psi^{1:N}$ for a set of optimal paths $\psi^*_{0:T} \in \Psi^{1:N}$ such that $\psi^*_{0:T} = [\psi^1_{0:T}, \psi^2_{0:T}, \dots, \psi^N_{0:T}]$ to maximize an information-theoretic objective function:

$$\psi^*_{0:T} = \operatorname{argmax} I(\psi_{0:T}), \text{ s.t. } c(\psi^i_{0:T}) \leq B^i \forall i \in [1, N], \quad (1)$$

where $I : \Psi \rightarrow \mathbb{R}^+$ is the information gained upon executing the trajectory $\psi^i_{0:T}$ and $c : \Psi^i_{0:T} \rightarrow \mathbb{R}^+$ maps the path $\psi^i_{0:T}$ to its execution cost.

While traversing the path $\psi^i_{0:t}$, the robot transitions between two consecutively executed actions over a straight line. Observations are collected at each waypoint in the path and are used to update the utility Gaussian process (u_{util}, P_{util}) and the occupancy map. Upon communication with a nearby robot, the communication Gaussian process (u_{comm}, P_{comm}) is updated with the waypoints visited by the communicating robots. Hence, due to the successive nature of the executed actions in the planned path, we model the MRIPP problem as a sequential decision-making process. Towards the MRIPP objective, we define a function $\zeta : \mathcal{A} \times \Psi^{1:N} \rightarrow \mathbb{R}^+$ as the number of new targets observed upon executing an action \mathbf{a}_t^i by robot i after following the path $\psi^i_{0:t-1}$.

We define the information obtained by robot i as:

$$I(\psi_{0:T}) = \sum_{t=1}^T \zeta(\mathbf{a}_t^i, \Psi^{1:N}_{0:t-1}), \quad (2)$$

and we aim to plan $\psi^*_{0:T}$ to maximize the information gain. The above formulation models the diminishing information gain by not considering the targets that have been observed during an earlier exploration of the region.

C. Reward Structure

In order to maximize the number of discovered targets, each robot needs to balance exploration of environment (always choosing the action with most uncertainty) with exploitation (always choosing action with maximum informativeness) of the obtained observations. Moreover, inter-robot communication is necessary to keep track of the regions previously explored by other robots. This aids in avoiding planning of sub-optimal actions leading to re-exploration. Inspired by previous works [3], [32], we propose a new reward structure that not only considers the exploration-exploitation trade-off but also encourages inter-robot communication. Upon communication with other robots, information is exchanged about the regions previously explored by the robots. At each time-step t , the robot i has executed the action \mathbf{a}_{t-1}^i , collected observations, communicated with

the nearby robots, and receives a reward $r_t^i \in \mathbb{R}^+$. The reward function consists of an exploratory term $r_{e,t}^i$, an informative term $r_{u,t}^i$, and a communication term $r_{c,t}^i$ so that:

$$r_t^i = \alpha r_{e,t}^i + \beta r_{u,t}^i + \gamma r_{c,t}^i \quad (3)$$

where:

$$\begin{aligned} r_{e,t}^i &= \frac{\operatorname{Tr}(P_{util}^-) - \operatorname{Tr}(P_{util}^+)}{\operatorname{Tr}(P_{util}^-)}, \\ r_{c,t}^i &= \frac{\operatorname{Tr}(P_{comm}^-) - \operatorname{Tr}(P_{comm}^+)}{\operatorname{Tr}(P_{comm}^-)}, \\ r_{u,t}^i &= \zeta(\mathbf{a}_{t-1}, \psi_{0:t-2}) \end{aligned} \quad (4)$$

where the constants α and β balance the exploration-exploitation trade-off and γ rewards inter-robot communication. $\operatorname{Tr}(\cdot)$ is the matrix trace operator. Here, P^- and P^+ indicate the prior and posterior covariance matrices of the Gaussian processes. The reduction in variance of the utility Gaussian process estimates the exploration of the environment due to the robot's own executed actions. Similarly, the variance reduction of the communication Gaussian process estimates the exploration knowledge gained from other robots. The number of new targets observed measures the information gained upon execution of action \mathbf{a}_t^i by robot i . Scaling the reward by $\operatorname{Tr}(P^-)$ stabilizes the actor-critic network training [3], [31].

