# Graph Neural Networks for Decentralized Multi-Robot Target Tracking

Lifeng Zhou,<sup>1\*</sup> Vishnu D. Sharma,<sup>2\*</sup> Qingbiao Li,<sup>3</sup> Amanda Prorok,<sup>3</sup> Alejandro Ribeiro,<sup>1</sup> Pratap Tokekar,<sup>2</sup> Vijay Kumar<sup>1</sup>

Abstract-The problem of decentralized multi-robot target tracking asks for jointly selecting actions, e.g., motion primitives, for the robots to maximize target tracking performance with local communications. One major challenge for practical implementations is to make target tracking approaches scalable for large-scale problem instances. In this work, we propose a general-purpose learning architecture towards collaborative target tracking at scale, with decentralized communications. Particularly, our learning architecture leverages a graph neural network (GNN) to capture local interactions of the robots and learns decentralized decision-making for the robots. We train the learning model by imitating an expert solution and implement the resulting model for decentralized action selection involving local observations and communications only. We demonstrate the performance of our GNN-based learning approach in a scenario of active target tracking with large networks of robots. The simulation results show our approach nearly matches the tracking performance of the expert algorithm, and yet runs several orders faster with up to 100 robots. Moreover, it slightly outperforms a decentralized greedy algorithm but runs faster (especially with more than 20 robots). The results also exhibit our approach's generalization capability in previously unseen scenarios, e.g., larger environments and larger networks of robots.

## I. INTRODUCTION

Multi-robot target tracking finds a wealth of applications in robotics. Typical examples include monitoring [1], patrolling [2], surveillance [3], and search and rescue [4]. Such applications ask for teams of robots that act as mobile sensors to jointly plan their actions to optimize tracking objectives (e.g., the number of tracked targets or the uncertainty reduction in the targets' positions). In this paper, we focus on tracking objective functions that are submodular, i.e., the functions that have the diminishing returns property. Examples of such functions include information-theoretic metrics such as entropy and mutual information [5] and the geometric metrics such as the visibility region [6]. The problems of maximizing submodular functions are generally NP-hard. The most well-known approach for tackling these problems

\*These authors contributed equally.

<sup>2</sup>V. D. Sharma and P. Tokekar are with the Department of Computer Science, University of Maryland, College Park, MD, USA (email: vishnuds@umd.edu, tokekar@umd.edu).

 $^{3}Q$ . Li and A. Prorok are with the Department of Computer Science and Technology, University of Cambridge, Cambridge, United Kingdom (email: {ql295, asp45}@cam.ac.uk).

This research was sponsored by the Army Research Lab through ARL DCIST CRA W911NF-17-2-0181, National Science Foundation under grant number 1943368, and ONR under grant number N00014-18-1-2829.

is the greedy algorithm that runs in polynomial time and vields a constant factor approximation guarantee [7], [8].

The greedy algorithm cannot be directly implemented in the scenarios where the robots can only communicate locally due to a limited communication range. To address the issue of local communication, some decentralized versions of the greedy algorithm were designed, where only neighboring information is utilized to choose actions for the robots for optimizing submodular objectives [9]-[12]. For example, building on the local greedy algorithm [8, Section 4], Atanasov et al. designed a decentralized greedy algorithm that achieves 1/2 approximation bound for multi-sensor target tracking [9]. Specifically, the algorithm greedily selects an action for each robot in sequential order, given all the actions selected so far. However, with limited communication, the robots may not have access to all the previously selected actions. To this end, a few decentralized submodular maximization algorithms were devised to execute a sequential greedy algorithm over directed acyclic graphs that may not be connected [11], [12]. Other decentralized greedy approaches include the ones that utilize a consensus-based mechanism to bring robots to an agreement by communicating local greedy selections with neighbors over multiple hops [10], [13]. However, these algorithms may take a considerable amount of time to reach a consensus.

In this paper, we aim to explore learning-based methods to learn policies by imitating expert algorithms [14] (e.g., the greedy algorithm [8]) for multi-robot target tracking. Particularly, we choose the graph neural network (GNN) as the learning paradigm given its nice properties of decentralized communication architecture that captures the neighboring interactions and the transferability that allows for the generalization to previously unseen scenarios [15], [16]. Also, GNN has recently shown success in various multi-robot applications such as formation control [17], [18], path finding [19], and task assignment [20]. Specifically, Tolstaya et al. implemented GNN to learn a decentralized flocking controller for a swarm of mobile robots by imitating a centralized flocking controller with global information. Similarly, Li et al. applied GNN to find collision-free paths for large networks of robots from start positions to goal positions in the obstacle-rich environments [19], [21]. Their results demonstrated that, by mimicking a centralized expert solution, their decentralized path planner exhibits a nearexpert performance, utilizing local observations and neighboring communication only, which also can be well generalized to larger teams of robots. The GNN-based approach was also investigated to learn solutions for the combinatorial

<sup>&</sup>lt;sup>1</sup>L. Zhou, A. Ribeiro, and V. Kumar are with the GRASP Laboratory, University of Pennsylvania, Philadelphia, PA, USA (email: {lfzhou, aribeiro, kumar}@seas.upenn.edu).

optimization in a multi-robot task scheduling scenario [20]. Contributions. To this end, we design a GNN-based learning framework that enables robots to communicate and share information with neighbors and selects actions for the robots to optimize target tracking performance. We train such a learning network to perform as close as possible to the greedy algorithm by imitating the behavior of the greedy algorithm. Different from classical algorithms (e.g., greedy algorithms), GNN can be seamlessly integrated with other networks such as convolutional neural network (CNN) or multilayer perceptron (MLP) to process richer data representations like images or sensor measurements. While, classical algorithms may not be able to either handle such data modalities or handle them efficiently. For example, (decentralized) greedy algorithms cannot directly use the raw image data or sensor measurements for decision making. Therefore, we devise a GNN-based learning framework that takes raw sensor measurements as inputs and leverages GNN for scalable and fast feature sharing to generate robots' actions. Specifically, we make the following contributions:

- We formulate the problem of applying GNN for multirobot target tracking with local communications (Problem 1);
- We design a GNN-based learning framework that processes robots' local observations and aggregates neighboring information to select actions for the robots (i.e., the solution to Problem 1).
- We demonstrate the performance of the GNN-based learning framework such as near-expert behavior, transferability, and fast running time in the scenario of active target tracking with large networks of robots. Specifically, our methods covers around 89% targets compared to the expert algorithm and runs several orders faster.

