

Multi-Agent Reinforcement Learning for Visibility-based Persistent Monitoring

Jingxi Chen,¹ Amrish Baskaran,¹ Zhongshun Zhang, and Pratap Tokekar

Abstract—The Visibility-based Persistent Monitoring (VPM) problem seeks to find a set of trajectories (or controllers) for robots to persistently monitor a changing environment. Each robot has a sensor, such as a camera, with a limited field-of-view that is obstructed by obstacles in the environment. The robots may need to coordinate with each other to ensure no point in the environment is left unmonitored for long periods of time. We model the problem such that there is a penalty that accrues every time step if a point is left unmonitored. However, the dynamics of the penalty are unknown to us. We present a Multi-Agent Reinforcement Learning (MARL) algorithm for the VPM problem. Specifically, we present a Multi-Agent Graph Attention Proximal Policy Optimization (MA-G-PPO) algorithm that takes as input the local observations of all agents combined with a low resolution global map to learn a policy for each agent. The graph attention allows agents to share their information with others leading to an effective joint policy. Our main focus is to understand how effective MARL is for the VPM problem. We investigate five research questions with this broader goal. We find that MA-G-PPO is able to learn a better policy than the non-RL baseline in most cases, the effectiveness depends on agents sharing information with each other, and the policy learnt shows emergent behavior for the agents.

I. INTRODUCTION

Multi-robot systems are often tasked with monitoring a large environment [1]–[6]. This includes applications such as patrolling where regularly revisiting the environment is crucial. When the robots have limited field-of-view, they may be able to see only a part of the environment at any given time. Therefore, the robots need to coordinate their actions so that they can monitor the environment effectively. In this paper, we investigate whether a team of agents can *learn* to cooperate for effective Visibility-based Persistent Monitoring (VPM).

We consider a scenario where a team of agents are tasked with persistently monitoring a 2D environment with known stationary obstacles. The agents have a limited field-of-view sensor whose visibility is obstructed by the obstacles. Any point that is not monitored accrues penalty at a rate unknown to the agents. The goal is to find trajectories for all agents so as to minimize the collective penalty.

The VPM problem is a variant of the well-studied visibility-based coverage problem. In such problems, the goal is to find a placement of robots such that they collectively see all points in the environment or maximize the area that

they cover while taking into account the geometric visibility constraints. In addition to coverage, the VPM problem also requires agents to revisit the previously visited positions in order to consistently cover the whole environment. As such, the problem we study is a generalization of the coverage problem. This makes it a more challenging problem to solve, especially if the dynamics of how the penalty accrued is unknown. Likewise, VPM is a generalization of the standard Persistent Monitoring (PM) problem [2], [7], [8] where in order to *monitor* a location, the robot needs to visit that location. The standard PM problem does not take visibility constraints into account.

We present a deep neural network architecture termed MA-G-PPO (Multi-Agent Graph Attention Proximal Policy Optimization) for solving this problem. The network takes as input the observations for each agent and extracts a lower dimensional feature representation using Convolutional Neural Network (CNN) [9], shares the features with other agents through Graph Attention (GAT) layer [10], and uses Proximal Policy Optimization (PPO) [11] for learning the policy. While this architecture is similar to the ones that have been recently proposed for multi-agent coordination [12], this is the first work to present a Multi-agent Reinforcement Learning (MARL) approach for the VPM problem. Nevertheless, our main contribution is an empirical investigation of the efficacy of MARL for VPM.

Specifically, we investigate the following five *research questions* in the context of the VPM problem: (1) Does MARL perform better than non-RL baselines for the VPM problem? (2) How does the information available to the agents affect the performance? (3) Are there emergent behaviors due to MARL? (4) How much does introducing communication improve the performance in MARL? (5) How well does MARL generalize to the number of agents?

We present our findings for each of these questions along with some conjectures that can inform future work. We hope that these findings can inform other applications of multi-agent RL to multi-robot cooperative tasks.

II. RELATED WORK

The stationary visibility-based coverage problem for a team of mobile agents is called Watchman Route Problem (WRP) [13], which is the mobile version of the Art Gallery Problem (AGP) [14]. The goal in the WRP is to design paths for agents to minimize the time required to visually cover the entire environment. There are variants of these problems that take into account practical restrictions on mobile robots

¹Chen and Baskaran contributed equally.

The authors are with the University of Maryland, College Park, USA {ianchen, amrishb, zszhang, tokekar}@umd.edu. This work is supported by the National Science Foundation under Grant No. 1943368 and the Office of Naval Research under Grant No. N000141812829.

(e.g., time to stop and take measurements [15]). In general, visibility-based coverage problems are NP-hard.

Several deterministic methods have been proposed for solving multi-agent persistent monitoring. [16] proposes a method of observing stochastic events at geographical nodes with a group of mobile agents by using a receding horizon sequential greedy algorithm to determine a sub-optimal policy with a polynomial cost and guaranteed bound on optimality. This is facilitated by showing that the reward function is a monotone sub-modular set function. There is also existing work on 2D PM [7] that focuses on monitoring a set of target points in the environment. They proposed a planner for robots to reduce the uncertainty on all target points. While there is significant work on the PM problem, there are major differences compared to the VPM problem we study here: The environment in PM problems is typically represented as a graph with vertices or a set of points on the plane that need to be visited whereas here we have a continuous 2D environment that needs to be persistently monitored. Furthermore, in the 2D PM problems, there is no consideration of agents' visibility unlike VPM, where visiting a point may not be necessary.

