

RTAW: An Attention Inspired Reinforcement Learning Method for Multi-Robot Task Allocation in Warehouse Environments

Aakriti Agrawal*, Amrit Singh Bedi*, and Dinesh Manocha*

Abstract—We present a novel reinforcement learning based algorithm for multi-robot task allocation problem in warehouse environments. We formulate it as a Markov Decision Process and solve via a novel deep multi-agent reinforcement learning method (called RTAW) with attention inspired policy architecture. Hence, our proposed policy network uses global embeddings that are independent of the number of robots/tasks. We utilize proximal policy optimization algorithm for training and use a carefully designed reward to obtain a converged policy. The converged policy ensures cooperation among different robots to minimize total travel delay (TTD) which ultimately improves the makespan for a sufficiently large task-list. In our extensive experiments, we compare the performance of our RTAW algorithm to state of the art methods such as myopic pickup distance minimization (greedy) and regret based baselines on different navigation schemes. We show an improvement of upto 14% (25-1000 seconds) in TTD on scenarios with hundreds or thousands of tasks for different challenging warehouse layouts and task generation schemes. We also demonstrate the scalability of our approach by showing performance with up to 1000 robots in simulations.

I. INTRODUCTION

Recent advancements in artificial intelligence (AI) and robotics have contributed to the widespread use of automation in warehouse environments. A recent case study by Honeywell¹ indicated more than a 50% rise in e-commerce sales after COVID-19 amidst existence of human labor shortage, thus proving the need of automating fulfillment and distribution centers/warehouses. Warehouse automation, widely being used by online retailers like Amazon, Alibaba, etc, can be performed by using multiple robots and a decision-making system. We are considering the task of picking and delivering a parcel from point A to point B [1]. By designing an appropriate multi-robot decision-making system, we can improve the task efficiency and throughput. The problem of designing such a system can be further divided into two parts: *task allocation* [2] and *robot navigation* [3]. Generally, greedy approaches are used to solve task allocation and multi-agent planning for navigation [4], [5].

We mainly deal with the problem of task allocation in multi-robot settings, which is popularly known as multi-robot task allocation (MRTA). This problem is widely studied in the literature using techniques from combinatorial optimization [2]. The problem of computing the optimal solution is NP-hard in general, and most practical solutions are based on greedy techniques [5], [6]. A comprehensive survey of

MRTA problems is given in [7], which divides the problem into multiple categories depending on the robot types, task types, and task assignment to robots.

In this paper, we address the problem corresponding to the ST-SR-IA category, where we consider single-task robots (ST) which execute only a single task at any given time. All tasks are single robot tasks (SR) and task assignment is instantaneous (IA) where the task is instantaneously allocated at each decision timestep. This class of problems can be solved in polynomial time at each decision timestep [2]. Prior methods focus on developing a deterministic solution for task allocation for a given set of tasks and robots [8]. We are interested in a sequential-lifelong version of the problem where new tasks are continuously generated and task is allocated to a robot as soon as it becomes available. Our goal is to maximize the sum of utilities over time, that is, assign tasks at each instant in such a way that the total time taken to execute a specific number of tasks, *makespan* as well as extra travel time *TTD* is minimized.

The existing state-of-the-art task allocation methods are mainly designed for multi-agent pickup and delivery (MAPD) and are primarily coupled with grid based navigation schemes, which are mostly designed for centralized systems [5]. Our method when combined with decentralised navigation algorithms can give rise to decentralised task-allocation and navigation with obstacle avoidance [9]. But the current practical warehouses uses only grid based navigation. Thus, this work focuses on the commonly occurring system in warehouses which is based on centralized task allocation and navigation.

Main Results: The novel components of our work include:

- 1) We formulate (first time in literature) the MRTA problem in warehouse settings as a reinforcement learning problem. This requires careful design of state, rewards, and actions for the problem to utilize any existing method to solve the problem.
- 2) The selection of reward for MRTA problem is quite challenging and we propose to use total travel delay (TTD) as the reward after evaluating different rewards like task length, total time of robot travel, etc. Interestingly, only the TTD based rewards exhibit converging behavior.
- 3) We propose a novel policy architecture that can handle a variable number of robots and tasks easily by forming useful global encoded representation using a mechanism inspired from attention. This is important because we can utilize the similar model to accommodate changes

*The Department of Computer Science, University of Maryland, College Park, MD, USA {agraval15, amritbd, dmanocha}@umd.edu

*This work was supported in part by ARO Grants W911NF1910069, W911NF2110026 and U.S. Army Cooperative Agreement W911NF2120076

¹<https://tinyurl.com/344p5e4y>

in the robot number during the execution process. with different number of robots. This also makes RTAW scalable to thousands of robots (Sec. III-C) in our simulations with constant space complexity and linear time complexity with respect to number of agents and tasks (cf. [10] Appendix VII).

- 4) We show extensive evaluations and compare against the state-of-the-art greedy baseline called myopic pickup distance minimization (MPDM) and regret-based task-selection (RBTS) for various navigation schemes, layouts and robots-number, and show an improvement of up to 14% (25 - 1000 sec) for 500 tasks. Our approach is trained with only single-agent A*, but works with other navigation schemes (SIPP and CBS). Hence it can be decoupled from the navigation system.