Our reward generation method ensures that each robot receives the reward reflecting the contribution of its actions towards the global MRIPP objective -

- Robot i will not receive informative reward $r_{u,t}^i$ for the new targets that have been observed by another robot.
- At each timestep t during training, we utilize a single global utility Gaussian process for all robots. The term $r_{e,t}^i$ is calculated for robot i 's action considering no action has been executed by any other robots.
- The communication reward $r_{c,t}^i$ depends on the exploration knowledge gained by robot i from other robots. Each robot has a separate instance of communication Gaussian process to generate this reward component to ensure the reward received depends on the communication performed due to execution of its own action only.

D. Coordination Graph

The MRIPP problem considered in this work requires each robot i in our multi-robot system to reason about the distribution of targets in the environment to optimize the MRIPP objective as described in (1) and the regions explored by other robots for inter-robot collision avoidance. Inspired by the dynamic graph approach for single-robot path planning [3], we propose a novel coordination graph that enables our approach to model the distribution of targets in the robot's local neighborhood, plan actions to visit regions in the environment not explored by other robots, and avoid inter-robot collisions and collisions with newly discovered static obstacles. Our policy neural network relies on the coordination graph to predict the next action to execute.

Each robot rebuilds its coordination graph at every timestep to account for the newly obtained observations. Our coordination

graph for a robot i at timestep t is a fully-connected graph $\mathcal{G}_t^i = (\mathcal{N}_t^i, \mathcal{E}_t^i)$. The node set \mathcal{N}_t^i defines the set of collision-free candidate actions. The edge set \mathcal{E}_t^i defines the collision-free paths from the robot's current pose to each candidate action and each edge in the set is associated with the cost of executing the given candidate action. We construct the feature matrix \mathcal{M}_t^i for robot i corresponding to its coordination graph as the input to our policy neural network. The features consist of the candidate actions, the utility and uncertainty values of candidate actions regressed from the Gaussian process modeling the utility, and the probability and uncertainty values queried from the communication Gaussian process. The n^{th} row of \mathcal{M}_t^i relates to the n^{th} candidate action of robot i at timestep t :

$$\mathbf{M}_t^i(n) = [\mathbf{a}_{n,t}^i, u_{\text{utility}}(\mathbf{a}_{n,t}^i), P_{\text{utility}}(\mathbf{a}_{n,t}^i, \mathbf{a}_{n,t}^i), \\ u_{\text{comm}}(\mathbf{a}_{n,t}^i, t), P_{\text{comm}}(\mathbf{a}_{n,t}^i, \mathbf{a}_{n,t}^i)],$$

where $\mathbf{a}_{n,t}^i = [x_{n,t}^i, y_{n,t}^i, z_{n,t}^i, d_{n,t}^i]^\top$, $u_{\text{utility}}(\mathbf{a}_{n,t}^i)$ and $P_{\text{utility}}(\mathbf{a}_{n,t}^i, \mathbf{a}_{n,t}^i)$ are the regressed utility and uncertainty values for candidate action $\mathbf{a}_{n,t}^i$, and $u_{\text{comm}}(\mathbf{a}_{n,t}^i)$ and $P_{\text{comm}}([\mathbf{a}_{n,t}^i], [\mathbf{a}_{n,t}^i])$ are the regressed probability and confidence values from the communication Gaussian process modeling the locations of other robots at timestep t .

E. Policy Neural Network

Our coordination graph models each robot's collision free action space and aids the deep reinforcement learning policy in reasoning about the robot's current knowledge of the environment. As the utility Gaussian process can only model the greedy action selection through utility regression, deep reinforcement learning is essential for achieving the balance in short-term exploitation of obtained information and the long-term exploration of the unknown environment.