#### **II. PROBLEM FORMULATION**

We present the problem of *decentralized action selection* for multi-robot target tracking (see Fig. 1). Particularly, at each time step, the problem asks for selecting an action for each robot to optimize a target tracking objective using local information only. Specifically, the robots' actions are their candidate motion primitives, the team objective can be the number of targets covered, and for each robot, the local information is the set of targets covered by it and its neighbors. The goal of this work is to design GNN to learn such decentralized planning for the robots by *imitating* an expert algorithm. We start with introducing the framework of decentralized target tracking and GNN, and then formally define the problem.

## A. Decentralized Target Tracking

a) Robots: We consider a team of N robots, denoted by  $\mathcal{V} = \{1, 2, \dots, N\}$ . At a given time step, the relative position between any two robots i and j in the environment is denoted by  $\mathbf{p}_{ij}^r, i, j \in \mathcal{V}$ . The global positions of the robots are not required.



Fig. 1: Multi-robot target tracking: a team of aerial robots, mounted with down-facing cameras, aims at covering multiple targets (depicted as colorful dots) on the ground. The red arrow lines and the blue dotted lines show inter-robot observations and communications. The red squares represent the fields of view of the robots' cameras.



Fig. 2: Robots' observations, motion primitives, and communications: at a given time step, each robot observes the robots within its sensing range  $r_s$ ; e.g., robot 1 observes robot 2 and robot 3. Each robot has a set of motion primitives (depicted as dotted arrow curves), from which it can choose one to cover some targets using its camera. For example, robot 1 has 3 motion primitives,  $\{a_{1,1}, a_{1,2}, a_{1,3}\}$ , and it follows  $a_{1,1}$  to cover 2 targets,  $\{t_3, t_4\}$ . Robot 2 has 4 motion primitives,  $\{a_{2,1}, a_{2,2}, a_{2,3}, a_{2,4}\}$ , and it chooses  $a_{2,3}$  to cover 3 targets,  $\{t_1, t_2, t_3\}$ . However, in combination, these two motion primitives jointly cover 4 targets,  $\{t_1, t_2, t_3, t_4\}$ . In addition, each robot communicates with those robots within its communication range  $r_c$ ; e.g., robot 1 communicates with robot 2.

b) Action set: We denote a set of available actions for each robot i as  $\mathcal{A}_i, i \in \mathcal{V}$ . At a given time step, the robot can select at most 1 action from its available action set; e.g., in a motion planning scenario,  $\mathcal{A}_i$  is robot i's motion primitives, and the robot can select only 1 motion primitive to execute at each time step. For example, in Fig. 2, there are 2 robots, where robot 1's action set is  $\mathcal{A}_1 = \{a_{1,1}, a_{1,2}, a_{1,3}\}$  (and robot 1 selects action  $a_{1,1}$  to execute), and robot 2's action set is  $\mathcal{A}_2 = \{a_{2,1}, a_{2,2}, a_{2,3}, a_{2,4}\}$  (and robot 2 chooses  $a_{2,3}$ to execute). Denote the joint action set of all robots as  $\mathcal{A} \triangleq \bigcup_{i \in \mathcal{V}} \mathcal{A}_i$ . Also, denote a valid selection of actions for all robots as  $\mathcal{U} \subseteq \mathcal{A}$ . For example, we have  $\mathcal{U} = \{a_{1,1}, a_{2,3}\}$ for the two robots in Fig. 2.

c) Observation: We consider each robot i is equipped with a sensor (e.g., a LiDAR sensor) to measure the relative positions to the robots within its sensing range (see Fig. 2). Without loss of generality, we assume all robots' sensors

have the same sensing range  $r_s$ . For each robot i, we denote the set of robots within its sensing range as  $\mathcal{V}_i^s$ . Then the sensor observation of each robot i, i.e., the relative positions between robot i and  $\mathcal{V}_i^s$ , can be represented by  $\mathcal{Z}_i^r = \{\mathbf{p}_{ij}^r\}_{j \in \mathcal{V}_s^s}$ .

In addition, each robot is mounted with a camera that perceives a part of the environment within its field of view (see Fig. 1). Using the camera, the robot can observe some objects and measure the relative positions to them in the environment, once it selects an action to execute. For example, in Fig. 2, when robot 1 selects motion primitive  $a_{1,1}$ , it can sweep and cover a set of targets  $\{t_3, t_4\}$ , and the corresponding observation  $Z_{1,1}$  is the relative positions to targets  $\{t_3, t_4\}$ . Notably, each action of the robot corresponds to an observation. Thus, given a time step, we denote the (*possible*) camera observation<sup>1</sup> of each robot i by  $\mathcal{Z}_{i}^{t}$ , which is the collection of the observations by the robot's available actions  $A_i$ . For example, in Fig. 2, the camera observation of robot 1 is  $\mathcal{Z}_1^{t} = \{\mathcal{Z}_{1,1}, \mathcal{Z}_{1,2}, \mathcal{Z}_{1,3}\}$ . Particularly, we denote those objects that can be covered by robot i as  $\mathcal{T}_i$  and the corresponding relative positions as  $\{\mathbf{p}_{ij}^t\}_{j\in\mathcal{T}_i}$ . Then the robot *i*'s camera observation can be represented by  $\mathcal{Z}_i^t = {\mathbf{p}_{ij}^t}_{j \in \mathcal{T}_i}$ . Finally, we define the observation of each robot *i* by  $Z_i$ , which contains the observations of the sensor and camera on it, i.e.,  $\mathcal{Z}_i = \{\mathcal{Z}_i^r, \mathcal{Z}_i^t\}$ .