A. Related Works

Multi-robot task allocation (MRTA) has been widely studied in the robotics and related areas such as taxi-matching [11], food-delivery [12], search and rescue etc. MRTA problems are typically solved using (1) market-based approaches and (2) optimization-based approaches for different environments [13]. Market-based approaches are based on auction systems where each robot bids for a task with a base selling price determined by a broker. The negotiation process is based on market theory, where the team seeks to optimize an objective function based on robots' abilities to perform particular tasks [14]. These solutions [15], [16] can be efficient, robust, and scalable but difficult in terms of designing appropriate cost and revenue functions [17]. In contrast, optimization-based approaches outperform market-based approaches in handling large-scale MRTA scenarios (such as 50 tasks and 15 robots) [13]. In [18], the idea is to utilize a hybrid optimization approach combining a tabular search with a random search method. The authors in [19] combine a tabular search with a noising method to solve for static tasks, where tasks are not continuously entering but are predefined. Further, [20] proposes a simultaneous approach to solve the navigation and task allocation problems for a multi-robot system, where simulated annealing and ant colony optimization methods are applied. In contrast to the existing above mentioned approaches, our problem is sequential-lifelong single item task-allocation in warehouse settings and we aim to develop an algorithm that can be easily deployed with any navigation techniques. In this work, many auction and optimisation-based approaches boil down to greedy methods or modifications of them, which form our baselines.

Taxi-Matching Problem: The problem of task allocation in this work is closely related to the taxi-matching problem, where the goal is to match drivers to customers while reducing the driver wait time, increasing the driver's income, and reducing the driving miles [21]. The problem of matching pick-up and delivery tasks to drivers was initially solved using combinatorial optimization [11]. Recently, the same problem has been solved via multi-agent RL techniques [22], which divide the city map into many dispatching regions

(grids) and then learn a policy to assign drivers for a request from the given grid point (state).

Both taxi-matching and MRTA problems are sequential decision-making problems in terms of allocating tasks. However, warehouses have a fixed and comparatively small layout with a smaller and fixed number of robots deployed. In contrast, ride-sharing platforms tackle thousands of vehicles and millions of orders over a large area like a city. Therefore ride-sharing allocation techniques focus on the distribution of tasks and agents rather than computing their exact coordinate positions.

Thus, the multi-agent formulation for task allocation using reinforcement learning in [22], [23] for taxi-matching applications cannot be directly applied to warehouse settings.

Task Allocation in Warehouses: Task allocation in warehouses has also been studied [1], where it is formulated as an integer linear program and solved using a table-working method. Moreover, [24] describe an optimal task allocation for heterogeneous robots using navigation path, speed, and current location of each robot and test its validity in simulated environments. These methods are best suited for static settings, where all the robots are available at the same time to solve the matching problem. In practice, different robots take different times to complete a task. Therefore, our goal here is to design dynamic task allocation solutions that assign a task to a robot whenever it becomes available.

The problem of task allocation in warehouses is also briefly studied in multi-agent pickup and delivery problem. The current state-of-the-art [5] uses greedy solution for task allocation. It continuously allocates and de-allocates tasks based on marginal cost or regret on marginal cost. Such techniques cannot be decoupled from the navigation part. But authors in [25] proposed a decoupled solution that assigns tasks to agents first in a greedy manner and then utilizes a Conflict-Based Search (CBS) [4] strategy to plan collision-free paths. We use this commonly used greedy approach (also called MPDM) as one of the baselines.

II. RL BASED FORMULATION

In the existing literature, multi-robot task allocation problem in warehouse environments is solved by standard optimization methods [7]. But in practice, tasks appear at different time instances, and robots' availability also varies with time, which makes the problem challenging to solve via standard optimization methods. To address this issue, for the first time in this work, we formulate the multi-robot task allocation problem in warehouse environments as a reinforcement learning problem. The goal is to allocate tasks to multiple robots such that the tasks are performed in an efficient manner (i.e. resources like battery are better utilised and throughput is maximized). The formulation is challenging because it requires to define a Markov Decision Processes (MDP) for the problem. An MDP here is defined by a tuple $\mathcal{M} := (\mathcal{S}, \mathcal{A}, R, P, \gamma)$, where \mathcal{S} denotes the finite state space, \mathcal{A} denotes the finite action space, and $R(s, a)$ denotes the reward for state s and action a in \mathcal{S} and \mathcal{A} , respectively. Further, P represent the transition probability

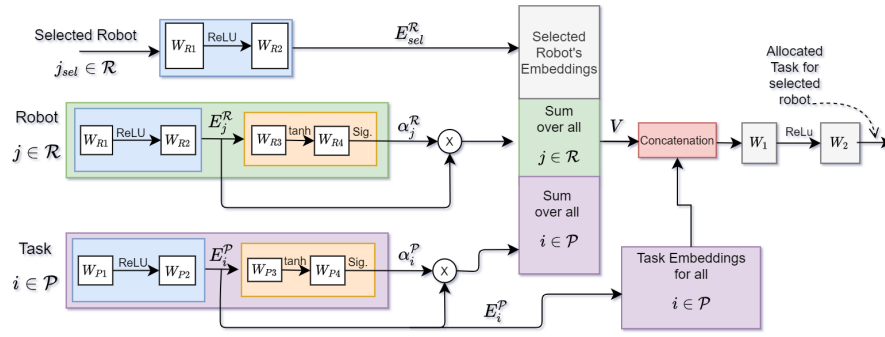


Fig. 1: We highlight our policy network architecture that is designed to work with variable number of robots and tasks in any warehouse layout. Our formulation takes the feature of robots and tasks as input to generate fixed dimensional embeddings. Next, these embeddings are mixed to generate a global state for the system, which is used to perform the action of choosing tasks for the selected robot.

