

Learning Scalable Policies over Graphs for Multi-Robot Task Allocation using Capsule Attention Networks

Steve Paul¹, Payam Ghassemi², and Souma Chowdhury^{3,†}

Abstract—This paper presents a novel graph reinforcement learning (RL) architecture to solve multi-robot task allocation (MRTA) problems that involve tasks with deadlines and workload, and robot constraints such as work capacity. While drawing motivation from recent graph learning methods that learn to solve combinatorial optimization (CO) problems such as multi-Traveling Salesman and Vehicle Routing Problems using RL, this paper seeks to provide better performance (compared to non-learning methods) and important scalability (compared to existing learning architectures) for the stated class of MRTA problems. The proposed neural architecture, called Capsule Attention-based Mechanism or CapAM acts as the policy network, and includes three main components: 1) an encoder: a Capsule Network based node embedding model to represent each task as a learnable feature vector; 2) a decoder: an attention-based model to facilitate a sequential output; and 3) context: that encodes the states of the mission and the robots. To train the CapAM model, the policy-gradient method based on REINFORCE is used. When evaluated over unseen scenarios, CapAM demonstrates better task completion performance and >10 times faster decision-making compared to standard non-learning based online MRTA methods. CapAM’s advantage in generalizability, and scalability to test problems of size larger than those used in training, are also successfully demonstrated in comparison to a popular approach for learning to solve CO problems, namely the purely attention mechanism.

I. INTRODUCTION

Multi-Robot Task Allocation (MRTA) involves coordinating a set of tasks among a team of cooperative robotic systems such that the decisions are free of conflict and optimize a quantity of interest [1]. The potential real-world applications of such MRTA problems are immense considering that multi-robotics is one of key emerging directions of robotics research and development, and task allocation is fundamental to most multi-robotic or swarm-robotic operations. Examples include disaster response [2], last-mile delivery [3], environment monitoring [4], and reconnaissance [5]. We specifically focus on a class of MRTA problems that falls into the Single-task Robots, Single-robot Tasks, and Time-extended Assignment (SR-ST-TA) class defined in [1], [6], that include characteristics such as tasks with time deadlines and workload, and robots with constrained work capacity. Based on iTax taxonomy as defined in [1], these

problems fall into the In-schedule Dependencies (ID) category. The MRTA problems considered here involves *mildly-heterogeneous* robots (with different task completion rates), and where the task selections are assumed to take place in a decentralized manner, under full observability of the task environment and peers’ decisions.

There exists a notable body of work (mostly non-learning approaches) in solving MRTA problems related to these applications, e.g., graph-based methods [2], integer-linear programming (ILP) approaches [7], [8], and auction-based methods [9]. However, critical challenges continue to persist, especially regarding the scalability with number of robots and tasks, and the ability to adapt to complex problem characteristics without tedious hand-crafting of underlying heuristics. Existing solutions to multi-Traveling Salesman Problem (mTSP) and Vehicle Routing Problem (VRP) in the literature [10], [11] have addressed analogical problem characteristics of interest to MRTA, albeit in a disparate manner; these characteristics include tasks with time deadlines, and multiple tours per vehicle, with applications in the operations research and logistics communities [12], [13]. Integer linear programming or ILP based mTSP-type formulations and solution methods have also been extended to task allocation problems in the multi-robotics domain [14]. Although the ILP-based approaches can in theory provide optimal solutions, they are characterized by exploding computational effort as the number of robots and tasks increases [8], [15]. For example, for the studied SR-ST problem, the cost of solving the exact ILP formulation of the problem, even with a linear cost function (thus an ILP), scales with $\mathcal{O}(n^3m^2h^2)$, where n , m , and h represent the number of tasks, the number of robots, and the maximum number of tours per robot, respectively [16]. As a result, most online MRTA methods, e.g., auction-based methods [9], [17], [18], metaheuristic methods [19], [20], and bi-graph matching methods [16], [21], [22], use some sort of heuristics, and often report the optimality gap at least for smaller test cases compared to the exact ILP solutions.

In recent years, learning approaches based on Graph Neural Networks or GNN (which can model non-Euclidean data) are being increasingly used to solve planning problems with a Combinatorial Optimization (CO) formulation, e.g., TSP, VRP, Max-Cut, Min-Vertex, and MRTA [23]–[33]. These existing studies are however limited in three key aspects: **1)** They address simplified problems that often exclude common real-world factors such as resource and capacity constraints [23], [25], [26], [31]). **2)** They are mostly focused on smaller sized problems (≤ 100 tasks) [30], [32], [34], [35], with

Authors ^{1,2,3} are with the department of Mechanical and Aerospace Engineering, University at Buffalo, Buffalo, NY, USA {stevepau, payamgha, soumacho}@buffalo.edu

[†] Corresponding Author, soumacho@buffalo.edu

This work was supported by the Office of Naval Research (ONR) award N00014-21-1-2530 and National Science Foundation (NSF) award CMMI 2048020. Any opinions, findings, conclusions, or recommendations expressed in this paper are those of the authors and do not necessarily reflect the views of the ONR or the NSF.

their scalability remaining unclear. **3)** They rarely provide evidence of generalizing to problem scenarios that are larger in size than those used for training. This latter capability is particularly critical since real-world MRTA problems often involve simulating episodes whose costs scale with the number of tasks and robots, making re-training efforts burdensome. Moreover the time-sensitive nature of many real-world applications of MRTA [36], such as in disaster response [22], call for a near real-time performance, while involving a large-sized decision space [37]. Therefore, near real-time computable good feasible solutions are of essence.