d) Communication: Each robot  $i \in \mathcal{V}$  communicates only with those robots within a prescribed communication range. Without loss of generality, we consider all robots have the same communication range  $r_c$  (see Fig. 2). That way, we introduce an (undirected) communication graph at a given time step as  $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \mathcal{W})$  with nodes the robots  $\mathcal{V}$ , edges  $\mathcal{E} \subseteq \mathcal{V} \times \mathcal{V}$  the communication links, and weights of the edges  $\mathcal{W} : \mathcal{E} \to \mathbb{R}$  denoting the strength of communications. The graph  $\mathcal{G}$  is distance-based and  $(i, j) \in \mathcal{E}$  if and only if  $\|\mathbf{p}_{ij}^r\|_2 \leq r_c$ . We denote the 1-hop neighbors of robot i by  $\mathcal{N}_i$ , which are the robots within the range  $r_c$ . We denote the adjacency matrix of graph  $\mathcal{G}$  by  $\mathbf{S} \in \mathbb{R}^{N \times N}$  with  $[\mathbf{S}]_{ij} = s_{ij} = 1$  if  $(i, j) \in \mathcal{E}$  and 0 otherwise. Notably, the connectivity of graph  $\mathcal{G}$  is *not* required.

e) Objective function: We consider a target tracking objective function  $f : 2^{\mathcal{A}} \to \mathbb{R}$  to be monotone nondecreasing and submodular in the robots' actions  $\mathcal{U}$ . For example, f can be the number of targets covered [22]. As shown in Fig. 2, the number of targets covered by the selected actions (motion primitives),  $\mathcal{U} = \{a_{1,1}, a_{2,3}\}$ , is  $f(\mathcal{U}) = 4$ .

## B. Graph Neural Networks

a) Graph Shift Operation: We consider each robot  $i, i \in \mathcal{V}$  has a feature vector  $\mathbf{x}_i \in \mathbb{R}^F$ , indicating the processed information of robot *i*. By collecting the feature vectors  $\mathbf{x}_i$  from all robots, we have the feature matrix for the robot team  $\mathcal{V}$  as:

$$\mathbf{X} = \begin{bmatrix} \mathbf{x}_{1}^{\mathsf{T}} \\ \vdots \\ \mathbf{x}_{N}^{\mathsf{T}} \end{bmatrix} = [\mathbf{x}^{1}, \cdots, \mathbf{x}^{F}] \in \mathbb{R}^{N \times F}, \quad (1)$$

<sup>1</sup>We call it as the possible observation, since the robot can select at most 1 action to execute at a time step.

where  $\mathbf{x}^f \in \mathbb{R}^N$ ,  $f \in [1, \dots, F]$  is the collection of the feature f across all robots  $\mathcal{V}$ ; i.e.,  $\mathbf{x}^f = [\mathbf{x}_1^f, \dots, \mathbf{x}_N^f]^\mathsf{T}$  with  $\mathbf{x}_i^f$  denoting the feature f of robot  $i, i \in \mathcal{V}$ . We conduct graph shift operation for each robot i by a linear combination of its neighboring features, i.e.,  $\sum_{j \in \mathcal{N}_i} \mathbf{x}_j$ . Hence, for all robots  $\mathcal{V}$  with graph  $\mathcal{G}$ , the feature matrix  $\mathbf{X}$  after the shift operation becomes  $\mathbf{SX}$  with:

$$[\mathbf{SX}]_{if} = \sum_{j=1}^{N} [\mathbf{S}]_{ij} [\mathbf{X}]_{j}^{f} = \sum_{j \in \mathcal{N}_{i}} s_{ij} \mathbf{x}_{j}^{f}, \qquad (2)$$

Here, the adjacency matrix S is called the *Graph Shift Operator* (GSO) [16].

b) Graph convolution: With the shift operation, we define the graph convolution by a linear combination of the shifted features on graph  $\mathcal{G}$  via K-hop communication exchanges [16], [19]:

$$\mathcal{H}(\mathbf{X};\mathbf{S}) = \sum_{k=0}^{K} \mathbf{S}^{k} \mathbf{X} \mathbf{H}_{k},$$
(3)

where  $\mathbf{H}_k \in \mathbb{R}^{F \times G}$  represents the coefficients combining F features of the robots in the shifted feature matrix  $\mathbf{S}^k \mathbf{X}$ , with F and G denoting the input and output dimensions of the graph convolution. Note that,  $\mathbf{S}^k \mathbf{X} = \mathbf{S}(\mathbf{S}^{k-1}\mathbf{X})$  is computed by means of k communication exchanges with 1-hop neighbors.

c) Graph neural network: Applying a point-wise nonlinearity  $\sigma : \mathbb{R} \to \mathbb{R}$  as the activation function to the graph convolution (eq. (3)), we define graph perception as:

$$\mathcal{H}(\mathbf{X};\mathbf{S}) = \sigma(\sum_{k=0}^{K} \mathbf{S}^{k} \mathbf{X} \mathbf{H}_{k}).$$
(4)

Then, we define a GNN module by cascading L layers of graph perceptions (eq. (4)):

$$\mathbf{X}^{\ell} = \sigma \left[ \mathcal{H}^{\ell}(\mathbf{X}^{\ell-1}; \mathbf{S}) \right] \quad \text{for} \quad \ell = 1, \cdots, L, \qquad (5)$$

where the output feature of the previous layer  $\ell - 1$ ,  $\mathbf{X}^{\ell-1} \in \mathbb{R}^{N \times F^{\ell-1}}$ , is taken as input to the current layer  $\ell$  to generate the output feature of layer l,  $\mathbf{X}^{\ell}$ . Recall that the input to the first layer is  $\mathbf{X}^0 = \mathbf{X}$  (eq. (1)). The output feature of the last layer  $\mathbf{X}^L \in \mathbb{R}^{N \times G}$ , obtained via *K*-hop communications and multi-layer perceptions, will be used to predict action set  $\mathcal{U}$  for all robots to the following problem.