For multi-agent systems, communication between agents can be modeled as a Graph Convolution Neural Network (GCN) [17], in this work the agents are sharing their local observation through a graph neural network and using the attention mechanism to assign weights for aggregating the received information, finally based on this aggregated information RL is used for decision making. The GAT [10] is an improvement over GCN, dealing with the dynamic communication graph in multi-agent systems. The work most closely related to ours is [18], [19]. Both papers focus on multi-robot coordination with supervised learning and reinforcement learning, respectively. However, these papers focus on simpler tasks such as navigation and static coverage. Here, we explicitly focus on VPM. Furthermore, we seek to understand the efficacy of MARL for VPM, something that has not been studied yet.

III. PERSISTENT MONITORING

In this section, we formally describe the persistent monitoring problem. We represent the environment as a 2D grid world K . Each cell in the environment is to be monitored by the agents. We associate a reward value, i.e., negative-valued penalty, with each cell. Let $R^{(k)}(t)$ be the reward associated with cell $k \in K$ at time t . This reward is a function of the last time any agent monitored the cell k . Let D be a non-negative decay rate for the reward. If the cell k is within the field of view of any agent at time t , then the reward will be reset to 0; otherwise, the reward for cell k decays with rate D . We define the reward associated with cell k as follows: $R^{(k)}(t+1) = 0$, if k is monitored at t , and $R^{(k)}(t+1) = \max\{R^{(k)}(t) - D, -R_{\max}\}$ otherwise.

All cells are initialized with a reward, $R^{(k)}(0) = 0$. Since $D \geq 0$, $R^{(k)}(t)$ will be non-positive. An equivalent interpretation of $R^{(k)}(t)$ is to think of it as the negative of the penalty associated with not viewing the cell k . The

longer the time spent between successive viewings, the larger the accrued penalty. R_{\max} is used to limit the penalty so that it does not grow unbounded. The objective of the persistent monitoring problem is to find an optimal policy π to maximize the cumulative reward collected by all agents in a finite horizon of length T : maximize $\sum_{t=0}^T \sum_{k \in K} R_{\pi}^{(k)}(t)$. We show the dependence on π to represent the fact that the reward collected at any time t is a function of the policy for the agents until time t . The standard coverage problem [20] is a special case where the D is set to 0 once a cell is viewed by some agent for the first time.

The policy maps the states to actions for each agent. Each agent can choose from one of four neighbors and staying in place. In general, we do not assume that D is known to the algorithm; thereby making this an RL problem.

Since the reward is a function of the history of the position of the agents, the state must contain not only the positions of the agents but also the reward values for each cell in the grid at the current time step. This ensures the Markov property. However, the state space is large and therefore, we resort to using RL to solve the problem. We investigate the effect of other types of state representation as discussed next, instead of the full grid map.

An instance of the simulated environment is shown in Figure 1. The figure shows the following elements: (1) Agents: represented in yellow dots. (Note: in the penalty scale, the yellow region is not to represent the agents.) (2) Static Obstacles in red (diagonal fill). (3) Free cells: Agents can only traverse through these free cells. The color of the free cells represent $R^{(k)}(t)$ ranging from blue (0) to red ($-R_{\max}$).

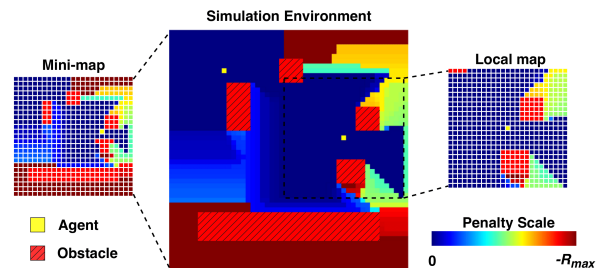


Fig. 1. From left: Mini-map, Simulation environment and Local map

We assume that each agent has a limited visibility coverage range with a square field-of-view of side length L centered at the agent. This is shown in the Figure 1 within the black dashed box. We consider both scenarios where agents have access to the global information of all cells or have only the local information of the area visible at the current time step. We investigate the effect of learning with: (1) Mini-map: As shown in Figure 1, The mini-map is a lower resolution form of the simulation environment. The mini-map contains the position of the current agent only with penalty values and position of the obstacles. (2) Local map: As shown in Figure 1, provides a higher resolution map of the environment around the agent at the current time step (in the ego-centric frame of the agent). The local map is essentially

the observation given by the sensor onboard the agent. We investigate the effect of learning with only mini-map, with only local map, and with both local and mini-map.

IV. PROPOSED APPROACH: MA-G-PPO

In this section, we describe the approach we use to solve the VPM problem discussed in the previous section. Our architecture uses the communication-based MARL structure [12] along with GAT [10] to model communication. We term this architecture as Multi-agent PPO with GAT (MA-G-PPO).