Along with the need to address scalability, we posit that a sequential decision-making approach is favored considering environmental and communication uncertainties and the possibility of new tasks to emerge during the operation [38], [39]. Therefore, we formulate the MRTA problem such that the decision taken by the robots on visiting the next node is performed in a sequential, decentralized, and asynchronous manner. To enable scalable policies, we design a novel **encoder-decoder** policy network, where the **encoder** is based on Graph Capsule Convolutional Neural Networks (GCAPCN) [40], which is hypothesized to incorporate local and global structural information with permutation invariance. The **decoder** is based on a Multi-head Attention mechanism (MHA) [23], [41] which fuses the encoded information and problem-specific **context** using matrix multiplication, in order to enable decentralized sequential decisions. The proposed network architecture is named Capsule Attention Mechanism or CapAM.

Contributions: The main contributions of this paper can be summarized as **1)** Formulate the general SR-ST class of MRTA problems as a Markov Decision Process or MDP over graphs with the multi-robot system’s state embedded as the *context* portion of the policy model, such that the task allocation policy can be learned using an RL approach. **2)** Explore how a policy network that learns the structural information of task-graphs results in better decision-making. **3)** Demonstrate this learning framework’s ability to generalize to larger-sized problems without the need to retrain.

Paper outline: The next section provides a brief overview of MRTA and its MDP formulation. Section III then describes our proposed new graph learning architecture that operates over this MDP. Section IV presents the settings and outcomes of numerical experiments on a large number of task allocation problems with time and capacity constraints [19], which are used to evaluate this new approach to MRTA. This evaluation study includes comparative analysis w.r.t. non-learning methods as well as a well-known attention mechanism based reinforcement learning method, and further parametric and scalability analysis. Finally, Section V provides concluding remarks and discusses the potential future extensions for our new graph learning-based approach to MRTA. Throughout this paper, **tasks** are also referred by the terms “**nodes**” or “**vertices**”.

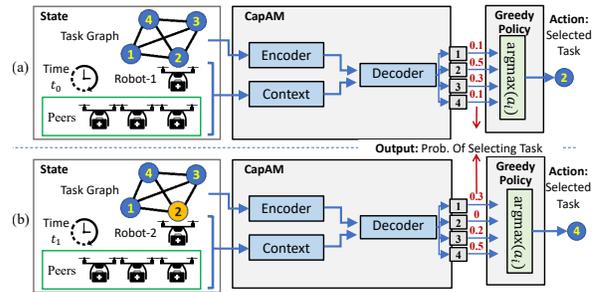


Fig. 1. Deployment of an MRTA policy based on the CapAM architecture. a) Robot-1 at t_0 . b) Robot-2 at t_1 ; here, the output for previously selected task (e.g., task 2 in (b)) is set at 0.

II. MRTA: PROBLEM DEFINITION AND FORMULATION

MRTA problems can be formulated as Integer Linear or Non-Linear Programming (ILP or INLP) problems. When tasks are defined in terms of location, the MRTA problem becomes analogous to the Multi-Traveling Salesmen Problem (mTSP) [42] and its generalized version, the Vehicle Route Planning (VRP) problem [43], albeit with additional constraints and operation-specific objective function.

A. MRTA problem description and formulation

The exact solution of the MRTA problem can be obtained by formulating it as the following INLP problem:

$$\min f_{\text{cost}} = \sum_{i=1}^N r_i, \begin{cases} r_i = \frac{t_i^f}{d_i}, & \text{if } t_i^f > d_i \\ 0, & \text{otherwise} \end{cases} \quad (1)$$

subject to

$$s_i \in S \quad \forall i \in [1..N] \quad (2)$$

$$s_i \neq s_j \quad \forall i \neq j \quad (3)$$

Here t_i^f is the time at which task i was completed, d_i is the time deadline of task i , and $S = [s_1, s_2, \dots, s_N]$ is the sequence of all the tasks/nodes that were visited, and N is the number of tasks available. The minimum cost function that can be achieved using Eq. 1 is 0, which corresponds to the case where all the tasks are successfully completed. A detailed formulation of the exact ILP constraints that describe the MRTA problem can be found in this recent work on MRTA [19]. Note that, in our paper, we use a slightly different objective/cost function as compared to [19], but the intention of the objective function of both of these works remains the same, which is to maximize the number of successfully completed tasks (thereby allowing fair comparison with methods reported in the literature). Here, we craft the objective function (Eq. 1) such that only missed tasks ($t_i^f > d_i$) contribute to the cost function. This objective function can be tailored according to user/application preferences, without affecting the proposed learning framework. The constraints in Eqs. 2 and 3 are such that each task must be visited exactly once by any robot.

B. An “MDP over Graph” Representation of MRTA

The task space of an MRTA problem can be represented as a complete un-directed graph $\mathcal{G} = (V, E, \mathbf{A})$, which contains a set of nodes/vertices (V), a set of edges (E) that connect

the vertices to each other, and the weighted adjacency matrix \mathbf{A} that gives the extent to which two nodes are connected. Each node is a task, and each edge connects a pair of nodes. For MRTA with N tasks, the number of vertices and the number of edges are N and $N(N-1)/2$, respectively. Node i is assigned a 3-dimensional feature vector denoting the task location and time deadline, i.e., $\delta_i = [x_i, y_i, d_i]$. Here we consider a weighted adjacency matrix without self-loop ($\alpha_{ii} = 0, \alpha_{i,j} \in \mathbf{A}$, where i and $j \in [1, N]$). The weights in the adjacency matrix $\alpha_{ij} = 1/(1 + |\delta_i - \delta_j|)$ take a real value between 0 and 1, in a manner such that if the features of two nodes i and j are relatively closer to each other, then α_{ij} is higher. Since we formulate MRTA as an undirected graph, the adjacency matrix \mathbf{A} is symmetric. In this paper, the elements of \mathbf{A} ($\alpha_{i,j}, i \neq j$) are computed as the inverse of the normalized Euclidean distance of the node properties.