matrix where $P_a(s, s')$ is the probability of transition to state $s' \in \mathcal{S}$ from state $s \in \mathcal{S}$ after taking action $a \in \mathcal{A}$. Here, $\gamma \in (0, 1)$ is the discount factor used to trade the near term reward with the long-term rewards in the RL formulation.

MDP model \mathcal{M} for the multi-robot task allocation problem: Let the time instance be $t \in \mathbb{N}$ and we consider a warehouse environment with the M number of robots collectively denoted by $\mathcal{R} = \{1, 2, \dots, M\}$. A warehouse environment is represented by a specific $K \times K$ grid layout (shown in Fig. 2). In this work, we do not restrict ourselves to any particular layout, and our approach is applicable to all general layouts. In fact, models trained on dense layouts have shown to give good performance on sparse layouts as well. Our test layout and task generation are described in section 4. After the task allocation at t , the corresponding task is removed from the queue and a new task is added. \mathcal{P} denotes the dynamic task list. To make the problem tractable, we assume that at each t , only one robot is available to be allocated one of the tasks. We accomplish this by adding a small Gaussian noise of mean 0.25 and variance 0.25 to the task completion time. This noise value is chosen to be in (1/100)th range of the task-completion-time range. This ensures that no two robots become available at the same point in time. All the other robots are currently performing the previously allocated tasks. We move from t to $t + 1$ as soon as a robot becomes available after completing a task.

State: Each task $i \in \mathcal{P}$ is represented by a tuple (o_i, d_i, k_i, l_i) , where o_i is the origin of the task i , d_i is the destination of task i , k_i denotes the distance between the selected robot's current position to o_i , and l_i is the length (distance between o_i and d_i) of the task. The origin o_i and destination d_i for each task i are the coordinates sampled from the 2D layout. Each robot $j \in \mathcal{R}$ is also represented by a tuple (p_j, r_j) , where p_j is position ((X, Y) coordinates) in the 2D map and r_j is the time left to complete the allocated task (i.e., the task completion time). Let j_{sel} denotes the available robot for task allocation. Then, the state s encapsulates the robots' positions, their task completion times, and the task list \mathcal{P} , which consists of tuples (o_i, d_i, k_i, l_i) for all the tasks and the selected robot j_{sel} . Collectively, we can write the state as $s :=$

$\{(p_j, r_j)_{\forall j \in \mathcal{R}}, (o_i, d_i, k_i, l_i)_{\forall i \in \mathcal{P}}, j_{sel}\}$ and the collection of all possible states s constitutes the state space \mathcal{S} .

Action: An action a in the environment is to select which task to execute from the list of N available tasks in the queue. We define a policy $\pi : \mathcal{S} \rightarrow \mathcal{D}(\mathcal{A})$ (here \mathcal{D} denotes the set of possible distributions defined over \mathcal{A}) which takes the current state as input and outputs a distribution across the all possible actions $a \in \mathcal{A}$. Note that in our model, \mathcal{A} constitutes the list of all available tasks from which we need to select a particular task for the selected robot j_{sel} . During the training time, we select the action $a \sim \pi(\cdot | s)$ where $\pi(\cdot | s)$ is the probability distribution across the tasks. During the test time, we select the action in a deterministic manner by choosing a task to accomplish with maximum probability in $\pi(\cdot | s)$. Further, we discuss the policy architecture in detail in Sec. II-A.

Reward: Designing an appropriate reward for the RL formulation is crucial because this reward eventually decides the behavior of the policy. We define the reward as the negative of the cost, which is the time a robot takes to reach the origin o_i of the task from its current location. In this work, we use the navigation schemes to calculate the approximate cost of the path by making the robot run through that path. This makes sense because during the time of travel from the robot's current location to the task origin, the robot is not serving any task, and hence we aim to minimize this time. Specifically, this is defined as total travel delay (TTD), which is the extra time the robot travels when it is not executing any task. Minimizing TTD would ultimately minimize the makespan in the lifelong task allocation. We note that other reward formulations are also possible (which we tried but didn't work) such as the negative of total *task-length* defined as the time taken to move from the origin to the destination of task.

Another possibility of reward is the negative of total travel time which is sum of TTD and task-length. That is, the time to travel from robot position to task origin plus time to complete the task. We call it *Reward-Travel Time*. We observe there exists unique minima for TTD. For task-length based rewards, it becomes non-stationary and thus hard to learn. We compare the different reward schemes in Table I.

Reward	Equation	Converged	Minimized TTD	Minimized Makespan
TTD (proposed)	$\text{dist}(p_j, o_i)$	Yes	Yes	Yes
Task length	$\text{dist}(o_i, d_i)$	No	No	No
TTD + Task length	$\text{dist}(p_j, o_i) + \text{dist}(o_i, d_i)$	Yes	No	No

TABLE I: The table shows a comparison of different rewards. Our reward (distance between robot position and task origin) minimizes both TTD and makespan and shows good convergence properties. This is not the case for the other two reward formulations, shown in second and third rows.