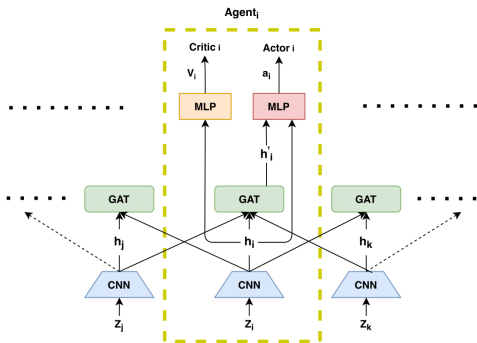


Fig. 2. MA-G-PPO architecture.

Figure 2 shows the schematic of MA-G-PPO. The structure in the yellow dashed box is replicated for all agents. Each agent gets a copy of the same structure (with identical weights). This allows us to generalize the setup by replicating this structure for additional agents. The structure consists of the following modules that we describe in detail: CNN for extracting features from the inputs for each agent; GAT to share information across agents; and PPO for finding the optimal policy for each agent.

A. Local Feature Extraction

The input to the network is the observation z_i for all agents i at time t . We use a CNN to extract a feature vector h_i from z_i . This feature vector is then passed to the GAT and shared with other agents. We investigate three types of observations z_i : (i) only the local map; (ii) only the global mini-map; and (iii) the combined local and mini-map (as shown in Figure 1). We represent the local map and the mini-map as an image where the cells corresponding to the obstacle, agent and free cells are 150, 200 and penalty ranging from 0 to $-R^{max}$ respectively. In case (iii), the input z_i is a two channel image consisting of the local map and mini-map images. As shown in Section V, case (iii) performs the best which is not surprising since it contains more information. This embedding layer runs locally on each agent.

B. Graph Attention Network

After the local feature extraction, we use GAT in order to share the local feature vectors with neighboring agents. This is represented as a weighted edge between the agents (nodes) i and j . Each agent aggregates all the incoming feature vectors into, h'_i , as shown in the Figure 2.

The learnable attention value $\alpha_{i,j}$ gives how much weight agent i gives to the feature vector from j . The attention weight value on the edge between i and j is calculated using:

$$\alpha_{i,j}^m = \frac{\exp(\text{LEAKYRELU}(\mathbf{a}(\mathbf{W}^m h_i, \mathbf{W}^m h_j)))}{\sum_{k \in N_i} \exp(\text{LEAKYRELU}(\mathbf{a}(\mathbf{W}^m h_i, \mathbf{W}^m h_k)))} \quad (1)$$

Here, $h_i \in \mathbb{R}^F$ is the feature vector of node i ; N_i is the set of neighboring nodes of node i ; $W \in \mathbb{R}^{F' \times F}$ is a learnable linear transformation matrix; $\mathbf{a} : \mathbb{R}^{F'} \times \mathbb{R}^{F'} \rightarrow \mathbb{R}$ is a single layer feedforward neural network which produces a real value; LEAKYRELU is the activation function applied to the output of \mathbf{a} . We are using multi-head attention for m heads as it can capture more types of relationships and makes training stable [17].

After a node i aggregates feature vectors from its neighbors N_i , the aggregated vector h'_i will be:

$$h'_i = \frac{1}{M} \sum_{m=1}^M \left(\sum_{j \in N_i} \alpha_{i,j}^m \mathbf{W}^m h_j \right) \quad (2)$$

This way of updating h'_i is permutation-invariant to the order of nodes in its neighborhood N_i . This is because the attention weight $\alpha_{i,j}$ between agent i and its neighbor j is a function of the features h_i and h_j .

In our simulations, we consider a fully connected network. That is N_i is the set of all other agents except i . However, any other neighborhood relation can directly be incorporated in the architecture. For example, N_i can be based on the communication range of the agents.

C. Multi-agent PPO

With the aggregated feature vector h'_i from GAT along with its own feature vector h_i are fed as the input to the next layers for each agent. Here, we choose PPO, a state-of-the-art actor-critic method for learning the optimal policy. Policy gradient actor-critic methods, in particular PPO, generally outperforms DQN in the context of path planning tasks [21].

The PPO portion of the network consists of an actor and a critic head. The parameters of the actor are denoted by θ_1 and those of the critic are denoted by θ_2 . Since the VPM problem is a cooperative task, the reward for all the agent is the same. The shared reward at time t is the summation of reward values for all unoccupied cells: $r(t, \theta_1) = \sum_{k \in K} R_{\pi(\theta_1)}^{(k)}(t)$.

Let $G(t, \theta_1) = \sum_{t=0}^T \gamma^t r(t, \theta_1)$ be the discounted sum of rewards collected from the start ($t = 0$) to the end ($t = T$) of an episode where γ is the discount factor. Let $V^i(s_t^i, \theta_2)$ be the estimate of the expected return for the state s_t^i for agent i , $V^i(s_t^i, \theta_2) = \mathbb{E}[G(t)|s_t^i]$. This estimate is provided by the critic network. The advantage value for agent i at time t is: $A_t^i(\theta_1, \theta_2) = G(t, \theta_1) - V^i(s_t^i, \theta_2)$. The surrogate loss in PPO for agent i is: $\text{surr}^i(\theta_1, \theta_2) = \mathbb{E}_t[\min_{\theta_1, \theta_2} \{ratio_t A_t^i, clip(ratio_t, 1 - \epsilon, 1 + \epsilon) A_t^i\}]$. Here, the $clip$ and $ratio$ are as defined in [11]. We minimize the following loss function:

$$L(\theta_1, \theta_2) = \frac{1}{S} \sum_S \frac{1}{N} \sum_{i=1}^N (\text{surr}^i) \quad (3)$$

where N is number of agents, S is the size of the mini-batch.