The MDP can be defined in a decentralized manner to capture the myopic task-assignment process of each individual robot, which can be expressed as a tuple $\langle \mathcal{S}, \mathcal{A}, \mathcal{P}_a, \mathcal{R} \rangle$. Figure 1 shows how this MDP is executed to perform the sequential task selection by each robot. As evident from Fig. 1, the components of the MDP can then be defined as follows: **State Space (\mathcal{S}):** A robot at any decision-making instance uses a state $s \in \mathcal{S}$, which contains the following information: 1) elapsed mission time, 2) robot's current location, 3) constraints of this robot such as its work capacity, 4) the planned tasks of its peers, and 5) the constraints of its peers. The state space also includes the environment, defined by the location and time deadline of tasks, which are expressed here as node features in the graph. Here we assume that each robot can broadcast its information, which includes the task it is currently engaged in and its work capacity, to its peers without the need for a centralized system for communication; this is well aligned with modern communication capabilities [32]. **Action Space (\mathcal{A}):** The set of actions is represented as \mathcal{A} , where each action $a \in \mathcal{A}$ is defined as the index of the selected task, $\{1, \dots, N\}$. In MRTA problems that include a depot, task 0 can be considered as the depot. The task status, i.e., *active*, *completed*, and *missed* (i.e., deadline is passed), is used to properly decode the actions generated by the policy model, with the latter being in the form of probability of selecting each available task. Note that even under full observability, as the action space of a single robot includes only its own degree of affinity for selecting each available task, the decentralized decision-making formulation differs from an ideal centralized formulation, and remains readily scalable. **Transition (\mathcal{P}_a):** The transition is an event-based trigger. An event is defined as the condition that a robot reaches its selected task or visits the depot. As environmental uncertainties and communication issues (thus partial observation) are not considered in this paper, only deterministic state transitions are allowed. **Reward (\mathcal{R}):** A delayed reward is considered here, estimated at the end of the multi-robot simulation, i.e., when there are no more active tasks. The actual reward function is typically application-dependent. For example, in some MRTA problems, task completion is more important, while in another, decreasing

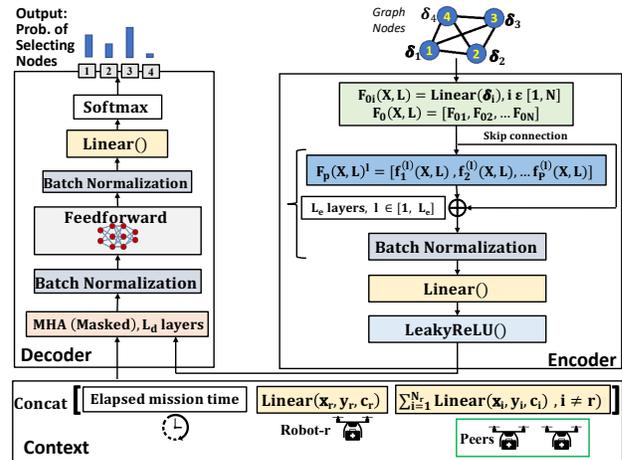


Fig. 2. CapAM architecture including context, encoder & decoder.

the total effort (e.g., distance traveled) by robots is more important, which can then be estimated using the weight matrix.

III. LEARNING FRAMEWORK

In this work, we implement an RL algorithm on an **encoder-decoder** architecture to learn the optimal MRTA policies over the graph formulation described earlier. An illustration of this architecture is shown in Fig. 2. The **encoder** represents each node as a continuous vector infusing both its own properties as well as its local and global structural properties within the graph representation of the task space. The decoder sequentially computes output probabilities for all the available tasks/nodes using the information from the encoder and the current state (**context**), by applying the MHA (as shown in left part of Fig. 2). The sequentially computed output probabilities for each node indicate the relative goodness of selecting that node as the next one to be visited by the concerned robot. A detailed description of the architecture of the encoder and decoder is discussed in the following sections using the representative SR-ST type MRTA problem described in section II-B.

A. Graph Capsule Convolutional Network based Encoder

The main purpose of the robot encoder is to represent useful information related to a node/task as a learnable continuous vector or tensor, which can then be used by the learning algorithm. For a graph node, the information includes the properties of the node itself (e.g., the coordinates and time deadline of the node), the local neighborhood information for the node, and also the global structural information. General solutions to CO problems such as CVRP, MTSP and MRTA should be permutation invariant, which means that the order by which each node is numbered should not affect the optimal solution. Hence the node encoding must also be permutation invariant. In this work, we are exploring how a Graph Capsule Convolutional Neural Network (GCAPCN) can be implemented for learning local and global structures with the node properties, with permutation invariant node embedding. GCAPCN is a class of Graph Neural Networks (GNN), introduced in [40] to address the drawbacks (e.g., permutation invariance) of Graph Convolutional Neural

Networks (GCN), and to enable the encoding of global information based on **capsule networks** presented in [44]. The advantage of GCAPCN lies in capturing more local and global information, compared to conventional aggregation operations used in GNN such as summation or standard convolution operations.