At each timestep t , each robot i in our multi-robot system utilizes an attention-based neural network [31] to model the planning policy $\pi(\mathcal{G}_t^i, \psi_{0:t-1}^i, \tilde{B}^i)$. The planning policy outputs the probability distribution over the candidate actions in the robot's candidate action set \mathcal{A}_t^i . The policy relies on the feature matrix \mathbf{M}_t^i of the current coordination graph \mathcal{G}_t^i , path executed so far $\psi_{0:t-1}^i$ and the remaining budget \tilde{B}^i . The network structure consists of an encoder and a decoder module. The encoder models the information distribution obtained from observations and the environment explored so far by learning the dependencies among the candidate actions in \mathcal{G}_t^i , forming the context over collected observations. The decoder utilizes the learned context features from the encoder, the planning state, and the budget mask to output the probability distribution over the set of candidate actions \mathcal{A}_t^i . The planning state consists of the path executed by the robot so far $\psi_{0:t-1}^i$ and the remaining budget \tilde{B}^i . The budget mask aids in filtering out candidate actions leading to the violation of the budget constraint. Additionally, the decoder module estimates the value function of the current state. The estimated value, executed actions, coordination graphs, planning states, and rewards generated by actions of the robots throughout the training episode are collected in the experience buffer for

on-policy actor-critic reinforcement learning under centralized training and decentralized execution paradigm. In this work, we use proximal policy optimization [33] due to its stability and sample efficiency. During deployment, at each time step and for each robot, we execute the most informative action.

V. RESULTS

A. Setup

Environment: We test our approach in an urban monitoring scenario consisting of buildings and windows. We represent the environment internally as bounded within a scale-agnostic unit cube. Each robot in our multi-robot system maintains an occupancy map for collision avoidance with static obstacles. We initialize the occupancy map as unknown space and update the free space or occupied space based on obtained observations. Our training environment consists of regularly spaced buildings with the windows generated randomly on the buildings. However, our test environments consist of buildings generated at random locations.

Hyperparameters: We tune the hyperparameters of our Gaussian processes in a small representative environment. We use the Matérn 1/2 kernel function for the Gaussian processes. For the reward structure defined in (1), we choose $\alpha = 20.0$ and $\delta = 0.02$ so that both the exploratory $r_{e,t}^i$ and utility reward $r_{u,t}^i$ terms lie numerically in the range $[0,1]$. In order to promote MRIPP objective over the inter-robot communication, we use $\gamma = 1.0$ to keep the numerical value of $r_{c,t}^i$ lower than the other terms.

Robot Configuration: We consider each robot as a UAV platform equipped with an RGB-D camera having 90° field of view. We model the reduction in the confidence of target identification with increasing distance by limiting the camera sensing range to 24% of the environment size. The UAVs can communicate at a maximum communication distance $\rho = 0.3$. The sensor view-point set \mathcal{D} is discretized as $\{0, \frac{\pi}{2}, \pi, \frac{3\pi}{2}\}$ radians. However, our approach supports extension to finer discretizations by extending the set \mathcal{D} .

Training: We generate multiple training episodes to populate our experience buffer. Each training episode consists of a multi-UAV system with a total budget B . Our policy is trained in a structured environment and then transferred to a randomized environment for testing. While we fix the number of buildings during training, the number of windows is varied in $[200, 250]$. The start action for each robot is $\mathbf{a}_0 = (0.0, 0.0, 0.0, \frac{\pi}{2})$. We set $L = 80$ for each robot's coordination graph. Since we normalize the internal environment representation, our budget value B is unitless. For each training episode, B is a randomly generated real value in the range $[7.0, 9.0]$. Each episode is constrained to a maximum of 256 timesteps. To speed-up the training process, we run 36 parallel environment instances and train our policy network over 8 epochs with a batch size of 1024. We utilize Adam optimizer with a learning rate of 10^{-4} , decaying by a factor of 0.96 after every 512 optimization steps. The policy gradient epsilon-clip parameter is set to 0.2. We train our policy network on a computing cluster equipped with Intel(R) Xeon(R)