Notably, GNN can represent rich classes of mappings (functions) from the input feature to the output feature. Since the graph convolution is a linear operator (eq. (3)), the function's characteristics mainly depends on the property of the activation function  $\sigma$  (in eq. (5)). For example, if the activation function is concave (or submodular), the function represented by GNN is also concave (or submodular).

## C. Problem Definition

**Problem** (Decentralized action selection for multi-robot target tracking). At each time step, the robots V, by exchanging information with neighbors only over the communication graph  $\mathcal{G}$ , select an action to each robot  $i \in \mathcal{V}$  to maximize a submodular target tracking objective function f:

$$\max_{\mathcal{U}\subseteq\mathcal{A}} f(\mathcal{U})$$
  
s.t.  $|\mathcal{U}\cap\mathcal{A}_i| = 1$ , for all  $i \in \mathcal{V}$ . (6)

The constraint follows a partition matroid constraint to ensure that each robot selects 1 action per time step (e.g., 1 motion primitive among a set of motion primitives).

Eq. (6) can be interpreted as a submodular maximization problem with a partition matroid constraint and decentralized communication. This problem is generally NP-hard even if we assume the centralized communication [8]. That is because, finding the *optimal* action set requires to exhaustive search and evaluate the quality of all possible valid action sets  $\mathcal{U} \subset \mathcal{A}$ . Clearly, this exhaustive search method has combinatorial complexity and quickly becomes intractable as either the number of robots or the number of available actions of the robots increase. The most well-known approach for tackling this type of problem (with the centralized communication) is the (centralized) greedy algorithm [8]. The advantage of the centralized greedy algorithm is two-fold: (i), it is efficient as it runs in polynomial time (with  $O(|\mathcal{A}|^2)$ ) complexity); (ii), it achieves at least 1/2-approximation of the optimal. In addition, since the 1/2-approximation bound is computed based on the worst-case performance, the centralized greedy algorithm can typically perform much better (on average) in practice; e.g., it performs close to the exhaustive search (see Figure 4-c in Section IV-D). Since the centralized greedy algorithm is much cheaper than the exhaustive search and performs comparatively, we use the centralized greedy algorithm as our *expert* algorithm to generate ground-truth training data.

Particularly, we aim to train a learning network to perform as well as the centralized greedy algorithm for solving Eq. (6), while involving the communications among neighboring robots only. For such a learning network, GNN can be a great fit given its decentralized communication protocol. Hence, the goal is to utilize GNN to learn an action set  $\mathcal{U}$  to Eq. (6) by imitating the action set selected by the centralized greedy algorithm. More formally, we define the problem of this work as follows.

**Problem 1.** Design a GNN-based learning framework to learn a mapping  $\mathcal{M}$  from the robots' observations  $\{\mathcal{Z}_i\}_{i \in \mathcal{V}}$ and the communication graph  $\mathcal{G}$  to the robots' action set  $\mathcal{U}$ , i.e.,  $\mathcal{U} = \mathcal{M}(\{\mathcal{Z}_i\}_{i \in \mathcal{V}}, \mathcal{G})$ , such that  $\mathcal{U}$  is as close as possible to action set selected by the centralized greedy algorithm, denoted by  $\mathcal{U}^{g}$ .

We describe in detail our learning architecture for solving Problem 1 in the next section.

# **III. ARCHITECTURE**

We design a learning architecture that consists of three main components—the Individual Observation Processing (Section III-A), the Decentralized Information Aggregation (Section III-B), an the Decentralized Action Selection module (Section III-C). Next, we describe in detail these key components, illustrated in Fig. 3, as follows.

#### A. Individual Observation Processing

Recall that each robot *i*'s observation  $Z_i$  includes its sensor observation  $Z_i^r$  and its camera observation  $Z_i^t$  (see Section II-A). We process the observation  $Z_i$  to generate a feature vector  $\mathbf{x}_i$  for each robot *i* in two steps.

- Step 1 is a pre-processing step. Since the sensor observation  $Z_i^r$  stores the relative robot positions  $\{\mathbf{p}_{ij}^r\}_{j \in \mathcal{V}_i^s}$ , we reshape (handcraft) it as a vector  $\mathbf{x}_{i,1}^- := [\mathbf{p}_{i1}^T, \cdots, \mathbf{p}_{i|\mathcal{V}_i^s|}^T]^T$ . Similarly, for the camera observation  $Z_i^t$  that contains the relative positions to the objects covered,  $\{\mathbf{p}_{ij}^t\}_{j \in \mathcal{T}_i}$ , we reshape it as vector  $\mathbf{x}_{i,2}^- := [\mathbf{p}_{i1}^T, \cdots, \mathbf{p}_{i|\mathcal{T}_i|}^T]^T$ . Finally, by concatenating  $\mathbf{x}_{i,1}^-$  and  $\mathbf{x}_{i,2}^-$ , we generate a pre-processed feature vector  $\mathbf{x}_i^-$  for robot *i*, i.e.,  $\mathbf{x}_i^- = [(\mathbf{x}_{i,1}^-)^T, (\mathbf{x}_{i,2}^-)^T]^T$ .<sup>2</sup>
- Step 2 is a post-processing step where the pre-processed feature vector x<sub>i</sub><sup>-</sup> is fed into a multi-layer perceptron (MLP) module to generate the robot's feature vector x<sub>i</sub>, i.e., x<sub>i</sub> = MLP(x<sub>i</sub><sup>-</sup>).

The feature vectors of the robots are then exchanged and fused through neighboring communications (Section III-B).