1) *Graph Capsule Convolutional Neural Networks*: Let $X \in \mathbb{R}^{N \times |\delta_i|}$, be the node feature matrix, where $|\delta_i|$ is the input dimension for each node i . The standard graph Laplacian is defined as $L = D - A \in \mathbb{R}^{N \times N}$, where D is the degree matrix and A is the adjacency matrix of the graph. A capsule vector is computed using a Graph Capsule function based on different order of statistical moments, as shown in the equations later in this section. We first compute a feature vector F_{0i} for each node by linear transformation of the node properties δ_i , as $F_{0i} = \delta_i \cdot W_0 + b_0$ for all $i \in [1, N]$, where $W_0 \in \mathbb{R}^{h_0 \times |\delta_i|}$, $b_0 \in \mathbb{R}^{h_0 \times 1}$, and h_0 is the length of the feature vector. For an MRTA problem with time constraints, $\delta_i = [x_i, y_i, d_i]$, where x_i , y_i , and d_i are the x coordinate, y coordinate, and the time deadline of location/node i . Note that this method for encoding is not limited to a specific MRTA problem, and can also be extended or modified for MRTA problems with nodes having other features such as workload. Each feature vector $F_{0i}, i \in [1, N]$ is then passed through a series of Graph capsule layers, where the output from the previous layers is used to compute a matrix $f_p^{(l)}(X, L)$ using a graph convolutional filter of polynomial form as given by:

$$f_p^{(l)}(X, L) = \sigma \left(\sum_{k=0}^K L^k (F_{(l-1)}(X, L)^{\circ p}) W_{pk}^{(l)} \right) \quad (4)$$

Here L is the graph Laplacian, p is the order of the statistical moment, K is the degree of the convolutional filter, $F_{(l-1)}(X, L)$ is the output from layer $l-1$, $F_{(l-1)}(X, L)^{\circ p}$ represents p times element-wise multiplication of $F_{(l-1)}(X, L)$. Here, $F_{(l-1)}(X, L) \in \mathbb{R}^{N \times h_{l-1} p}$, $W_{pk}^{(l)} \in \mathbb{R}^{h_{l-1} p \times h_l}$. The variable $f_p^{(l)}(X, L) \in \mathbb{R}^{N \times h_l}$ is a matrix where each row is an intermediate feature vector for each node $i \in [1, N]$, infusing nodal information from $L_e \times K$ hop neighbors, for a value of p . The output of layer l is obtained by concatenating all $f_p^{(l)}(X, L)$, as given by:

$$F_l(X, L) = [f_1^{(l)}(X, L), f_2^{(l)}(X, L), \dots, f_P^{(l)}(X, L)] \quad (5)$$

Here P is the highest order of statistical moment, and h_l is the node embedding length of layer l . We consider all the values of h_l (where $l \in [0, L_e]$) to be the same for this paper. Equations 4 and 5 were computed for L_e layers, where each layer uses the output from the previous layer ($F_{l-1}(X, L)$). Adding more layers helps in learning the global structure, however, this can affect the performance by increasing the number of learnable parameters (compared to the size of the problem), leading to over-fitting. The output from the final layer is then passed through a feed-forward layer so that the final feature vector has the right dimension (h_l) to be fed into the decoder as shown in Fig. 2

B. Attention-based Decoder and Context

The main objective of the MHA-based decoder is to use the information from the encoder and the current state

as context or query, and thereof choose the best task by calculating the probability value of getting selected for each (task) node. In this case, the first step is to feed the embedding for each node (from the encoder) as **key-values** (\mathcal{K}, V), since inputs for MHA are key-value pairs [45]. The key K and value V for each node is computed by two separate linear transformations of the node embedding obtained from the encoder. The next step is to compute a vector representing the current state, also known as the **context** (as shown in bottom left of Fig. 2). In our case studies, the context for the multi-head attention or MHA layer in the decoder consists of the following five features: 1) elapsed mission time; 2) work capacity of the robot taking decision; 3) current location of the robot taking decision; 4) current destination of robot's peers; and 5) work capacity of peer; all concatenated to a single vector of length h_q , which then undergoes a linear transformation to get a vector of length h_l also called the query Q . It is important to note that the **context** can be modified to account for other features such as remaining ferry-range and payloads of robots, thus promoting the flexibility to apply this architecture to a wider range of MRTA problems than those specifically studied in this paper. Figure 2 illustrates the structure of the decoder.

Now the attention mechanism can be described as mapping the query (Q) to a set of key-value (\mathcal{K}, V) pairs. The inputs, which are the query (Q) is a vector, while \mathcal{K} and V are matrices of size $h_l \times N$ (since there are N nodes). The output is a weighted sum of the values V , with the weight vector computed using the compatibility function expressed as:

$$\text{Attention}(Q, \mathcal{K}, V) = \text{softmax}(Q^T \mathcal{K} / \sqrt{h_l}) V^T \quad (6)$$

Here h_l is the dimension of the key of any node i ($k_i \in \mathcal{K}$). In this work, we implement a MHA layer in order to determine the compatibility of Q with \mathcal{K} and V . The MHA implemented in this work is similar to the decoder implemented in [23] and [45]. As shown in [45] the MHA layer can be defined as:

$$\text{MHA}(Q, \mathcal{K}, V) = \text{Linear}(\text{Concat}(\text{head}_1 \dots \text{head}_{h_e})) \quad (7)$$

Here $\text{head}_i = \text{Attention}(Q, \mathcal{K}, V)$ and h_e (taken as 8 here) is the number of heads. The feed-forward layer is implemented to further process the mapping that results from the MHA layer, and transform it to a dimension that is coherent with the number of nodes in the task-graph (N). The interjecting batch normalization layers serve to bound values of a specific batch using the mean and variance of the batch. The final softmax layer outputs the probability values for all the nodes. Here, the next task to be done is then chosen based on a greedy approach, which means that the node with the highest probability will be chosen. The nodes which are already visited will be masked (by setting their probability as 0) so that these nodes are not available for selection in the future time steps of the simulation of the multi-robot operation.