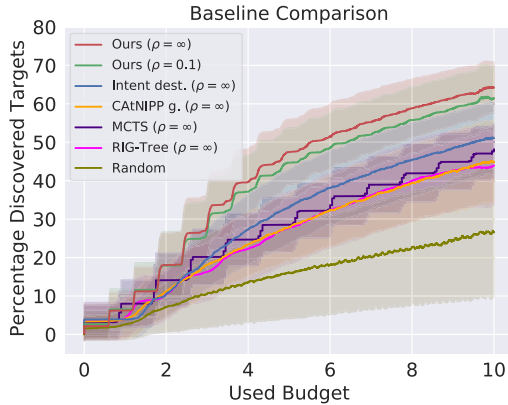


Fig. 3. Comparison of our approach with other baselines in an urban environment. Our performance metric is the percentage of targets discovered during the episode. The solid lines represent the mean values across 250 trials, while the shaded areas denote the standard deviations.

CPUs @ 3.60GHz and one NVIDIA A30 Tensor Core GPU. We require $\sim 120,000$ environment interactions for our policy to converge. All our tests are conducted on the same compute cluster.

B. Baseline Comparison

In this section, we compare the performance of our approach against state-of-the-art learning and non-learning baselines. We utilize fixed random seeds to generate 25 test environments. We enable global communications for the baselines and present the performance of our approach with both global communication ($\rho = \infty$) and restricted communication ($\rho = 0.1$). We deploy $N = 3$ UAVs with a total budget of $B = 10.00$ and run 10 trials corresponding to each randomly generated environment, leading to a total of 250 tests. Our baselines include: (i) Intent with destination modeling [22] as a zero-shot greedy policy (Intent dest.), (ii) CAtnIPP [31] with a zero-shot policy (CAtnIPP g.), (iii) non-learning Monte Carlo Tree Search [34] (MCTS), (iv) non-learning Rapidly exploring random Information Gathering Trees [35] (RIG-Tree), and (v) a random policy (random agent). As MCTS, and RIG-Tree are single robot planners, we extend them to multi-robot planning via sequential allocation [6]. We tune all the baselines for best performance in our problem setting. We consider the metric of percentage of targets discovered, as well as provide the average planning time per step. Additionally, as modification of these approaches to account for collision-free path planning is non-trivial, we allow the UAVs for each planner to ignore obstacle-UAV and inter-robot collisions, ensuring a fair comparison within this set of experiments.

Our results shown in Fig. 3 and Table I indicate that our method outperforms the considered baselines in terms of the number of discovered targets. Note that introduction of communication constraint leads to a drop in the performance of our approach, however we still outperform the considered baselines operating with global communications. This could be attributed to our coordination graph explicitly representing the regions

TABLE I
RESULTS OF OUR APPROACH COMPARED WITH OTHER LEARNING AND NON-LEARNING BASELINES IN AN URBAN MONITORING SCENARIO

Baseline	% targets	Time (s)
Our approach ($\rho = \infty$)	66.00 ± 5.59	1.58
Our approach ($\rho = 0.1$)	62.16 ± 8.25	1.58
Intent dest. ($\rho = \infty$)	52.31 ± 9.29	25.39
CAtnIPP g. ($\rho = \infty$)	45.31 ± 10.59	3.84
MCTS ($\rho = \infty$)	48.50 ± 7.81	153.49
RIG-Tree ($\rho = \infty$)	47.38 ± 10.43	60.46
Random agent	31.83 ± 15.23	0.06

TABLE II
ABLATION STUDY FOR MODELING OF EXPLORED REGIONS

Approach	% targets
With modeling explored regions	66.00 ± 5.59
Without modeling explored regions	61.97 ± 7.34

TABLE III
ABLATION STUDY FOR REWARD STRUCTURE

Approach	% targets
With communication reward	66.00 ± 5.59
Without communication reward	62.92 ± 7.01

explored by other robots, leading to planning actions visiting unexplored regions that provide performance improvement in our approach over the baselines that do not model the regions visited by other robots. Note that deep reinforcement learning based methods are significantly more time-efficient than non-learning approaches, justifying their use over non-learning methods for real-time applications.