V. EXPERIMENTAL EVALUATION

We conducted a set of experiments to evaluate the proposed approach for solving the VPM problem. Our goal for the experiments was to answer the five questions posed in the introduction with the overarching goal of assessing the efficacy of MARL methods for solving the VPM problem. In the following, we first describe our experimental setup followed by a discussion of the investigation we carried out for the five research questions.

Experimental Setup: We carried out simulations in a custom-built environment in Python.¹ The training and testing are performed on a computer with Ubuntu 16.04 LTS Operation system, an Intel(R) Xeon(R) Silver CPU, and a GeForce RTX 2080 Ti GPU.

The MA-G-PPO network was implemented in Pytorch. We use the Deep Graph Library (DGL) [22] for graph neural network operations.

The training is conducted with $T = 1000$ steps per episode and testing with $T = 2000$ steps per episode on 50×50 maps and field-of-view $L = 25$. In the following sections, we will use notation N_{train} as the number of agents used during the training and N_{test} as number of agents during the deployment/testing. For updating penalty value on cells, decay rate $D = 1$ and the maximum penalty value on a cell $R^{max} = 400$. The local map and the mini-map are 25×25 . The size of the feature vectors are $F = F' = 128$. We use three attention heads, i.e., $m = 3$. All N_{test} values are averaged over 10 test runs for the learned policy.

Actual Training Time: Here is a sample training time to show the practical time constraint on the total number of training episodes we can run. For 4 agents, the typical MA-G-PPO training time for 10k episodes is almost 16 hours.

Recall that the reward values as defined in Section III are non-positive. For ease of understanding we view it as a minimization problem and look at the absolute value of the rewards i.e. minimizing total penalty.

We now discuss the results of our investigation of the five research questions.

A. Does MARL perform better than non-RL baselines for the VPM problem?

We compare our MA-G-PPO with three non-RL baselines: Greedy Centralized Search (GCS), Travelling Salesman Problem-Cyclic (TSPC) and Random policy.

The GCS selects a destination cell for every *unassigned* agent at every time instance as follows. We first find the subset of cells, say S , that have non-zero penalties. Then, we sort the cells in the descending order of penalties. Next, we iteratively choose the cell in S with the highest penalty and add it to a candidate list, say C . Once a cell c is added to C , we iterate through S and remove from it any cell that is within a predefined $L2$ distance from c . We repeat this until S is empty. Then, we greedily assign every *unassigned*

¹The code of MA-G-PPO is available at https://github.com/raaslab/rl_multi_agent.

agent, in a predefined order, to the nearest cell in C such that no cell is assigned to more than one agent. Then, the agent follows the shortest path using Dijkstra to the assigned cell. An agent is reassigned when it reaches its assigned cell or if the assigned cell is no longer present in the set S .

In a context where the location of nodes are determined and the routes between them are undetermined, [23] proposed an optimal cyclic patrolling scheme for persistent monitoring. Similar approaches have been devised for other variants of PM [23]–[27]. We cannot directly use these strategies as a baseline since they do not take the visibility into account. Instead, we take inspiration from these methods and devise a baseline that does find a cyclic tour for all robots to follow. To find this tour, we first find points that collectively sees all points in the environment. Finding the shortest tour in an environment with obstacles is NP-complete [28]. However, we can find an approximation of the shortest tour. In our case, we first find a set of points such that visiting these points is sufficient to ensure all points in the environment are seen at least once (we do this manually for the test environments). Then, we find a TSP tour that visits all the points using Machine Learning Randomized Optimization and SEARCH (mlrose) solver [29]. The agents are then evenly placed along this tour and they execute the entire tour simultaneously.

The random policy samples an action for each agent from a uniform distribution at every time instance.

	$N_{train} = N_{test}$				$N_{train} = N_{test}$				
	2	4	6	10	2	4	6	10	
MA-G-PPO	141.8	44.0	12.5	4.0	159.4	63.6	41.5	4.4	750 500 250 0
GCS	191.7	61.5	17.7	4.0	115.8	50.4	23.4	10.6	
TSPC	127.7	36.5	15.9	9.1	70.0	26.3	13.8	6.3	
Random	822.8	520.1	332.1	116.7	604.8	482.0	286.7	189.6	
	Open map				8 Room map				

Fig. 3. Comparison of cumulative penalty (unit: 10^6) between four algorithms on two maps. Note that TSPC finds an approximation of the shortest tour with optimal preset points, or the near-optimal policy.

Figure 3 shows the results of comparison between the four algorithms for a variety of N_{train} and N_{test} values. Not surprisingly, MA-G-PPO, GCS and TSPC always outperform the random policy. On open map, for MA-G-PPO, GCS and TSPC, as number of agents decrease from 10 to 2, the gap between their performance increases. When there are fewer agents, MA-G-PPO performs much better than GCS. But with 10 agents MA-G-PPO and GCS perform identically. This is not surprising. As number of agents increases, the solution of problem becomes closer to static coverage. TSPC follows a cyclic policy irrespective of number of agents, hence its performance degrades compared to MA-G-PPO and GCS as the solution moves towards a static coverage solution.