# B. Decentralized Information Aggregation

Each robot *i* communicates its feature (or processed information) with its neighbors  $\mathcal{N}_i$  over multiple communication hops. As shown in Section II-B, for each robot *i*, we use GNN to aggregate and fuse the feature vectors through *K*-hop communication exchanges among neighbors (eq. (3)). Thus, the output of the GNN (i.e.,  $\mathbf{X}^L$  in eq. (5)) is a hyperrepresentation of the fused information of the robots and their *K*-hop neighbors. The output is then taken as input to the action selection module, described in Section III-C, to generate an action for each robot. Notably, since only neighboring information is exchanged and fused, GNN renders a decentralized decision-making architecture.

#### C. Decentralized Action Selection

We aim at selecting an action set  $\mathcal{U}$  for the robots  $\mathcal{V}$  (1 action per robot) to maximize the team tracking performance  $f(\mathcal{U})$ . To this end, we use a MLP for each robot *i* to train an action selection module. More specifically, each robot applies a MLP that takes the aggregated features as input and selects an action for the robot as output. We consider all robots carry the same MLP, resembling a weight-sharing scheme. The actions of the robots are selected based on a supervised learning approach, as in Section III-D.

<sup>&</sup>lt;sup>2</sup>We assign the same dimension to feature vector  $\mathbf{x}_i^-$  for all the robots. In particular, for each robot *i*, we use relative positions of 10 nearest robots in  $\{\mathbf{p}_{ij}^r\}_{j \in \mathcal{V}_i^s}$  and relative positions of 20 nearest targets in  $\{\mathbf{p}_{ij}^t\}_{j \in \mathcal{T}_i}$ . If robot *i* has less than 10 nearby robots measured or covers less than 20 targets, we simply add some dummy values (e.g., -1) in the feature vector to maintain the same dimension.



Fig. 3: Our decentralized learning architecture includes three main modules: (i), the Individual Observation Processing module processes the local observations and generates a feature vector for each robot; (ii), the Decentralized Information Aggregation module utilizes a GNN to aggregate and fuse the features from K-hop neighbors for each robot; (iii), the Decentralized Action Selection module selects an action for each robot by imitating an expert algorithm.

## D. Supervised Learning

We train our learning architecture by a supervised learning approach, i.e., to mimic an expert algorithm (the centralized greedy algorithm). Specifically, during the training stage, we have access to the action set  $\mathcal{U}^{g}$  selected by the centralized greedy algorithm, the corresponding observations on the robots  $\{\mathcal{Z}_i\}_{i\in\mathcal{V}}$ , and the corresponding communication graph  $\mathcal{G}$ . Thus, the training set  $\mathcal{D}$  can be constructed as a collection of these data, i.e.,  $\mathcal{D} := \{(\{\mathcal{Z}_i\}_{i\in\mathcal{V}}, \mathcal{G}, \mathcal{U}^{g})\}$ . Over the training set  $\mathcal{D}$ , we train the mapping  $\mathcal{M}$  (defined in Problem 1) so that a cross entropy loss  $\mathcal{L}(\cdot, \cdot)$ , representing the difference between the output action set  $\mathcal{U}$  and the greedy action set  $\mathcal{U}^{g}$  is minimized. That is,

$$\min_{\text{PM, }\{\mathbf{H}_{\ell,k}\}, \text{ MLP }} \sum_{(\{\mathcal{Z}_i\}_{i\in\mathcal{V}}, \mathcal{G}, \mathcal{U}^g)\in\mathcal{D}} \mathcal{L}(\mathcal{M}(\{\mathcal{Z}_i\}_{i\in\mathcal{V}}, \mathcal{G}), \mathcal{U}^g),$$
(7)

where we optimize over the learnable parameters of the processing method, named PM (e.g., CNN, MLP) to process the robots' observations, the set of learnable parameter matrices  $\{\mathbf{H}_{\ell,k}\}_{l \in \{1, \dots, L\}, k \in \{0, \dots, K\}}$  in GNN to aggregate and fuse the neighboring information over multiple communication hops, and the learnable parameters of MLP to select actions for the robots.

Notably, by the decentralized learning architecture where the parametrization is operated locally on each robot, the number of learnable parameters is independent of the number of the robots N. This decentralized parametrization offers a perfect complement to the supervised learning which requires the availability of expert solutions that can be costly in the large-scale settings. For example, even though the centralized greedy algorithm is much more efficient than the exhaustive search, it still takes considerable time to generate a solution when the number of robots or their actions is large, given its running time grows quadratically in the number of robots' actions (i.e., with  $O(|\mathcal{A}|^2)$  complexity). However, leveraging the decentralized parametrization, we only need to train over the small-scale cases and generalize the trained models such as the distributed computation (i.e., PM and MLP) and the neighboring information exchange (i.e., GNN) to the largerscale settings, as it will be demonstrated in Section IV-D. In other words, once trained, the learned models can be implemented in other cases, including those with different communication graphs and varying numbers of robots.

# **IV. PERFORMANCE EVALUATION**

We present the evaluations of our method in scenarios of *active target tracking with large networks of robots*. In particular, we compare our method with other baseline algorithms in terms of both running time and tracking quality. In these comparisons, we test the trained learning models in the cases with various team sizes. The code of implementations is available online.<sup>3</sup> The experiments are conducted using a 32-core, 2.10Ghz Xeon Silver-4208 CPU and an Nvidia GeForce RTX 2080Ti GPU with 156GB and 11GB of memory, respectively.

Next, we first describe the active target tracking scenario, the specifications of the learning architecture, and the compared baseline algorithms. Then we present the evaluations.