C. Training Algorithm

The training algorithm used here is REINFORCE. Other relevant details for training are shown in Table 1. This training process can readily be further advanced in the future

through the adoption of advanced state-of-the-art policy gradient algorithms such as Proximal Policy Optimization [46] and the Actor-Critic method [47]. For each epoch of the training, two sets of data (MRTA operation scenarios) are used, which are the training set and the validation set. The training data set is used to train the model (θ_{CapAM}) while the validation set is used to update the baseline model ($\theta_{\text{CapAM}}^{\text{BL}}$). Both the training data and validation data generated from the distribution of MRTA scenarios are explained later in section IV-A.1. The size of the training data and the validation data used in this paper is mentioned in Table II.

IV. EXPERIMENTAL EVALUATION

A. Comparison with baselines

1) *Details on dataset for training and testing:* We consider the (NP-hard) MRTA problem with time and capacity constraints also known as Task Allocation Problem with Time and Capacity (TAPTC), as described in [19].

Each training sample has 100 tasks, and are located randomly within a 100×100 grid map. Each task i has a time deadline $50 \leq d_i \leq 600$, and a workload $10 \leq w_i \leq 30$. Each sample has n_r number of robots where $2 \leq n_r \leq 7$. The initial positions of the robots in a sample is also chosen randomly within the grids. Each robot j has a work capacity of c_j where $1 \leq c_j \leq 3$. All the robots move at a speed of 1 unit. A task i is considered to be completed only if a robot j visits node i and spends a time of w_i/c_j . All the training samples are generated such that all the associated variables (mentioned above) follow a uniform distribution within their respective bounds. The average training time per epoch for CapAM is ~ 10 mins. In order to evaluate and compare the generalizability of CapAM-trained models with that of baselines, we used the dataset in [19] (found in [48]). This dataset, which has a similar distribution of scenarios as used in training CapAM, consists of 96 test cases with 100 tasks and a varying number of robots (2, 3, 5, 7 robots), and the speed of every robot is considered to be the same (1 unit/s). The 96 cases are divided into ones with tight and slack deadlines. Each category can be further divided into 4 sub-categories based on the fraction of tasks (25%, 50%, 75%, and 100%) that have normally distributed tasks deadlines. For example, 25% indicates that there are 25 tasks with deadlines normally distributed between the limits d_{low} and d_{high} , while the remaining 75 tasks have a deadline of d_{high} . For group 1 (tight deadline), the value of d_{high} is half of that for group 2 (slack deadline). Further description of the test cases can be found in [19].

2) *Test for Generalizability:* We compare our method with the following four baseline methods over the stated test cases:

i) **Iterated Local Search (ILS):** This is an online metaheuristic iterated search algorithm [20], where the output of

TABLE I
SETTINGS FOR MODEL TRAINING IN
CAPAM AND AM-RL

DETAILS	VALUES
<i>Algorithm</i>	REINFORCE
<i>Baseline</i>	Rollout
<i>Epochs</i>	100
<i>Training samples</i>	500,000
<i>Validation samples</i>	10,000
<i>Optimizer</i>	Adam
<i>Learning step size</i>	0.0001
<i>Training frequency</i>	500 SAMPLES

TABLE III
COMPARISON OF COMPUTING TIME TO GENERATE ENTIRE SOLUTION
SEQUENCE, AVERAGED OVER 96 CASES

METHOD	TIME (SEC)
ILS	1.342
EILS	0.992
BiG-MRTA	0.301
CapAM	0.085

one iteration is partially used as the input to the next iteration. During each iteration, the best solution is improved by a perturbation step, followed by a local search.

ii) **Enhanced Iterated Local Search (EILS):** EILS is also an online metaheuristic iterated search method [19], with an improved perturbation step as compared to [20].

iii) **Bi-Graph MRTA (BiG-MRTA):** BiG-MRTA [16] is an online method based on the construction and matching of a bipartite graph. In this method, a bipartite graph is constructed to connect robots to tasks, with the weights of connecting edges determined by an incentive model as a function of the tasks' features and robots' states. This decomposes the problem and yield a measure of robot-task pairing suitability. A maximum weighted matching problem is then solved by each robot to identify the optimal task assignments that maximize a net incentive for the team.

iv) **Attention Mechanism based RL (AM-RL):** The AM-RL method [23] consists of an encoder-decoder architecture purely based on attention mechanism. To implement this method for our problem, it is adapted to a multi-robot setting by making the following changes: 1) The node properties defined in Section II-B are used here; 2) The context for the MHA in the decoder is modified to be the same as that in CapAM; and 3) The cost function used for training is changed to that in Eq. (1). We set the parameters which correspond to the best performing model for a CVRP with 100 locations. This AM-RL model is trained using the same sample distribution and settings as respectively described in Sections IV-A.1. The average training time per epoch is found to be ~ 11 mins for AM. The architectural difference between AM-RL and CapAM is mainly in the encoder, with their performance comparison intended to show the impact of better encoding for learning over graphs.