This can be seen in Figure 4. We show the policy learned by MA-G-PPO with 4 and 10 agents on the open map. The shaded trail shows the trajectories followed by the agents. As we can see, with 4 agents, the learned policy has the agents

cover larger distances effectively carrying out a tour of the environment. But with 10 agents, we see that the agents all remain in almost the same positions. When there are fewer agents, there is a need for careful planning. This is where MA-G-PPO outperforms GCS by a larger margin. Regardless of the number of agents, we see that MA-G-PPO is able to learn a good solution that outperforms GCS.

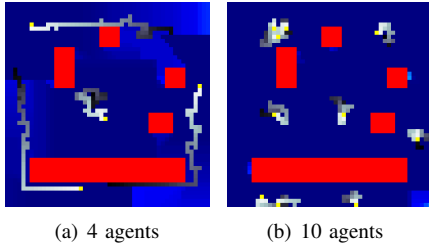


Fig. 4. Learned policy figures for 4 and 10 agents in the open map.

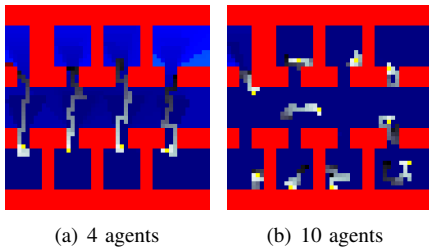


Fig. 5. Learned policy for 4 and 10 agents in the 8-Room map.

We also carried out the same experiment in a more structured and larger environment (Figure 5) with 8 rooms. Here, we find that GCS performs better than the policy learned by MA-G-PPO after 30k episodes. We suspect that 30k episodes are not sufficient for learning a better policy in this more structured environment. Part of the reason is that the optimal policy here requires precise movement (each agent must monitor two of the rooms, move along a vertical line, and move in a way so as to be out of phase with respect to each other). After 30k episodes, the agent learns a policy that has the first two properties but not the third. However, as we see from the trend shown in Figure 6 the policy has not converged after 30k episodes for the map with 8 rooms. In fact, in the open map case, the policy is learnt quickly and we see marginal improvement over the 30k episodes. But with the larger map, the policy consistently improves. We suspect that by training longer, we would end up with a policy that does outperform GCS as we see in smaller environments. In the case of a static solution, MA-G-PPO performs better than TSPC.

From the results, we conclude that MA-G-PPO effectively plans the path for the agents for VPM ranging from scenarios requiring only a nearly-static deployment to those requiring carefully chosen trajectories. We study the structure of these trajectories further in Section V-C.

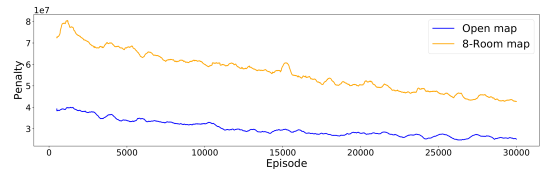


Fig. 6. Rolling average of 500 episodes of the training performance for $N_{train} = 4$.

B. How does the information available to the agents affect the performance?

We are interested in understanding the effect of the available information on the performance of the agent. We consider three options for the input z_i for each agent: the local map and mini-map independently and together.

We trained the policy in each of the three cases for a total of 30k episodes. We plot the total penalty accumulated across each episode for each of the inputs used, in Figure 7. We see that training with only mini-map learns faster and is more stable than using only local map. This is not surprising since the mini-map has full access to the global information but just at a lower resolution. However, there is a diminishing return to training with just the mini-map. Global information at a lower resolution is not enough to plan trajectories through complex structures that do not show up on the mini-map.

We see that the local and mini-map together far outperforms learning with just one of the prior maps, right from the outset. It must be noted however, that despite having only local map, the network is still able to perform well (even outperforming global mini-map information) given sufficient training time. Due to the high computation cost and the risk of over-fitting, a more faster and stable option is chosen. Hence, in the rest of the simulations, we use the networks trained on both local and mini-map.

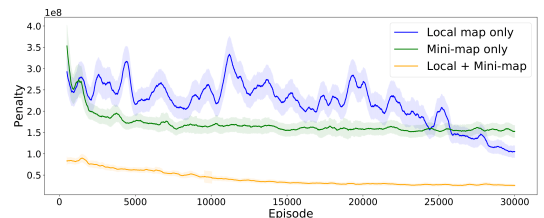


Fig. 7. Training performance with three input options to MA-G-PPO using rolling average over 500 episodes and the $\pm 0.5std$ error band.

C. Are there emergent behaviors due to MARL?

The persistent monitoring problem has a strong underlying geometric structure. We investigated whether MA-G-PPO can find structure in the policy learned and lead to emergent behaviors. In particular, we are interested in understanding and interpreting the behaviors learned by the agent, if any. We investigated this through two types of environments. The first environment is structured with four rooms as shown in Figure 8(a). There are two agents. MA-G-PPO learns a natural policy: each agent is “assigned” to monitor two of the rooms. Figure 8(a) shows the trail of the agents paths.