State-Transitions: At current state s_t , we have a task list in queue \mathcal{P} and one available robot (say $i \in \mathcal{R}$). At this state, we take action a_t according to a policy $\pi(\cdot | s_t)$ that decides which task would be allocated to the free robot. The selected task is removed from the list, new task gets added to the list, and the robot becomes unavailable. Our algorithm will wait until next robot becomes available which would mark our new state s_{t+1} . Note that the state transitions happens at irregular intervals.

A. Network Architecture using Attention Based Embeddings

The objective of an RL problem is to maximize the expected value of the cumulative reward returned under policy π_θ which is parametrized by θ . Mathematically, the objective is given as:

$$\max_{\theta} J(\theta), \quad (1)$$

where $J(\theta) := V_{\pi_\theta}(s_0)$. Here, $V_{\pi_\theta}(s_0)$ is the value function for any arbitrary initial state s_0 given as

$$V_{\pi_\theta}(s_0) = \mathbb{E}_{\pi} \left[\sum_{t=0}^{\infty} \gamma^t R(s_t, a_t) \mid s_0 = s, a_t \sim \pi_\theta(\cdot | s_t) \right].$$

The formulation in (1) gives rise to class of policy gradient algorithms [26]. The explicit form of the policy gradient is given by

$$\nabla_{\theta} J(\theta) = (1/1 - \gamma) \mathbb{E}_t [\nabla_{\theta} \log \pi_{\theta}(a_t | s_t) \cdot A_t], \quad (2)$$

where \mathbb{E}_t is the empirical average over a finite batch of samples, and A_t denotes the estimate of the advantage function $A_t = Q_{\pi_{\theta_t}}(s_t, a_t) - V_{\theta_t}(s_t)$. We use the popular proximal policy optimization (PPO) algorithm to solve the problem due to its advantages as compared to other RL methods such as data efficiency and reliable performance [27]. A generalized advantage estimator [28] is used to estimate A_t in this work. Here, $Q_{\pi_{\theta_t}}(s_t, a_t)$ denotes the state action value function.

Attention Inspired Policy Architecture: To this end, the most important part is to design the policy π_θ which is mostly defined as a deep neural network (DNN) in the literature [29] and θ corresponds to the weights of the DNN. The design of a specific neural network architecture for the policy is crucial and challenging because it will control the class of policies over which we will be optimizing the non-convex objective of (1). A naive architecture would be to simply concatenate all robot and task information (locations, availability etc.) and pass that as input to the network. The input layer size would then depend upon the number of robots and tasks. Moreover, the number of trainable parameters would keep increasing with the increase in number of tasks and robots. This kind of architecture is not practical due to the scalability

issues with respect to the number of agents and tasks in the given warehouse environment.

To address the issue of scalability, we take motivation from the problem of multi-driver vehicle dispatching problem as solved in [21] and propose an attention mechanism inspired policy architecture shown in Fig. 1. The proposed architecture is able to handle arbitrary number of robots and tasks in any warehouse layout via utilizing global embeddings. For each robot $j \in \mathcal{R}$, let us denote the features $\mathcal{F}_j^{\mathcal{R}} := \{p_j, r_j\}$, which consists of position coordinates $p_j = (x_j, y_j)$ and time to task completion r_j , are passed through one layer of 16 nodes and ReLU activation. This is followed by an another layer with 16 nodes to generate robot embeddings denoted by $E_j^{\mathcal{R}}$. This is important to extract the right information from the data input in the form of embeddings which is then used for further processing instead of the direct data input. Selected robot's embeddings $E_{sel}^{\mathcal{R}}$ are also generated in a similar manner. Mathematically, we have

$$\begin{aligned} E_j^{\mathcal{R}} &= W_{R2} \cdot \text{ReLU}(W_{R1} \cdot \mathcal{F}_j^{\mathcal{R}}), \\ E_{sel}^{\mathcal{R}} &= W_{R2} \cdot \text{ReLU}(W_{R1} \cdot \mathcal{F}_{sel}^{\mathcal{R}}). \end{aligned} \quad (3)$$

where W_{R1} and W_{R2} are the weights of the neural network. In a similar manner, for each task $i \in \mathcal{P}$, the input features $\mathcal{F}_i^{\mathcal{P}} = (o_i, d_i, k_i, l_i)$ are passed through one layer of 16 nodes with ReLU activation. This is followed by another 16 node layer to generate task embeddings $E_i^{\mathcal{P}}$ as $E_i^{\mathcal{P}} = W_{P2} \cdot \text{ReLU}(W_{P1} \cdot \mathcal{F}_i^{\mathcal{P}})$. The robot embeddings $E_j^{\mathcal{R}}$ as well as task embeddings $E_i^{\mathcal{P}}$ are passed through a 16 node layer with tanh activation and 1 node layer with sigmoid activation to compute scalar outputs $\alpha_j^{\mathcal{R}}$ and $\alpha_i^{\mathcal{P}}$, respectively, for a robot and the tasks given by

$$\alpha_j^{\mathcal{R}} = \text{sig}(W_{R4} \cdot \tanh(W_{R3} \cdot E_j^{\mathcal{R}})) \quad (4)$$

$$\alpha_i^{\mathcal{P}} = \text{sig}(W_{P4} \cdot \tanh(W_{P3} \cdot E_i^{\mathcal{P}})) \quad (5)$$