TABLE II

The computation time to generate the entire sequence of task assignments, as required by the non-learning baselines and the learnt policies of CapAM and AM, are compared in Table III

– reported as an average over the 96 test cases. The real-time performance advantage of the learnt models, which are >10 times faster than the non-learning online methods (ILS/EILS/BiG-MRTA) is readily evident here. Figure 3 shows the box plots for task completion performance over the 96 test cases divided into 4 categories (each with 24 cases) based on the number of robots (2,3,5, and 7). As seen

COMPARISON OF COMPUTING
TIME TO GENERATE ENTIRE
SOLUTION SEQUENCE,
AVERAGED OVER 96 CASES

METHOD	TIME (SEC)
ILS	1.342
EILS	0.992
BiG-MRTA	0.301
AM-RL	0.085

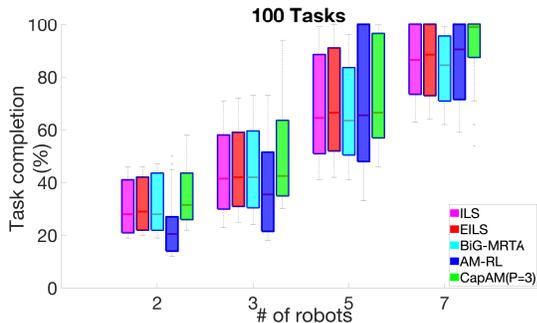


Fig. 3. Comparison of the performance (on test cases) of CapAM with other baseline methods. Here CapAM has $K = 2$, $P = 3$, and $L_e = 1$.

from Fig. 3, CapAM obtains a better performance compared to the baseline methods. To provide statistical evidence of the advantage of CapAM over the baseline methods, we performed a pairwise T-test of CapAM vs each baseline on the 24 cases in each category, with results summarized in Table IV; here, the null hypothesis is that the difference between the values of the two sets has a mean equal to 0. Considering a 5% significance level, Table IV shows that the **p-value** from the T-test is generally less than 0.05, which indicates the rejection of the null hypothesis – CapAM’s performance on the test cases is thus significantly better than of the baselines.

TABLE IV

P-VALUE FOR PAIRWISE (CAPAM VS. BASELINES) T-TEST

# robots	CapAM vs			
	ILS	EILS	BiG-MRTA	AM-RL
2	2.03e-4	0.002	0.003	1.6e-11
3	0.0307	0.079	0.038	2.4e-8
5	0.024	0.163	8.8e-4	0.42
7	0.025	0.049	0.003	0.031

B. Parametric and Scalability Analysis

Here, parametric analyses is performed w.r.t. the key parameters of P and L_e in the encoder of CapAM. Higher values of P promotes greater information infusion for a given neighborhood; with increasing L_e or K (and fixed P), information from (great number of) $L_e \times K$ hop neighbours is considered in the node embeddings. Increasing K , P and L_e however increases the number of learnable parameters, which can lead to over-fitting and deteriorating performance for unseen cases. Our numerical experiments with higher values of K and L_e corroborated this expectation, and hence we fixed $K = 2$ and $L_e = 1$. Hence this section focuses on the impact of P , particularly in generalizing, and scaling to larger-sized problems beyond training – both of which are premised as major advantages of learning local structural information. To this end, we train 3 different models on the 100-task problem, with varying P from 1 to 3, and fixed K and L_e . Then the CapAM (and for comparison AM-RL) models that are trained for 100 tasks and 2-7 robots are implemented on two sets of unseen problems – one with the same size as in training, and another with numbers of tasks and robots scaled up by a factor of 2 (training tasks are drawn from the same distribution).

Table V shows the effect of P on scalability of CapAM, which is also compared to AM-RL. Results are reported in terms of task completion % averaged over 24 samples in

each category (based on # tasks and robots). It can be seen that all CapAM models with any P perform significantly better than AM-RL. For the problems of the same size as training (100-task cases), the task completion % gap between CapAM and AM-RL is between 3% and 13.5%. For the problems of double the size as training (200-task cases), the task completion % gap between CapAM and AM-RL is between 1.9% and 17.6%. In both cases, the gap (CapAM vs. AM-RL) is the highest when task-to-robot ratios are high. These observations directly demonstrate the effectiveness of the novel encoder used in CapAM, which is designed to capture local structural information much better than in attention mechanisms. On comparing the CapAM models with different P (in Table V), it can be seen that ones with $P = 1$ perform slightly poorer compared to corresponding ones with $P = 2$ and $P = 3$, with the difference between the latter two being marginal. Thus increasing P beyond 2 may not provide further performance improvements.

TABLE V

PERFORMANCE OF CAPAM (IN TERMS OF % OF TASK COMPLETED) WITH $P = 1, 2, 3$ FOR UNSEEN PROBLEM SETS WITH THE SAME AND TWICE THE NUMBER OF TASKS AND ROBOTS AS IN TRAINING.

# tasks	# robots	CapAM($P=3$)	CapAM($P=2$)	CapAM($P=1$)	AM-RL
100	2	34.6	36.2	32.6	22.7
	3	51.1	50.3	49.4	38.0
	5	74.9	74.9	73.2	70.2
	7	94.8	94.1	91.2	85.6
200	4	28.2	27.1	24.4	10.7
	6	46.7	47.0	45.2	29.4
	10	67.8	68.6	66.5	61.4
	14	77.8	77.5	76.2	74.3

V. CONCLUSION

In this paper, we proposed a new graph neural network architecture called CapAM, to learn policies for MRTA problems involving tasks with time deadlines and robots with work capacity. This new architecture incorporates an encoder based on capsule networks and a decoder based on the attention mechanism, along with a context module to feed the state of the robots and the environment. To learn the features of the encoder and decoder, the problem has been posed as an RL problem and solved using the policy gradient algorithm, REINFORCE. In addition, the proposed architecture is found to provide effective MRTA policies over varying task size and robot-team size. The new CapAM architecture demonstrated better performance compared to other state-of-the-art baselines, which include both a learning based method (AM-RL) and non-learning based methods (ILS, EILS, BiG-MRTA). The CapAM architecture with its capsule based node embedding showed that learning local and global structural information of the task graph results in better generalizability over unseen test cases, as observed from comparing its performance with AM-RL (that uses a different node embedding). The computational cost analysis showed that the trained CapAM model takes only a few milliseconds to yield a task-assignment decision, making it highly suited for operations that require time-sensitive online decisions. The advantage of using local neighborhood information was also empirically evident from the scalability analysis on the defined MRTA problem, where CapAM demonstrated superior performance when applied to graphs

with a larger number of tasks/nodes and robots in the team.