The agents oscillate between the two rooms. Interestingly, there is a phase difference between the two oscillations. We see this in the y -coordinates plotted in Figure 8(b) for both agents over time. We also observe that the learned policy is periodic and with a fixed phase difference between the two agents. We believe by learning a phase difference, the agents take turns monitoring the middle corridor ensuring that the penalty values there do not reach a higher value.

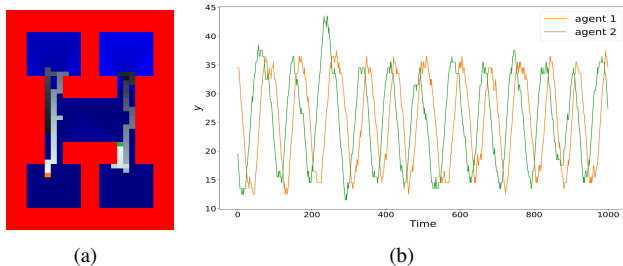


Fig. 8. (a) Trajectory of learned policy for two agents in this 4-Room Map. (b) Phase difference between the agents in y -coordinates.

We observe a similar emergent behavior in the second type of environment which is an open one (Figure 4(a)). Here, we train with four agents. MA-G-PPO again learns a policy with an intuitive structure: one agent at the center remains nearly stationary throughout the episode and the other three agents move along a tour around the periphery of the environment. To observe the structure, we plot $\sin(\theta_i)$ for each of the three outer agents in Figure 9(a). Here θ_i are their angular coordinates with respect to the center of the environment. We can observe that the three agents follow a nearly periodic policy with a fixed phase difference as before.

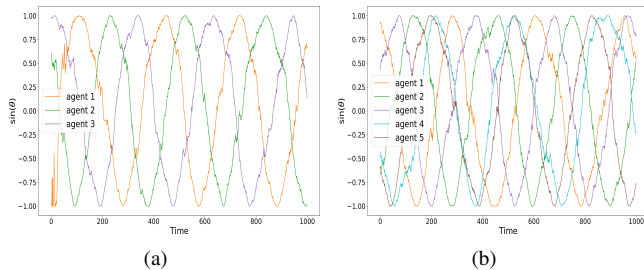


Fig. 9. In open map. (a) polar-coordinate plot for $N_{test} = 4$ (c.f. Figure 4(a)). (b) polar-coordinate plot for $N_{test} = 6$.

In [30], Alamdari et al. introduced the notion of *kernels* for infinite horizon persistent monitoring paths. A kernel is a finite length path which when repeated forms the infinite horizon solution to the original problem. We conjecture that MA-G-PPO is able to learn the kernel in structure environments. The two sets of results shown here provide evidence of such a learned kernel (i.e., periodic paths).

D. How much does introducing communication improve the performance in MARL?

To investigate this question, we compared two versions: one with communication (MA-G-PPO) and one without it (MA-PPO). The latter is the same as the former but without

GAT. We train both models with 2, 4, 6, and 10 agents on the same map. Both models are trained separately and then deployed with various number of agents.

We show the comparison between the learned policies in Figure 10. First consider the diagonal entries in the table that correspond to the case where $N_{train} = N_{test}$. We observe that in all cases, the architecture with communication outperforms the one without it. As the mini-map of an agent does not contain location information of other agents, it cannot predict or coordinate with each other. GAT allows agents to communicate their global position and observation with each other, in the form of aggregated information. We conjecture that the aggregated information facilitates agents' coordination on a global scale.

E. How well does MARL generalize to the number of agents?

In the results discussed so far, we assume that the number of agents during training are the same as the number of agents in deployment, i.e., $N_{train} = N_{test}$. In practice, that may not be the case. For example, some robots may malfunction during deployment and retraining a policy from scratch for the new number of robots may be infeasible. Therefore, we also investigated how well does MA-G-PPO generalize when $N_{train} \neq N_{test}$. We also compare this to the generalization abilities of MA-PPO.

This corresponds to the off-diagonal entries in Figure 10. Consider each column in the table (corresponding to the same number of agents N_{test} deployed with varying number of agents during training). Not surprisingly, the penalty is lowest when $N_{train} = N_{test}$. But we also observe that MA-G-PPO generalizes better than MA-PPO. For example, consider the second column corresponding to $N_{test} = 4$. When $N_{train} = N_{test}$, both MA-G-PPO and MA-PPO perform close to each other (44.0 vs 46.1). However, as the gap between N_{train} and N_{test} increases, so does the gap between MA-G-PPO and MA-PPO. In the extreme case when $N_{test} = 4$ and $N_{train} = 10$, MA-G-PPO still yields a penalty of 98.5 compared to the much higher penalty of 313.7 with MA-PPO. A similar trend is seen in other cases.

One possible way of explaining this is that the communication help agents learn the structure of the solution (periodicity and phase difference) as described earlier. For MA-G-PPO, as we can see in Figure 9(a), the emergent behavior learned when $N_{test} = 4$ translates to the case shown in Figure 9(b) when $N_{test} = 6$. In both cases, $N_{train} = 4$. We observe the same structure in the policy: one agent in the center, the remaining ones performing a tour around the periphery with the same period and fixed phase difference. Figure 9(b) shows that even if we train on four agents and deploy on six agents, the structure in the policy (one agent at the center, five along the periphery) remains similar to the $N_{test} = 4$ case. The phase difference becomes smaller with more agents.