# A. Multi-Robot Active Target Tracking

We consider N aerial robots that are tasked to track M mobile targets on the ground. Each robot uses its sensor to obtain the relative positions to those robots that are within its sensing range  $r_s$ . The sensing range  $r_s$  is set to be  $r_s = 20$  units. The camera on each robot i has a square field of view  $d_o \times d_o$ , and each robot i has 5 available motion primitives,  $\mathcal{A}_i = \{\text{forward, backward, left, right, idle}\}^4$ . Once the robot selects a motion primitive. If the robot selects the idle motion primitive, it stays still (i.e.,  $d_m = 0$ ). Hence, each motion

<sup>3</sup>https://github.com/VishnuDuttSharma/ deep-multirobot-task

<sup>&</sup>lt;sup>4</sup>The designed learning architecture is generic and can be extended to handle other action spaces (e.g., robots have different action sets as shown in Figure 2) or robot trajectories as long as the centralized expert algorithm it learns from operates with the corresponding action spaces or robot trajectories.

primitive corresponds to a rectangular tracking region with length  $d_m + d_o$  and width  $d_o$ . The tracking width (or the side of the field of view) is set to be  $d_o = 6$  units. The flying length  $d_m$  is set to be  $d_m = 20$  units for all robots selecting the non-idle motion primitive. We set the robot's communication range as  $r_c = 10$  units and the communication hop as K = 1 (i.e., each robot *i* communicates only with its 1-hop neighbors  $\mathcal{N}_i$ ). We let robots fly at different heights so that collisions do not occur during their movement. The objective function is considered to be the number of targets covered, given all robots selecting motion primitives.

## B. Supervised Learning Specification

We apply the centralized greedy algorithm [8] as the expert algorithm to generate a ground-truth data set. In each problem instance, the size of the environment, a square, depends on the number of robots, i.e., its side length is specified as  $100 \times \sqrt{N/20}$ . We randomly generate the positions of the robots and targets in the environment, and utilize the centralized greedy algorithm to select an action set for the robots. Notably, each instance includes the robots' observations  $\{Z_i\}_{i \in \mathcal{V}}$ , the communication graph  $\mathcal{G}$  (represented by its adjacency matrix S), and the greedy (or ground-truth) action set  $\mathcal{U}^{g}$ . The ground-truth data set comprises 120,000 instances for varying numbers of robots and the corresponding environments. In particular, we scale the size of the environment proportionally to the number of robots but keep the target density the same. Here, the target density is captured by the percentage of the cells in the grid occupied by the targets. We set the target density as 2.5% in all cases. The robot's observations include the relative positions of 10 nearest robots within its sensing range and 20 nearest targets that can be covered. We experimentally found these hyperparameters such as the target density and the numbers of nearest robots and targets are suitable across different environmental scales. The data is randomly shuffled at training time and divided into a training set (60%), a validation set (20%), and a testing set (20%).

Our learning architecture consists of a 3-layer MLP with 32, 16, and 8 hidden layers as the Individual Observation Processing module, a 2-layer GNN with 32 and 128 hidden layers as the Decentralized Information Aggregation module, and a single layer MLP as the Decentralized Action Selection module. For each robot, the network outputs a probability of selecting each action in the action set. All the robots have identical action sets. This learning network is implemented in PyTorch v1.6.0 and accelerated with CUDA v10.1 APIs. We use a learning rate scheduler with cosine annealing to decay the learning rate from  $5 \times 10^{-3}$  to  $10^{-6}$  over 1500 epochs with batch size 64. This architecture and training parameters are selected from multiple parameter search experiments.

# C. Compared Algorithms

We compare our method, named GNN with three other algorithms. The algorithms differ in how they select the robots' motion primitives. The first algorithm is the centralized greedy algorithm (the expert algorithm), named Centrl-gre. The second algorithm is an optimal algorithm, named Opt which attains the optimal solution for eq. (6) by exhaustive search. Particularly, for N robots, each with 5 motion primitives, Opt needs to evaluate  $5^N$  possible cases to find the optimal solution. Evidently, Opt is viable only for small-scale cases. Hence, Opt is used for comparison only when the number of robots is small (e.g., N < 10). The third algorithm is a random algorithm, named Rand, which randomly (uniformly) selects one motion primitive for each robot. The fourth algorithm is a decentralized greedy algorithm, named Decent-gre, which applies the standard greedy algorithm [7] to select an action with the maximal marginal gain for each robot among the robot and its 1-hop neighbors<sup>5</sup>. With the same settings, these four algorithms are compared in terms of the team's tracking quality, i.e., the number of targets covered by all robots, and the running time, across 1000 trials. Notably, the running time for the Decent-gre is calculated as the maximum time taken by any robot. For GNN, it is the inference time on GPU.

## D. Evaluations

a) Small-scale comparison: We first compare GNN with Opt, Centrl-gre, Decent-gre and Rand in small-scale cases, as shown in Figures 4-(a) & (c). Since Opt is only feasible for small-scale scenarios, we set the number of robots N = 4, 6, 8, 10. Particularly, we train GNN with the number of robots N = 20 and test its performance over N = 4, 6, 8, 10.

We observe GNN has a superior running time as shown in Figure 4-(a): it runs considerably faster than both Centrl-gre and Opt: around 0.5 order faster than the former and 1.5 orders faster than the latter with 4 robots. It also runs slightly faster than Decent-gre<sup>6</sup>. This superiority becomes more significant as the number of robot increases. In addition, GNN has an average running time less than 1 ms, regardless of the number of robots, which is due to its decentralized decision-making protocol. Despite the faster running time, GNN retains a tracking performance close to Centrl-gre and Opt, slightly better than Decent-gre, and better than Rand: it covers in average more than 89% of the number of targets covered by Centrl-gre and more than 85% of that by Opt, and covers more targets than both Decent-gre and Rand, as shown in Figure 4-(c)<sup>7</sup>.

Figures 4-(a) & (c) also demonstrate the generalization capability of GNN in smaller-scale scenarios: even though it is trained with 20 robots, it maintains both fast running time and the tracking performance close to Centrl-gre and Opt with smaller number of robots, e.g., N = 4, 6, 8, 10. Another interesting observation is that Centrl-gre covers

<sup>&</sup>lt;sup>5</sup>Notably, Decent-gre is different from the distributed (sequentially) greedy algorithms in [11], [12] where each robot selects an action based on the actions of its previous neighbors (i.e., its neighbors that have already selected actions). This is to ensure Decent-gre and GNN are compared with the same communication setting.