C. Ablation Studies

We study the impact of our communication Gaussian process and our new reward structure on the performance of our approach via the metric of percentage discovered targets.

Communication Gaussian Process: To evaluate the impact of modeling the unexplored regions of the environment on the performance of our approach, we compare our approach trained with and without the communication Gaussian process. Our results in Table II show that the performance improves when the communication Gaussian process is included. Furthermore, an unpaired t-test conducted between the two groups ($n = 250$ each) yielded a p-value of 3.91×10^{-9} , indicating a statistically significant difference well below the conventional threshold of 0.05. Hence, we conclude that our policy neural network learns to reason about the unexplored regions during planning.

Communication Reward: To evaluate the impact of the communication reward term $r_{c,t}^i$ in (4) on the performance of our approach, we compare the performance of our policies trained with $\gamma = 0.0$ and $\gamma = 1.0$. Our results in Table III show that the performance improves upon inclusion of the communication reward. Again, the difference between the two groups of 250 tests each is statistically significant, with a p-value of 5.11×10^{-6} , significantly less than 0.05, providing strong evidence against

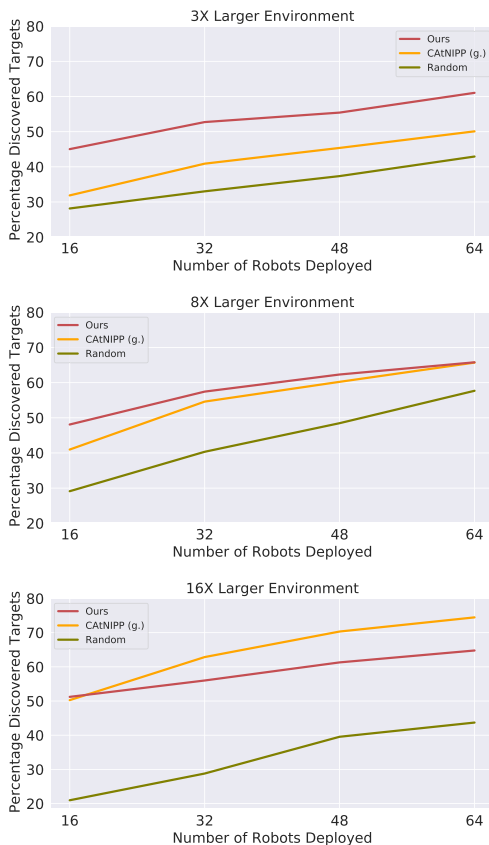


Fig. 4. Our approach outperforms other baselines in terms of percentage discovered targets in a $3\times$, $8\times$, and $16\times$ larger environments. The x-axis indicates the number of robots in the multi-robot system during the test. The solid lines represent the mean values across 100 trials.

the null hypothesis. Hence, our new reward structure promotes inter-robot communication and leads to improved performance.

D. Scalability

We compare the ability of our approach to scale to larger environments and more number of robots N in the multi-robot system with other approaches. Our policy learned in the small training environment with $N = 3$ robots is evaluated in larger environments and varying number of robots not seen during training. We present results for test environments that are approximately $3\times$, $8\times$, and $16\times$ larger than the training environment. We consider $N \in \{16, 32, 48, 64\}$ and conduct 100 tests in each environment for every N .

We compare the performance of our approach, deployed with communication distance $\rho \in \{0.2, 0.15, 0.12\}$ in $3\times$, $8\times$, $16\times$ environment respectively, with (i) CATNIPP [31] with global observability and (ii) a random policy. We train CATNIPP in the training environment described in Section V-A. We do not consider the intent baseline due to its compute intensive nature for large multi-robot systems. We do not evaluate the non-learning methods, as the sequential allocation based time-intensive nature leads to infeasible computation times for large multi-robot systems.