In contrast, when we use MA-PPO to train the model, the policy we get is as Figure 11. The learned trajectory is that one agent stays in the upper-left corner and the remaining $N_{test} - 1$ agents move along a tour at the lower-right corner. The polar-coordinate plot is shown in Figure 12.

		N_{test}				N_{test}				
		2	4	6	10	2	4	6	10	
N_{train}	2	141.8	86.8	54.8	30.0	316.7	244.7	223.3	207.1	750 500 250 0
	4	185.5	44.0	21.8	6.4	212.4	46.1	27.2	11.4	
	6	370.1	50.2	12.5	7.4	266.8	70.8	22.3	8.8	
	10	424.8	98.5	33.9	4.0	857.5	313.7	52.3	7.7	
		MA-G-PPO				MA-PPO				

Fig. 10. Average cumulative penalty (units: 10^6) per episode.

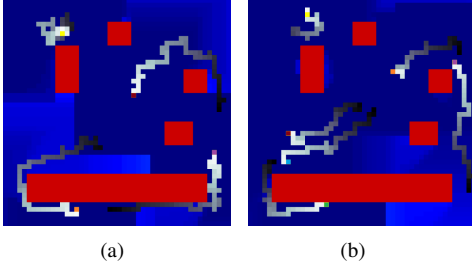


Fig. 11. Training MA-PPO with $N_{train} = 4$. (a) Trajectory for $N_{test} = 4$. (b) Trajectory for $N_{test} = 6$

For MA-PPO, the structure (especially the phase difference part) in the solution of $(N_{train} = 4, N_{test} = 4)$ does not generalize to $(N_{train} = 4, N_{test} = 6)$ case. This suggests a lack of emergent behavior in the case of MA-PPO as compared to MA-G-PPO which precludes generalization.

We conjecture that the communication is likely the reason for learning the structure in the solution which leads to MA-G-PPO outperforming MA-PPO.

VI. CONCLUSION

Our emphasis in this paper was to investigate a series of questions regarding the efficacy of using MARL for VPM. The experiment results show that given sufficient training time, the proposed approach may be better than TSPC baseline if there exists a nearly-static solution and GCS baseline otherwise. We also show that agents that use a combination of local and global information far outperform agents trained with only one type of information. Furthermore, the agent trained with higher-resolution local information eventually performs better than those with lower-resolution global information. We also see that communication in MA-G-PPO leads to better performance than no communication. One interesting feature we discover is that MA-G-PPO learns emergent behavior for agents. Specifically, we observe that the agents learn a periodic policy with a fixed phase difference in the agents' paths. As mentioned earlier, we conjecture that MA-

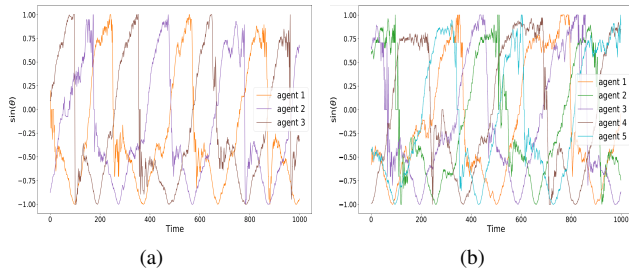


Fig. 12. In Open map. (a) Polar-coordinate plot for $N_{test} = 4$. (b) Polar-coordinate plot for deployment $N_{test} = 6$.

G-PPO is able to learn the kernels [30] of infinite horizon paths in some structured environments. Furthermore, this structure is preserved even when the number of agents that are deployed is not the same as the number of agents during training. This leads to better generalization. The results presented here provide evidence of this in Visibility based Persistent Monitoring settings.