 $<sup>^6</sup> The averaged running times of Opt, Centrl-gre, Decent-gre, and GNN with 4 robots are 18ms, 2.14ms, 0.77ms, and 0.72ms, respectively.$ 

<sup>&</sup>lt;sup>7</sup>In the small-scale experiment, the averaged number of the targets covered for Opt, Centrl-gre, Decent-gre, GNN, and Rand are 49.12, 47.21, 42, 42.67, and 34.23, respectively.



Fig. 4: Comparison of Opt, Centrl-gre, Decent-gre, GNN, and Rand in terms of running time (plotted in log scale) and the number of targets covered. (a) & (d) are for small-scale comparison averaged across 1000 Monte Carlo trials. (b) & (e) are for large-scale comparison averaged across 1000 Monte Carlo trials.

the similar number of targets as Opt, and yet runs several orders faster when  $N \ge 8$ . This demonstrates the rationality of choosing Centrl-gre as the expert algorithm in this target tracking scenario.

b) Large-scale comparison: We compare GNN with Centrl-gre, Decent-gre, and Rand in the large-scale scenarios where the number of robots is set as N =20, 30, 40, 50, 100. Opt is not included in this comparison due to its long evaluation time. For example, it takes almost two days to evaluate 1000 instances, each with  $5^{10}$  possible cases, for 10 robots. GNN is trained with the number of robots N = 20 and is tested over N = 20, 30, 40, 50, 100. The results are reported in Figures 4-(b) & (d). Similarly, we observe GNN runs around 1.5 to 2.5 orders faster than Centrl-gre with the running time less than 1 ms for all N = 20, 30, 40, 50, 100 (Figures 4-(b)). It also runs faster than Decent-gre, which runs slower as the number of robots increases. Although GNN runs faster, it achieves a tracking performance close to Centrl-gre, slightly better than Decent-gre, and better than Rand (Figure 4-(d)). Additionally, these results verify GNN's generalization capability in larger-scale scenarios: it is trained with 20 robots, and yet, can be well generalized to a larger number of robots, e.g., N = 30, 40, 50, 100.

Test Train	20	30	40	50
20	89.34%	88.98%	88.56%	88.61%
30	89.45%	88.93%	88.70%	88.66%
40	89.33%	88.78%	88.54%	88.62%
50	89.38%	88.87%	88.57%	88.72%

TABLE I: Percentage of the number of targets covered (the average across 1000 trials) by GNN trained and tested with varying numbers of robots.

c) Generalization evaluation: We further verify GNN's generalization capability by evaluating the tracking quality of GNN trained and tested with varying numbers of robots. Specifically, we train GNN with N = 20, 30, 40, 50 robots and test it on N = 20, 30, 40, 50 robots. The evaluation results are reported in Table I where the tracking quality is captured by the percentage of the number of targets covered with respect to the number of the targets covered by Centrl-gre. We observe that GNN trained and tested with the same and different number of robots cover the similar percentage of the targets (around 89%), which further demonstrates the generalization capability of GNN.

To summarize, in the evaluations above, GNN provides a significant computational speed-up, and, yet, still attains a target tracking quality that nearly matches that of Centrl-gre and Opt. Moreover, GNN achieves a slightly better tracking quality than Decent-gre but runs faster (especially with more than 20 robots) and thus scales better. In addition, GNN has a better tracking quality than Rand. Further, GNN exhibits the capability of being able to well generalize to previously unseen scenarios. Particularly, it can be trained in a smaller-scale environment, which typically has a cheaper computational overhead. Then the trained policies can apply to larger-scale environments.

## V. CONCLUSION AND FUTURE WORK

We worked towards choosing actions for the robots with local communications to maximize a team's tracking quality. We devised a supervised learning approach that selects actions for the robots by imitating an expert solution. Particularly, we designed a GNN-based learning network that maps the robots' individual observations and inter-robot communications to the robots' actions. We demonstrated the near-expert performance, generalization capability, and fast running time of the proposed approach.

This work opens up a number of future research avenues. An ongoing work is to extend the designed GNN-based learning architecture in a distributed fashion and analyze the number of messages shared and communication costs. In addition, we will incorporate the attention mechanism [23] that enables robots to learn when to communicate [24], who to communicate with [25], and what to communicate [21] to prioritize the information with higher contributions and reduce communication costs. A second research avenue is to learn resilient coordination that secures team performance against either the malicious team members [26] or adversarial outsiders [27], [28]. A third research direction is to explore decentralized reinforcement learning methods [29], [30] for multi-robot target tracking. Further, our learningbased approach can be extended to data-based learning for more complicated tasks which do not have well-performing, decentralized solutions yet [31]. The model could also be used to bootstrap the GNN-based planner for such tasks, which cannot be done using classical counterparts [31].

#### REFERENCES

- P. Tokekar, E. Branson, J. Vander Hook, and V. Isler, "Tracking aquatic invaders: Autonomous robots for monitoring invasive fish," *IEEE Robotics & Automation Magazine*, vol. 20, no. 3, pp. 33–41, 2013.
- [2] U. Zengin and A. Dogan, "Real-time target tracking for autonomous uavs in adversarial environments: A gradient search algorithm," *IEEE Transactions on Robotics*, vol. 23, no. 2, pp. 294–307, 2007.
- [3] B. Grocholsky, J. Keller, V. Kumar, and G. Pappas, "Cooperative air and ground surveillance," *IEEE Robotics & Automation Magazine*, vol. 13, no. 3, pp. 16–25, 2006.
- [4] V. Kumar and N. Michael, "Opportunities and challenges with autonomous micro aerial vehicles," in *Robotics Research*. Springer, 2017, pp. 41–58.
- [5] A. Krause, A. Singh, and C. Guestrin, "Near-optimal sensor placements in gaussian processes: Theory, efficient algorithms and empirical studies," J. Mach. Learn. Res., vol. 9, no. Feb, pp. 235–284, 2008.
- [6] H. Ding and D. Castanón, "Multi-agent discrete search with limited visibility," in *Decision and Control (CDC)*, 2017 IEEE 56th Annual Conference on. IEEE, 2017, pp. 108–113.