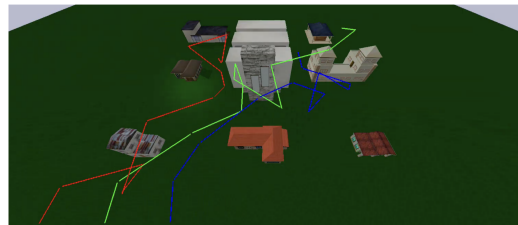


Fig. 5. Our approach implemented in an urban simulation environment. We place 3 UAVs and 161 targets in the environment and trace each UAV's path with colored tracelines.

TABLE IV
COMPARISON OF OUR DEEP REINFORCEMENT LEARNING-BASED APPROACH AGAINST BASELINES IN AN URBAN ENVIRONMENT

Approach	% targets
Our Approach	66.46 ± 11.08
Random Planner	28.94 ± 8.70

Fig. 4 demonstrates the results for this set of experiments. Our approach strongly outperforms the considered baselines in $3\times$ environment, slightly outperforms in $8\times$ environment, and is outperformed in $16\times$ environment. The success of our planner in $3\times$ and $8\times$ environments can be attributed to the modeling of unexplored regions, leading to planning of more informative paths by robots in our multi-robot system as compared to other approaches. However, in the $16\times$ environment, the instances of communication with robots located further away drastically reduces, causing our approach to be outperformed by the CATNIPP planner with global observability. Future work will explore more complex communication paradigms to mitigate this issue.

E. Simulation

We demonstrate the applicability of our deep reinforcement learning approach in an urban monitoring scenario. We use the gym-PyBullet-drones [36] simulator to accurately model UAV physics. Our simulation environment is built using the Houses3K dataset [37] and is bounded by a $60\text{ m} \times 60\text{ m} \times 30\text{ m}$ cuboid as shown in Fig. 5. We assume perfect localization and use ground truth target discovery. The 3 UAVs move at a maximum speed of 1 m/s.

We compare the performance of our approach with a random planner that reflects the performance lower bound. We do not implement other baselines considered in Section V-B as modifying these approaches for avoiding inter-robot collisions and consideration of the presence of unknown obstacles in the environment is a non-trivial task. Here, our evaluation metric is the percentage of windows discovered by the robots. To ensure every discovered target is counted only once, we record the coordinates of discovered targets. Our results are reported for missions with a budget of 7.0 units in Table IV. Our approach outperforms the random planner.

F. Implementation

We demonstrate the real-world applicability of our method on a multi-robot system for target discovery as illustrated in Fig. 1.

We carried out experiments on 4 Ryze Tello drones in a $7.62 \times 3.25 \times 2.4$ m³ arena containing randomly placed obstacles and 6 Aruco tags as targets.

VI. CONCLUSION

We present a novel deep reinforcement learning approach for the MRIPP problem in an unknown 3D environment. Our coordination graph-based approach models the unexplored regions of the environment for efficient target discovery. We present experimental results to support that: (i) our coordination graph encourages exploration of unknown regions of the environment, (ii) our reward function encourages inter-robot communication, (iii) our approach outperforms the state-of-the-art baselines in environments unseen during training, and (iv) our learned policy scales to larger multi-robot systems and more complex environments. We evaluate the performance of our approach in a UAV-based urban mapping scenario in a simulator, as well as conduct real robot experiments to demonstrate the practical applicability. Future research directions include extension to multi-robot pathfinding, task allocation, accounting for localization errors, and cooperative communication.