Future directions: Firstly, computational analysis of training feasibility when episodes must be simulated over more realistic virtual robotics environment remains an open direction of research across all learning-based approaches. Further, to enable transition of these MRTA methods to application, in the future we should consider dynamic tasks, environment uncertainties, and partially observable state spaces within the CapAM architecture.

REFERENCES

- [1] B. P. Gerkey and M. J. Mataric, "A formal analysis and taxonomy of task allocation in multi-robot systems," *The International Journal of Robotics Research*, vol. 23, no. 9, pp. 939–954, 2004.
- [2] P. Ghassemi and S. Chowdhury, "Decentralized Task Allocation in Multi-Robot Systems via Bipartite Graph Matching Augmented With Fuzzy Clustering," in *Volume 2A: 44th Design Automation Conference*. American Society of Mechanical Engineers, aug 2018, p. V02AT03A014.
- [3] J.-P. Aurambout, K. Gkoumas, and B. Ciuffo, "Last mile delivery by drones: An estimation of viable market potential and access to citizens across european cities," *European Transport Research Review*, vol. 11, no. 1, p. 30, 2019.
- [4] M. V. Espina, R. Grech, D. De Jager, P. Remagnino, L. Iocchi, L. Marchetti, D. Nardi, D. Monekosso, M. Nicolescu, and C. King, "Multi-robot teams for environmental monitoring," in *Innovations in Defence Support Systems-3*. Springer, 2011, pp. 183–209.
- [5] E. Olson, J. Strom, R. Morton, A. Richardson, P. Ranganathan, R. Goeddel, M. Bulic, J. Crossman, and B. Marinier, "Progress toward multi-robot reconnaissance and the magic 2010 competition," *Journal of Field Robotics*, vol. 29, no. 5, pp. 762–792, 2012.
- [6] E. Nunes, M. Manner, H. Mitiche, and M. Gini, "A taxonomy for task allocation problems with temporal and ordering constraints," *Robotics and Autonomous Systems*, vol. 90, pp. 55–70, 2017.
- [7] R. Nallusamy, K. Duraiswamy, R. Dhanalaksmi, and P. Parthiban, "Optimization of non-linear multiple traveling salesman problem using k-means clustering, shrink wrap algorithm and meta-heuristics," *International Journal of Nonlinear Science*, vol. 8, no. 4, pp. 480–487, 2009.
- [8] P. Toth and D. Vigo, *Vehicle routing: problems, methods, and applications*. SIAM, 2014.
- [9] M. B. Dias, R. Zlot, N. Kalra, and A. Stentz, "Market-based multirobot coordination: A survey and analysis," *Proceedings of the IEEE*, vol. 94, no. 7, pp. 1257–1270, 2006.
- [10] T. Bektas, "The multiple traveling salesman problem: an overview of formulations and solution procedures," *Omega*, vol. 34, no. 3, pp. 209–219, 2006.
- [11] K. Braekers, K. Ramaekers, and I. Van Nieuwenhuysse, "The vehicle routing problem: State of the art classification and review," *Computers & Industrial Engineering*, vol. 99, pp. 300–313, 2016.
- [12] N. Azi, M. Gendreau, and J.-Y. Potvin, "An exact algorithm for a vehicle routing problem with time windows and multiple use of vehicles," *European Journal of Operational Research*, vol. 202, no. 3, pp. 756–763, 2010.
- [13] D. Wang, M. Hu, and Y. Gao, "Multi-criteria mission planning for a solar-powered multi-robot system," in *ASME 2018 International Design Engineering Technical Conferences and Computers and Information in Engineering Conference*. American Society of Mechanical Engineers Digital Collection, 2018.
- [14] K. Jose and D. K. Pratihari, "Task allocation and collision-free path planning of centralized multi-robots system for industrial plant inspection using heuristic methods," *Robotics and Autonomous Systems*, vol. 80, pp. 34–42, 2016.
- [15] D. Cattaruzza, N. Absi, and D. Feillet, "Vehicle routing problems with multiple trips," *4OR*, vol. 14, no. 3, pp. 223–259, 2016.
- [16] P. Ghassemi, D. DePauw, and S. Chowdhury, "Decentralized Dynamic Task Allocation in Swarm Robotic Systems for Disaster Response: Extended Abstract," in *2019 International Symposium on Multi-Robot and Multi-Agent Systems (MRS)*. New Brunswick, NJ: IEEE, aug 2019, pp. 83–85. [Online]. Available: <https://ieeexplore.ieee.org/document/8901062/>
- [17] E. Schneider, E. I. Sklar, S. Parsons, and A. T. Özgelen, "Auction-based task allocation for multi-robot teams in dynamic environments," in *Conference Towards Autonomous Robotic Systems*. Springer, 2015, pp. 246–257.
- [18] M. Otte, M. J. Kuhlman, and D. Sofge, "Auctions for multi-robot task allocation in communication limited environments," *Autonomous Robots*, vol. 44, no. 3, pp. 547–584, 2020.
- [19] H. Mitiche, D. Boughaci, and M. Gini, "Iterated local search for time-extended multi-robot task allocation with spatio-temporal and capacity constraints," *Journal of Intelligent Systems*, 2019.
- [20] P. Vansteenwegen, W. Souffriau, G. Vanden Berghe, and D. Van Oudheusden, "Iterated local search for the team orienteering problem with time windows," *Computers & Operations Research*, vol. 36, no. 12, pp. 3281–3290, 2009, new developments on hub location. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S030505480900080X>
- [21] S. Ismail and L. Sun, "Decentralized hungarian-based approach for fast and scalable task allocation," in *2017 International Conference on Unmanned Aircraft Systems (ICUAS)*. IEEE, 2017, pp. 23–28.
- [22] P. Ghassemi and S. Chowdhury, "Multi-robot task allocation in disaster response: Addressing dynamic tasks with deadlines and robots with range and payload constraints," *Robotics and Autonomous Systems*, vol. 147, p. 103905, 2022. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0921889021001901>
- [23] W. Kool, H. Van Hoof, and M. Welling, "Attention, learn to solve routing problems!" in *7th International Conference on Learning Representations, ICLR 2019*, 2019.
- [24] T. D. Barrett, W. R. Clements, J. N. Foerster, and A. I. Lvovsky, "Exploratory combinatorial optimization with reinforcement learning," *arXiv preprint arXiv:1909.04063*, 2019.
- [25] E. Khalil, H. Dai, Y. Zhang, B. Dilkina, and L. Song, "Learning combinatorial optimization algorithms over graphs," in *Advances in Neural Information Processing Systems*, 2017, pp. 6348–6358.
- [26] Y. Kaempfer and L. Wolf, "Learning the multiple traveling salesman problem with permutation invariant pooling networks," *ArXiv*, vol. abs/1803.09621, 2018.
- [27] A. Mittal, A. Dhawan, S. Manchanda, S. Medya, S. Ranu, and A. Singh, "Learning heuristics over large graphs via deep reinforcement learning," *arXiv preprint arXiv:1903.03332*, 2019.
- [28] Z. Li, Q. Chen, and V. Koltun, "Combinatorial optimization with graph convolutional networks and guided tree search," in *Advances in Neural Information Processing Systems*, 2018, pp. 539–548.
- [29] A. Nowak, S. Villar, A. S. Bandeira, and J. Bruna, "A note on learning algorithms for quadratic assignment with graph neural networks," *stat*, vol. 1050, p. 22, 2017.
- [30] 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.
- [31] E. V. Tolstaya, J. Paulos, V. R. Kumar, and A. Ribeiro, "Multi-robot coverage and exploration using spatial graph neural networks," *ArXiv*, vol. abs/2011.01119, 2020.
- [32] Q. Sykora, M. Ren, and R. Urtasun, "Multi-agent routing value iteration network," in *37th International Conference on Machine Learning, ICML 2020*, 2020.
- [33] H. Dai, E. B. Khalil, Y. Zhang, B. Dilkina, and L. Song, "Learning combinatorial optimization algorithms over graphs," in *Advances in Neural Information Processing Systems*, 2017.
- [34] R. Paleja, A. Silva, L. Chen, and M. Gombolay, "Interpretable and personalized apprenticeship scheduling: Learning interpretable scheduling policies from heterogeneous user demonstrations," in *Advances in Neural Information Processing Systems*, 2020.
- [35] M. Strens and N. Windelinckx, "Combining planning with reinforcement learning for multi-robot task allocation," in *Adaptive Agents and Multi-Agent Systems II*, D. Kudenko, D. Kazakov, and E. Alonso, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2005, pp. 260–274.
- [36] K. Lin, R. Zhao, Z. Xu, and J. Zhou, "Efficient large-scale fleet management via multi-agent deep reinforcement learning," in *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 2018, pp. 1774–1783.
- [37] K. S. Shetye, T. J. Overbye, H. Li, J. Thekkemathote, and H. Scribner, "Considerations for interconnection of large power grid networks," in *2021 IEEE Power and Energy Conference at Illinois (PECI)*, 2021, pp. 1–8.
- [38] M. Liu, K. Sivakumar, S. Omidshafiei, C. Amato, and J. P. How, "Learning for multi-robot cooperation in partially observable stochas-