These scalar would provide different weightage to different robots or tasks. Next, we compute the weighted average of all robot embeddings as well as task embeddings given by

$$v^{\mathcal{R}} = \sum_j \alpha_j^{\mathcal{R}} E_j^{\mathcal{R}} \quad \text{and} \quad v^{\mathcal{P}} = \sum_i \alpha_i^{\mathcal{P}} E_i^{\mathcal{P}} \quad (6)$$

Further, the global task embeddings are computed using an attention-like mechanism [30] so that we generate the expected global information. The final vectors $v^{\mathcal{R}}$ and $v^{\mathcal{P}}$ are concatenated together with selected robot's embedding $E_{sel}^{\mathcal{R}}$ to generate one single vector $V = (v^{\mathcal{R}}, v^{\mathcal{P}}, E_{sel}^{\mathcal{R}})$ which denotes the state of the system. We remark that the dimension of this state of the system is independent of the number of tasks or agents in the system which is an important aspect of the proposed architecture. Moreover, to make an accurate prediction of action, the network requires both global and local information, which motivates us to concatenate selected robot (local embedding) and, $v^{\mathcal{R}}$ and $v^{\mathcal{P}}$ (global embeddings). Next, this one dimensional vector V is concatenated with two dimensional task embedding using broadcasting, which implies

$$\text{Output} = (V, E_i^{\mathcal{P}}), \quad \text{for } \forall i \in \mathcal{P}. \quad (7)$$

Then the Output is passed through one layer with 8 nodes and ReLU activation, then passed through another layer with one node, and finally passed through categorical distribution to choose the task to be executed.

$$Task^{chosen} = \text{Categorical}(W_2 \cdot \text{ReLU}(W_1 \cdot \text{Output})). \quad (8)$$

The output ($Task^{chosen}$) is in form of integer denoting the task number the selected robot is assigned to. We use categorical distribution during training time to sample the chosen task. During test time, we choose the task with the highest probability.

III. IMPLEMENTATION AND EVALUATION

In this section, we describe our implementation and analyze its performance in different settings by varying the number of robots, task generation schemes, layouts and navigation schemes. During training, we use two different navigation techniques to compute the cost for the RL environment: *direct navigation* and *A**. In the direct navigation approach, the Euclidean distance between the robot position and task origin is used to evaluate the cost at each instant t . In the *A** navigation, the *A** algorithm is used to calculate the distance. We compare our method (RTAW) with minimum pickup distance minimization (MPDM) and regret based task selection as our baselines. Both these baselines are used in sequential single-item auctions.

MPDM algorithm chooses the task that is closest to the robot. It is optimal if there is only one robot in the system and the task utility equals the distance of the robot to the task start position (see [11] for a more in depth discussion). MPDM is the only baseline used for comparison in the literature for decoupled task allocation and navigation [25]. The other state-of-the-art [5] solutions for warehouse environments are based on coupled task allocation and navigation. Therefore, they cannot be directly applied. We present a modified version of **RBTS** used in [5] inspired from [31], [32].

RBTS: For every task in the list, find the distance to the closest robot. Subtract the distance of the selected robot to each of these tasks with the previous quantity. Choose the task with the maximum value (i.e. maximum regret).

For all experiments, we run the test on a fixed task set by fixing the seed of the random number generator and report TTD values for them. We train our model on the Nvidia GeForce RTX 2080Ti GPU with 11GB of memory, for 4 different environments and batch size of 32. We run the training for 4 million iterations which takes around 12 hours with *A** navigation algorithm. The network architecture is not bulky with GPU usage less than 5-10 %. Refer to appendix of [10] for the training curves and hyper-parameters tuning. Next, we describe the different experimental settings and discuss the experimental results in detail.

A. First Experimental Setting

In appendix sec. VI-A of [10], we show a simple experiment as proof of concept with 2 agents and 5 tasks, to compare the three methods. In appendix sec. VI-B of [10], we show another simple experiment with two fixed

tasks both of which have the same starting point but with different destinations. The two tasks are chosen such that one is always longer than the other. Since the starting coordinates are same, greedy policy will randomly pick a task while RL policy learns to pick the smaller task. Thus, giving 50% and 55% improvement over MPDM for 10 and 100 robots, respectively.

B. Final Experimental Setting

Next, we raise the difficulty level further and shift focus to a more general setting where the tasks are generated randomly in the environment using designated method. We report the results in Table II for direct navigation and *A** navigation, respectively. We use the layouts as shown in 2 for this evaluation.

Task Generation: To the best of our knowledge, there are no publicly available datasets for the warehouse task allocation problem. Therefore, in the literature, the task generation for the experiments is synthetic and has been divided into three different ways random, clustered, and designated [3]. For the experiments in this section, we consider designated method for task generation where specific regions for pick-up and delivery are defined and tasks are randomly sampled from them. For our case, we define regions for pick-up and delivery using Gaussian distribution, and randomly sample our tasks from those distributions. This method is the closest to the practical task generation method and is most commonly used in literature [33], [34].