REFERENCES

- [1] J. Berger and N. Lo, "An innovative multi-agent search-and-rescue path planning approach," *Comput. Operations Res.*, vol. 53, pp. 24–31, 2015.
- [2] S. Yoon and C. Qiao, "Cooperative search and survey using autonomous underwater vehicles (AUVs)," *IEEE Trans. Parallel Distrib. Syst.*, vol. 22, no. 3, pp. 364–379, Mar. 2011.
- [3] A. Vashisth, J. Ruckin, F. Magistri, C. Stachniss, and M. Popovic, "Deep reinforcement learning with dynamic graphs for adaptive informative path planning," *IEEE Robot. Automat. Lett.*, vol. 9, no. 9, pp. 7747–7754, Sep. 2024.
- [4] Y. U. Cao, A. B. Kahng, and A. S. Fukunaga, "Cooperative mobile robotics: Antecedents and directions," *Robot Colonies*, vol. 4, no. 1, pp. 7–27, 1997.
- [5] J. Su, X. Zhu, S. Li, and W.-H. Chen, "AI meets UAVs: A survey on AI empowered UAV perception systems for precision agriculture," *Neurocomputing*, vol. 518, pp. 242–270, 2022.
- [6] A. Singh, A. Krause, C. Guestrin, W. Kaiser, and M. Batalin, "Efficient planning of informative paths for multiple robots," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2007, pp. 2204–2211.
- [7] A. Singh, A. Krause, C. Guestrin, and W. J. Kaiser, "Efficient informative sensing using multiple robots," *J. Artif. Intell. Res.*, vol. 34, pp. 707–755, 2009.
- [8] A. Singh, A. Krause, and W. J. Kaiser, "Nonmyopic adaptive informative path planning for multiple robots," in *Proc. 21st Int. Joint Conf. Artif. Intell.*, 2009, pp. 1843–1850.
- [9] S. Y. Luis, D. Shutin, J. M. Gómez, D. G. Reina, and S. T. Mariñ, "Deep reinforcement multi-agent learning framework for information gathering with local gaussian processes for water monitoring," *Adv. Intell. Syst.*, vol. 6, no. 8, 2024, Art. no. 2300850.
- [10] G. Hollinger and S. Singh, "Multi-robot coordination with periodic connectivity," in *Proc. IEEE Int. Conf. Robot. & Automat.*, 2010, pp. 4457–4462.
- [11] N. K. Yilmaz, C. Evangelinos, P. F. Lermusiaux, and N. M. Patrikalakis, "Path planning of autonomous underwater vehicles for adaptive sampling using mixed integer linear programming," *IEEE J. Ocean. Eng.*, vol. 33, no. 4, pp. 522–537, Oct. 2008.
- [12] A. Dutta, A. Ghosh, and O. P. Kreidl, "Multi-robot informative path planning with continuous connectivity constraints," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2019, pp. 3245–3251.
- [13] Y. Wei and R. Zheng, "Multi-robot path planning for mobile sensing through deep reinforcement learning," in *Proc. IEEE Conf. Comput. Commun.*, 2021, pp. 1–10.
- [14] G. A. Di Caro and A. W. Z. Yousaf, "Multi-robot informative path planning using a leader-follower architecture," in *Proc. IEEE Int. Conf. Robot. & Automat.*, 2021, pp. 10045–10051.
- [15] D. S. Diop, S. Y. Luis, M. P. Esteve, S. L. T. Mariñ, and D. G. Reina, "Decoupling patrolling tasks for water quality monitoring: A multi-agent deep reinforcement learning approach," *IEEE Access*, vol. 12, pp. 75559–75576, 2024.
- [16] A. M. Barrionuevo, S. Y. Luis, D. G. Reina, and S. L. T. Mariñ, "Informative deep reinforcement path planning for heterogeneous autonomous surface vehicles in large water resources," *IEEE Access*, vol. 12, pp. 71835–7185, 2024.
- [17] H. M. La, W. Sheng, and J. Chen, "Cooperative and active sensing in mobile sensor networks for scalar field mapping," *IEEE Trans. Syst., Man, Cybern.: Syst.*, vol. 45, no. 1, pp. 1–12, Jan. 2015.
- [18] N. Cao, K. H. Low, and J. M. Dolan, "Multi-robot informative path planning for active sensing of environmental phenomena: A tale of two algorithms," in *Proc. Int. Conf. Auton. Agents Multi-Agent Syst.*, 2013, pp. 7–14.
- [19] R. Cui, B. Gao, and J. Guo, "Pareto-optimal coordination of multiple robots with safety guarantees," *Auton. Robots*, vol. 32, pp. 189–205, 2012.