- [7] G. L. Nemhauser, L. A. Wolsey, and M. L. Fisher, "An analysis of approximations for maximizing submodular set functions-i," *Mathematical programming*, vol. 14, no. 1, pp. 265–294, 1978.
- [8] M. L. Fisher, G. L. Nemhauser, and L. A. Wolsey, "An analysis of approximations for maximizing submodular set functions—ii," in *Polyhedral combinatorics*. Springer, 1978, pp. 73–87.
- [9] N. Atanasov, J. Le Ny, K. Daniilidis, and G. J. Pappas, "Decentralized active information acquisition: Theory and application to multi-robot slam," in 2015 IEEE International Conference on Robotics and Automation (ICRA). IEEE, 2015, pp. 4775–4782.
- [10] R. K. Williams, A. Gasparri, and G. Ulivi, "Decentralized matroid optimization for topology constraints in multi-robot allocation problems," in 2017 IEEE International Conference on Robotics and Automation (ICRA). IEEE, 2017, pp. 293–300.
- [11] B. Gharesifard and S. L. Smith, "Distributed submodular maximization with limited information," *IEEE Transactions on Control of Network Systems*, vol. 5, no. 4, pp. 1635–1645, 2018.
- [12] D. Grimsman, M. S. Ali, J. P. Hespanha, and J. R. Marden, "The impact of information in greedy submodular maximization," *IEEE Transactions on Control of Network Systems*, 2018.
- [13] G. Qu, D. Brown, and N. Li, "Distributed greedy algorithm for multiagent task assignment problem with submodular utility functions," *Automatica*, vol. 105, pp. 206–215, 2019.
- [14] S. Choudhury, A. Kapoor, G. Ranade, S. Scherer, and D. Dey, "Adaptive information gathering via imitation learning," *arXiv preprint* arXiv:1705.07834, 2017.
- [15] L. Ruiz, F. Gama, and A. Ribeiro, "Graph neural networks: Architectures, stability, and transferability," *Proceedings of the IEEE*, 2021.
- [16] F. Gama, A. G. Marques, G. Leus, and A. Ribeiro, "Convolutional neural network architectures for signals supported on graphs," *IEEE Trans. on Signal Processing*, vol. 67, no. 4, pp. 1034–1049, 2019.
- [17] E. Tolstaya, F. Gama, J. Paulos, G. Pappas, V. Kumar, and A. Ribeiro, "Learning decentralized controllers for robot swarms with graph neural networks," in *Conference Robot Learning 2019*. Osaka, Japan: Int. Found. Robotics Res., 30 Oct.-1 Nov. 2019.
- [18] A. Khan, E. Tolstaya, A. Ribeiro, and V. Kumar, "Graph policy gradients for large scale robot control," in *Conference on Robot Learning*, 2020, pp. 823–834.
- [19] Q. Li, F. Gama, A. Ribeiro, and A. Prorok, "Graph neural networks for decentralized multi-robot path planning," in 2020 IEEE/RSJ Intl Conference on Intelligent Robots and Systems (IROS). IEEE, 2020.
- [20] Z. Wang and M. Gombolay, "Learning scheduling policies for multi-robot coordination with graph attention networks," *IEEE Robotics and Automation Letters*, vol. 5, no. 3, pp. 4509–4516, 2020.
  [21] Q. Li, W. Lin, Z. Liu, and A. Prorok, "Message-aware graph attention
- [21] Q. Li, W. Lin, Z. Liu, and A. Prorok, "Message-aware graph attention networks for large-scale multi-robot path planning," *IEEE Robotics* and Automation Letters, 2021.
- [22] P. Tokekar, V. Isler, and A. Franchi, "Multi-target visual tracking with aerial robots," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2014, pp. 3067–3072.
- [23] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, "Attention is all you need," arXiv preprint arXiv:1706.03762, 2017.
- [24] Y.-C. Liu, J. Tian, N. Glaser, and Z. Kira, "When2com: Multi-agent perception via communication graph grouping," in *Proceedings of the IEEE/CVF Conference on computer vision and pattern recognition*, 2020, pp. 4106–4115.
- [25] Y.-C. Liu, J. Tian, C.-Y. Ma, N. Glaser, C.-W. Kuo, and Z. Kira, "Who2com: Collaborative perception via learnable handshake communication," in 2020 IEEE International Conference on Robotics and Automation (ICRA). IEEE, 2020, pp. 6876–6883.
- [26] J. Blumenkamp and A. Prorok, "The emergence of adversarial communication in multi-agent reinforcement learning," arXiv preprint arXiv:2008.02616, 2020.
- [27] L. Zhou, V. Tzoumas, G. J. Pappas, and P. Tokekar, "Resilient active target tracking with multiple robots," *IEEE Robotics and Automation Letters*, vol. 4, no. 1, pp. 129–136, 2018.
- [28] —, "Distributed attack-robust submodular maximization for multirobot planning," in 2020 IEEE International Conference on Robotics and Automation (ICRA). IEEE, 2020, pp. 2479–2485.
- [29] K. Zhang, Z. Yang, H. Liu, T. Zhang, and T. Basar, "Fully decentralized multi-agent reinforcement learning with networked agents," in *Intl Conference on Machine Learning*. PMLR, 2018, pp. 5872–5881.
- [30] S. Omidshafiei, J. Pazis, C. Amato, J. P. How, and J. Vian, "Deep decentralized multi-task multi-agent reinforcement learning under partial observability," in *International Conference on Machine Learning*. PMLR, 2017, pp. 2681–2690.

[31] V. D. Sharma, L. Zhou, and P. Tokekar, "D2coplan: A differentiable decentralized planner for multi-robot coverage," *arXiv preprint arXiv:2209.09292*, 2022.