	Robot No.	MPDM	RBTS	RTAW	% Imp. MPDM	% Imp. RBTS
Direct	10	8312 ± 931	8376 ± 653	7924 ± 645	4.42 ± 3.27	5.36 ± 3.39
	100	13089 ± 751	13322 ± 905	12298 ± 747	6.04 ± 1.86	7.55 ± 4.96
Layout A	10	5662 ± 372	6194 ± 235	5165 ± 132	8.77 ± 4.96	16.6 ± 5.14
	100	7429 ± 213	7882 ± 408	7032 ± 393	5.33 ± 3.94	10.77 ± 2.84
Layout B	10	5591 ± 227	6098 ± 330	4989 ± 177	10.74 ± 1.28	18.1 ± 3.12
	100	7885 ± 173	8288 ± 488	7387 ± 355	3.34 ± 3.22	10.7 ± 5.2
Layout C	10	6205 ± 247	6503 ± 435	5861 ± 378	5.56 ± 4.2	9.74 ± 5.23
	100	7210 ± 201	7915 ± 338	6954 ± 164	3.53 ± 1.69	12 ± 4.63
Layout D	10	5920 ± 501	6675 ± 663	5719 ± 403	3.27 ± 3.23	14.09 ± 3.96
	100	7764 ± 764	8182 ± 712	7101 ± 357	8.11 ± 6.22	13 ± 4.56
Layout E	10	5536 ± 367	5916 ± 259	5281 ± 312	4.6 ± 2	10.7 ± 3
	100	7791 ± 354	8126 ± 349	7453 ± 207	4.33 ± 3.45	8.28 ± 5.1

TABLE II: TTD time for 500 tasks in seconds. We assume the robot moves 1 unit in 1 sec. We show results for Direct navigation scheme with no obstacles and for *A** navigation scheme on various layouts of grid size 60×60 . We generate results for 5 random seeds and report the mean and std of these 5 runs for each value. For eg. in the first entry of the table, [8312 ± 931], 8312 is mean of 5 runs and 931 is the std.

Layout Generation: In order to evaluate approach, we considered real-world warehouse environment layouts to test the performance of our algorithms. These are based on [3] and warehouse layouts generated through Asprilo. We evaluate the performance of the proposed algorithm on all the layouts in Fig. 2 and compare the performance for 10 and 100 robots summarized in Table II and IV.

The task queue length is 10 for all cases. As shown in Table IV, RTAW outperforms greedy (MPDM) as well as regret-based (RBTS) baselines for all layouts and navigation scenarios. For CBS [4] and SIPP [36] navigation schemes, we train the algorithm on *A** navigation and test on CBS

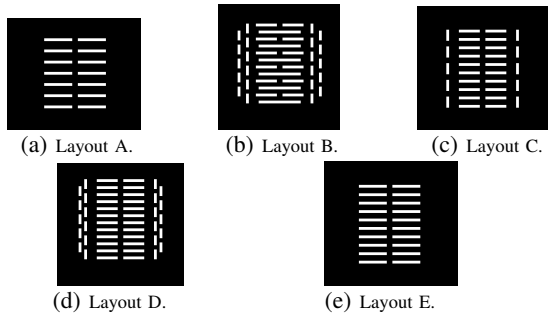


Fig. 2: This figure presents different real-world warehouse layouts generated with varying level of compactness using the Aspriolo framework for a fixed grid size of 60×60 [3], [35]. We use some of the most challenging layouts. Layout A-very low compactness, Layout B-high compactness, Layout C-low compactness, Layout D-high compactness, Layout E- low compactness.

and SIPP. These navigation algorithm takes a lot of time to solve path therefore training with these navigation schemes is infeasible. For the same limitation of these navigation algorithms, we show results only for 10 agents for various layouts. Our proposed architecture trained on one layout can be efficiently used for a similar layout with different compactness. For instance, trained policy network for layouts A and E can be used interchangeably. This is due to the similarity in task-distribution.

Robot Number	Grid Size	MPDM	RBTS	RTAW	MPDM (%)	RBTS (%)
500	Direct	547538	554681	503530	8.04	9.22
500	A* on Layout B	311403	301221	278590	10.54	7.5
1000	Direct	695816	681236	674089	3.12	1.04
1000	A* on Layout B	413406	423368	403112	2.49	4.78

TABLE III: TTD time in seconds for 5000 tasks for 500 and 1000 robots for direct and A* navigation scheme. We assume the robot moves 1 unit in 1 sec. Here we show output for only 1 run because with such large number of robots, baseline algorithm (RBTS) takes a very long to run.