- [20] A. Viseras and R. Garcia, "DeepIG: Multi-robot information gathering with deep reinforcement learning," *IEEE Robot. Automat. Lett.*, vol. 4, no. 3, pp. 3059–3066, Jul. 2019.
- [21] J. Westheider, J. Rückin, and M. Popović, "Multi-UAV adaptive path planning using deep reinforcement learning," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2023, pp. 649–656.
- [22] T. Yang, Y. Cao, and G. Sartoretti, "Intent-based deep reinforcement learning for multi-agent informative path planning," in *Proc. 2023 Int. Symp. Multi-Robot Multi-Agent Syst.*, 2023, pp. 71–77.
- [23] S. Yanes Luis, M. Perales Esteve, D. Gutiérrez Reina, and S. Toral Mariñ, "Deep reinforcement learning applied to multi-agent informative path planning in environmental missions," in *Mobile Robot: Motion Control and Path Planning*. Berlin, Germany: Springer, 2023, pp. 31–61.
- [24] G. Best, O. M. Cliff, T. Patten, R. R. Mettu, and R. Fitch, "DEC-MCTS: Decentralized planning for multi-robot active perception," *Int. J. Robot. Res.*, vol. 38, no. 2-3, pp. 316–337, 2019.
- [25] F. Venturini et al., "Distributed reinforcement learning for flexible UAV swarm control with transfer learning capabilities," in *Proc. 6th ACM Workshop Micro Aerial Veh. Netw., Syst., Appl.*, 2020, pp. 1–6.
- [26] R. Cui, Y. Li, and W. Yan, "Mutual information-based multi-AUV path planning for scalar field sampling using multidimensional rrt," *IEEE Trans. Syst., Man, Cybern.: Syst.*, vol. 46, no. 7, pp. 993–1004, Jul. 2015.
- [27] B. J. Julian, M. Angermann, M. Schwager, and D. Rus, "Distributed robotic sensor networks: An information-theoretic approach," *Int. J. Robot. Res.*, vol. 31, no. 10, pp. 1134–1154, 2012.
- [28] A. Asgharivaskasi, F. Girke, and N. Atanasov, "Riemannian optimization for active mapping with robot teams," *IEEE Trans. Robot.*, vol. 41, pp. 1077–1097, 2024. [Online]. Available: <https://api.semanticscholar.org/CorpusID:269448652>
- [29] M. Tzes, N. Bousias, E. Chatzipantazis, and G. J. Pappas, "Graph neural networks for multi-robot active information acquisition," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2023, pp. 3497–3503.
- [30] C. Rasmussen and C. Williams, *Gaussian Processes for Machine Learning*. Cambridge, MA, USA: MIT Press, 2006.
- [31] Y. Cao, Y. Wang, A. Vashisth, H. Fan, and G. A. Sartoretti, "Catnipp: Context-aware attention-based network for informative path planning," in *Proc. Conf. Robot Learn.*, 2023, pp. 1928–1937.
- [32] D. M. S. Tan, Y. Ma, J. Liang, Y. C. Chng, Y. Cao, and G. Sartoretti, "IR 2: Implicit rendezvous for robotic exploration teams under sparse intermittent connectivity," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2024, pp. 13245–13252.
- [33] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, "Proximal policy optimization algorithms," *arXiv:1707.06347*.
- [34] J. Ott, E. Balaban, and M. J. Kochenderfer, "Sequential Bayesian optimization for adaptive informative path planning with multimodal sensing," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2023, pp. 7894–7901.
- [35] G. A. Hollinger and G. S. Sukhatme, "Sampling-based robotic information gathering algorithms," *Int. J. Robot. Res.*, vol. 33, no. 9, pp. 1271–1287, 2014.
- [36] J. Panerati, H. Zheng, S. Zhou, J. Xu, A. Prorok, and A. P. Schoellig, "Learning to fly—A gym environment with PyBullet physics for reinforcement learning of multi-agent quadcopter control," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2021, pp. 7512–7519.
- [37] D. Peralta, J. Casimiro, A. M. Nilles, J. A. Aguilar, R. Atienza, and R. Cajote, "Next-best view policy for 3D reconstruction," in *Proc. Eur. Conf. Comput. Vis.*, Glasgow, U.K., 2020, pp. 558–573.