REFERENCES

- [1] J. Yu, S. Karaman, and D. Rus, "Persistent monitoring of events with stochastic arrivals at multiple stations," *IEEE Transactions on Robotics*, vol. 31, no. 3, pp. 521–535, 2015.
- [2] S. K. K. Hari, S. Rathinam, S. Darbha, K. Kalyanam, S. G. Manyam, and D. Casbeer, "Persistent monitoring of dynamically changing environments using an unmanned vehicle," 2019.
- [3] S. K. K. Hari, S. Rathinam, S. Darbha, K. Kalyanam, S. G. Manyam, and D. Casbeer, "The generalized persistent monitoring problem," in *2019 American Control Conference (ACC)*, 2019, pp. 2783–2788.
- [4] F. R. Inacio, D. G. Macharet, and L. Chaimowicz, "Persistent monitoring of multiple areas of interest with robotic swarms," in *2018 Latin American Robotic Symposium, 2018 Brazilian Symposium on Robotics (SBR) and 2018 Workshop on Robotics in Education (WRE)*, 2018, pp. 82–87.
- [5] S. L. Smith, M. Schwager, and D. Rus, "Persistent monitoring of changing environments using a robot with limited range sensing," in *2011 IEEE International Conference on Robotics and Automation*, 2011, pp. 5448–5455.
- [6] X. Lin and C. G. Cassandras, "An optimal control approach to the multi-agent persistent monitoring problem in two-dimensional spaces," in *52nd IEEE Conference on Decision and Control*, 2013, pp. 6886–6891.
- [7] C. Song, L. Liu, G. Feng, and S. Xu, "Optimal control for multi-agent persistent monitoring," *Automatica*, vol. 50, no. 6, pp. 1663 – 1668, 2014. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S0005109814001411>
- [8] S. Alamdari, E. Fata, and S. L. Smith, "Persistent monitoring in discrete environments: Minimizing the maximum weighted latency between observations," *The International Journal of Robotics Research*, vol. 33, no. 1, pp. 138–154, 2014. [Online]. Available: <https://doi.org/10.1177/0278364913504011>
- [9] Y. Bengio and Y. Lecun, "Convolutional networks for images, speech, and time-series," 11 1997.
- [10] P. Veličković, G. Cucurull, A. Casanova, A. Romero, P. Liò, and Y. Bengio, "Graph attention networks," 2018.
- [11] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, "Proximal policy optimization algorithms," 2017.
- [12] J. Jiang, C. Dun, T. Huang, and Z. Lu, "Graph convolutional reinforcement learning," 2020.
- [13] S. Carlsson and B. J. Nilsson, "Computing vision points in polygons," *Algorithmica*, vol. 24, no. 1, pp. 50–75, 1999.
- [14] J. O'Rourke, *Art gallery theorems and algorithms*. Oxford University Press Oxford, 1987.
- [15] P. Tokekar and V. Kumar, "Visibility-based persistent monitoring with robot teams," in *Intelligent Robots and Systems (IROS), 2015 IEEE/RSJ International Conference on*. IEEE, 2015, pp. 3387–3394.
- [16] N. Rezazadeh and S. S. Kia, "A sub-modular receding horizon approach to persistent monitoring for a group of mobile agents over an urban area," *IFAC-PapersOnLine*, vol. 52, no. 20, pp. 217 – 222, 2019, 8th IFAC Workshop on Distributed Estimation and Control in Networked Systems NECSYS 2019. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S2405896319320129>
- [17] J. Jiang, C. Dun, and Z. Lu, "Graph convolutional reinforcement learning for multi-agent cooperation," *CoRR*, vol. abs/1810.09202, 2018. [Online]. Available: <http://arxiv.org/abs/1810.09202>
- [18] J. Blumenkamp and A. Prorok, "The emergence of adversarial communication in multi-agent reinforcement learning," 2020.
- [19] Q. Li, F. Gama, A. Ribeiro, and A. Prorok, "Graph neural networks for decentralized multi-robot path planning," 2020.
- [20] E. Galceran and M. Carreras, "A survey on coverage path planning for robotics," *Robotics and Autonomous Systems*, vol. 61, no. 12, pp. 1258 – 1276, 2013. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S092188901300167X>

- [21] R. S. Sutton, D. McAllester, S. Singh, and Y. Mansour, "Policy gradient methods for reinforcement learning with function approximation," in *Proceedings of the 12th International Conference on Neural Information Processing Systems*, ser. NIPS'99. Cambridge, MA, USA: MIT Press, 1999, p. 1057–1063.
- [22] M. Wang, D. Zheng, Z. Ye, Q. Gan, M. Li, X. Song, J. Zhou, C. Ma, L. Yu, Y. Gai, T. Xiao, T. He, G. Karypis, J. Li, and Z. Zhang, "Deep graph library: A graph-centric, highly-performant package for graph neural networks," 2020.
- [23] J. Yu, S. Karaman, and D. Rus, "Persistent monitoring of events with stochastic arrivals at multiple stations," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2014.
- [24] Z. Zhang, J. Lee, J. M. Smereka, Y. Sung, L. Zhou, and P. Tokekar, "Tree search techniques for minimizing detectability and maximizing visibility," in *2019 International Conference on Robotics and Automation (ICRA)*. IEEE, 2019, pp. 8791–8797.
- [25] Z. Zhang, J. M. Smereka, J. Lee, L. Zhou, Y. Sung, and P. Tokekar, "Game tree search for minimizing detectability and maximizing visibility," *Autonomous Robots*, pp. 1–15, 2021.
- [26] C. Baykal, G. Rosman, K. Kotowick, M. Donahue, and D. Rus, *Persistent Surveillance of Events with Unknown Rate Statistics*. Cham: Springer International Publishing, 2020, pp. 736–751. [Online]. Available: https://doi.org/10.1007/978-3-030-43089-4_47
- [27] X. Lan and M. Schwager, "Planning periodic persistent monitoring trajectories for sensing robots in gaussian random fields," in *2013 IEEE International Conference on Robotics and Automation*, 2013, pp. 2415–2420.
- [28] P. Tokekar, A. K. Budhiraja, and V. Kumar, "Algorithms for visibility-based monitoring with robot teams," *arXiv preprint arXiv:1612.03246*, 2016.
- [29] G. Hayes, "mlrose: Machine Learning, Randomized Optimization and SEarch package for Python," <https://github.com/gkhayes/mlrose>, 2019, accessed: day month year.
- [30] S. Alamdari, E. Fata, and S. L. Smith, "Persistent monitoring in discrete environments: Minimizing the maximum weighted latency between observations," *The International Journal of Robotics Research*, vol. 33, no. 1, pp. 138–154, 2014.