Layout	SIPP				
	MPDM	RBTS	RTAW	MPDM (%)	RBTS (%)
A	692 ± 106	725 ± 94	645 ± 67	6.7 ± 6.4	11 ± 4
B	616 ± 33.7	636 ± 102	566 ± 70	8.4 ± 7.6	10.5 ± 7.6
C	684 ± 110	707 ± 138	626 ± 86	8.1 ± 4.95	10.6 ± 5.5
D	710 ± 131	706 ± 165	642 ± 140.5	10 ± 4.28	8.6 ± 4.7
E	643 ± 62.5	689 ± 110	595 ± 84	7.53 ± 5.5	13.7 ± 8.9
Layout	CBS				
	MPDM	RBTS	RTAW	MPDM (%)	RBTS (%)
A	644 ± 80	696 ± 63	615 ± 74	4.5 ± 1.6	11.6 ± 7.43
B	680 ± 114	675 ± 105	619 ± 99	8.9 ± 1.8	8.3 ± 3.7
C	655 ± 28	670 ± 45	625 ± 26.4	4.5 ± 3.8	6.6 ± 3.6
D	766 ± 121	817 ± 128	738 ± 125	3.73 ± 2.78	9.74 ± 4.3
E	618 ± 124	671 ± 127	591 ± 109	4.3 ± 4.2	11.98 ± 1.7

TABLE IV: TTD time for 50 tasks (in seconds) for 10 robots for SIPP and CBS navigation scheme respectively. We assume the robot moves 1 unit in 1 sec. The model is trained on only A* navigation scheme for various layouts (refer Fig. 2) of grid size of 60×60 .

C. Scalability

We show that RTAW is scalable by performing experiments on 500 and 1000 robots. For direct navigation, we used a 300×300 grid layout with no obstacles. For A* navigation we used a 256×256 grid size of layout C. The results are

summarized in Table III. We note that RTAW outperforms the baselines even for higher number of robots and tasks.

IV. ANALYSIS AND LIMITATIONS

Our final reward function based on TTD exhibits better behavior than previous reward functions based on task-length. This is because task-length is a fixed quantity and does not aid in learning. It makes it hard to compute the optimal function, whereas TTD has shown to exhibit good converging behavior. We observe that we do not get good improvement for all the seeds values. There are some cases, when our model performs similar to greedy strategy. But such cases are not frequent, as shown in our results. When we compute an average over multiple runs, we always see a clear improvement. Some hyper-parameters played an important role in formulating the problem and the performance of the learning method. We choose a high γ value as close as possible to 1 because we want to take actions that result in higher future reward and not emphasise only the immediate reward (which happens in the greedy method). We used entropy regularisation to aid the learning and deal with complex environment (i.e., larger robots, more tasks, compact layout). We reduce the entropy coefficient to help in the learning process. More details are given in the supplementary material.

Our RTAW algorithm has some limitations. One issue is handling the variable load capacity of the robots. Our current approach makes an assumption that a robot executes only one task (going from point A to point B) at a time. Another assumption we make is that robots always finish their task and are ready in terms of performing the next task right away. But this may not hold in practice due to robot's state, battery charging etc. Our current approach can handle variable number of robots so in-order to accommodate new robots coming in and leaving the system, we just need to maintain separate system that keeps track of the total number of robots. Currently, we assume that the robots are similar or homogeneous, that is, all robots travel with same velocity and acceleration, but this is may not be the case always in the real world.

V. CONCLUSION

We considered a problem of multi-robot task allocation (MRTA) in warehouse environment and proposed a novel and state-of-the-art RL-based solution we called RTAW. RTAW is general and can be applied to arbitrary warehouse layouts. We have evaluated the performance of RTAW on multiple layouts and under different experimental settings, where the number of robots is varied from 10 to 1000. We observe that our method outperforms prior techniques based on greedy and regret-based strategies and the computational complexity scales linearly with the number of robots. Furthermore, RTAW can be combined with any centralized or decentralized navigation approach.