- tic environments with macro-actions,” in *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2017, pp. 1853–1860.
- [39] K.-C. Chen, S.-C. Lin, J.-H. Hsiao, C.-H. Liu, A. F. Molisch, and G. P. Fettweis, “Wireless networked multirobot systems in smart factories,” *Proceedings of the IEEE*, vol. 109, no. 4, pp. 468–494, 2021.
- [40] S. Verma and Z. L. Zhang, “Graph capsule convolutional neural networks,” 2018.
- [41] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, “Attention is all you need,” *CoRR*, vol. abs/1706.03762, 2017. [Online]. Available: <http://arxiv.org/abs/1706.03762>
- [42] A. Khamis, A. Hussein, and A. Elmogy, “Multi-robot task allocation: A review of the state-of-the-art,” in *Cooperative Robots and Sensor Networks 2015*. Springer, 2015, pp. 31–51.
- [43] G. B. Dantzig and J. H. Ramser, “The truck dispatching problem,” *Management science*, vol. 6, no. 1, pp. 80–91, 1959.
- [44] G. E. Hinton, A. Krizhevsky, and S. D. Wang, “Transforming auto-encoders,” in *Artificial Neural Networks and Machine Learning – ICANN 2011*, T. Honkela, W. Duch, M. Girolami, and S. Kaski, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2011, pp. 44–51.
- [45] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, “Attention is all you need,” *CoRR*, vol. abs/1706.03762, 2017. [Online]. Available: <http://arxiv.org/abs/1706.03762>
- [46] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, “Proximal policy optimization algorithms,” 2017.
- [47] D. Bahdanau, P. Brakel, K. Xu, A. Goyal, A. Courville, R. L. J. Pineau, and Y. Bengio, “An actor-critic algorithm for sequence prediction,” in *5th International Conference on Learning Representations, ICLR 2017 - Conference Track Proceedings*, 2017.
- [48] <http://tinyurl.com/taptc15in>