REFERENCES

- [1] F. Xue, H. Tang, Q. Su, and T. Li, "Task allocation of intelligent warehouse picking system based on multi-robot coalition," *KSI Transactions on Internet and Information Systems*, vol. 13, no. 7, 2019. 1, 2
- [2] 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. [Online]. Available: <https://doi.org/10.1177/0278364904045564> 1
- [3] R. Stern, N. R. Sturtevant, A. Felner, S. Koenig, H. Ma, T. T. Walker, J. Li, D. Atzmon, L. Cohen, T. S. Kumar, et al., "Multi-agent pathfinding: Definitions, variants, and benchmarks," in *Twelfth Annual Symposium on Combinatorial Search*, 2019. 1, 5, 6
- [4] W. Hönig, S. Kiesel, A. Tinka, J. Durham, and N. Ayanian, "Conflict-based search with optimal task assignment," in *Proceedings of the International Joint Conference on Autonomous Agents and Multiagent Systems*, 2018. 1, 2, 5
- [5] Z. Chen, J. Alonso-Mora, X. Bai, D. D. Harabor, and P. J. Stuckey, "Integrated task assignment and path planning for capacitated multi-agent pickup and delivery," *IEEE Robotics and Automation Letters*, vol. 6, no. 3, pp. 5816–5823, 2021. 1, 2, 5
- [6] B. Gerkey and M. Mataric, "Sold!: auction methods for multirobot coordination," *IEEE Transactions on Robotics and Automation*, vol. 18, no. 5, pp. 758–768, 2002. 1
- [7] G. A. Korsah, A. Stentz, and M. B. Dias, "A comprehensive taxonomy for multi-robot task allocation," *The International Journal of Robotics Research*, vol. 32, no. 12, pp. 1495–1512, 2013. 1, 2
- [8] A. Frank, "On kuhn's hungarian method—a tribute from hungary," *Naval Research Logistics*, vol. 52, pp. 2–5, 2005. 1
- [9] A. Agrawal, S. Hariharan, A. S. Bedi, and D. Manocha, "Dc-mrta: Decentralized multi-robot task allocation and navigation in complex environments," in *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2022, pp. 11 711–11 718. 1
- [10] A. Agrawal, A. S. Bedi, and D. Manocha, "Rtaw: An attention inspired reinforcement learning method for multi-robot task allocation in warehouse environments," *arXiv preprint arXiv:2209.05738*, 2022. 2, 5
- [11] L. Zhang, T. Hu, Y. Min, G. Wu, J. Zhang, P. Feng, P. Gong, and J. Ye, "A taxi order dispatch model based on combinatorial optimization," in *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2017, p. 2151–2159. 2, 5
- [12] Z. Lou, W. Jie, and S. Zhang, "Multi-objective optimization for order assignment in food delivery industry with human factor considerations," *Sustainability*, vol. 12, no. 19, 2020. 2
- [13] M. Badreldin, A. Hussein, and A. Khamis, "A comparative study between optimization and market-based approaches to multi-robot task allocation," *Advances in Artificial Intelligence*, vol. 2013, p. 11, 09 2013. 2
- [14] R. Zlot and A. Stentz, "Market-based multirobot coordination for complex tasks," *The International Journal of Robotics Research*, vol. 25, no. 1, pp. 73–101, 2006. 2
- [15] A. Khamis, A. Elmogy, and F. Karray, "Complex task allocation in mobile surveillance systems," *Journal of Intelligent and Robotic Systems*, vol. 64, pp. 33–55, 10 2011. 2
- [16] M. B. Dias and A. Stentz, "A free market architecture for distributed control of a multirobot system," 2000. 2
- [17] M. 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. 2
- [18] W. Kmieciak, M. Wojcikowski, L. Koszalka, and A. Kasprzak, "Task allocation in mesh connected processors with local search meta-heuristic algorithms," in *Intelligent Information and Database Systems*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2010, pp. 215–224. 2
- [19] W.-H. Chen and C.-S. Lin, "A hybrid heuristic to solve a task allocation problem," *Computers and OR*, vol. 27, pp. 287–303, 03 2000. 2
- [20] D. K. Liu and A. K. Kulatunga, "Simultaneous planning and scheduling for multi-autonomous vehicles," in *Evolutionary Scheduling*, K. P. Dahal, K. C. Tan, and P. I. Cowling, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2007, pp. 437–464. 2
- [21] J. Holler, R. Vuorio, Z. Qin, X. Tang, Y. Jiao, T. Jin, S. Singh, C. Wang, and J. Ye, "Deep reinforcement learning for multi-driver vehicle dispatching and repositioning problem," in *2019 IEEE International Conference on Data Mining (ICDM)*. IEEE, 2019, pp. 1090–1095. 2, 4
- [22] O. de Lima, H. Shah, T.-S. Chu, and B. Fogelson, "Efficient ridesharing dispatch using multi-agent reinforcement learning," *arXiv preprint arXiv:2006.10897*, 2020. 2
- [23] M. Li, Z. Qin, Y. Jiao, Y. Yang, J. Wang, C. Wang, G. Wu, and J. Ye, "Efficient ridesharing order dispatching with mean field multi-agent reinforcement learning," in *The world wide web conference*, 2019, pp. 983–994. 2
- [24] N. Baras, A. Chatzisavvas, D. Ziouzos, and M. Dasygenis, "Improving automatic warehouse throughput by optimizing task allocation and validating the algorithm in a developed simulation tool," *Automation*, vol. 2, no. 3, pp. 116–126, 2021. [Online]. Available: <https://www.mdpi.com/2673-4052/2/3/7> 2
- [25] H. Ma, J. Li, T. Kumar, and S. Koenig, "Lifelong multi-agent path finding for online pickup and delivery tasks," *arXiv preprint arXiv:1705.10868*, 2017. 2, 5
- [26] K. Zhang, A. Koppel, H. Zhu, and T. Basar, "Global convergence of policy gradient methods to (almost) locally optimal policies," *SIAM Journal on Control and Optimization*, vol. 58, no. 6, pp. 3586–3612, 2020. 4
- [27] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, "Proximal policy optimization algorithms," *arXiv preprint arXiv:1707.06347*, 2017. 4
- [28] J. Schulman, P. Moritz, S. Levine, M. Jordan, and P. Abbeel, "High-dimensional continuous control using generalized advantage estimation," *arXiv preprint arXiv:1506.02438*, 2015. 4
- [29] L. Wang, Q. Cai, Z. Yang, and Z. Wang, "Neural policy gradient methods: Global optimality and rates of convergence," *CoRR*, vol. abs/1909.01150, 2019. [Online]. Available: <http://arxiv.org/abs/1909.01150> 4
- [30] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, "Attention is all you need," in *NeurIPS*, 2017. 4
- [31] F. A. Tillman and T. M. Cain, "An upperbound algorithm for the single and multiple terminal delivery problem," *Management Science*, vol. 18, no. 11, pp. 664–682, 1972. 5
- [32] S. K. X. Zheng, C. Tovey, R. Borie, P. Kilby, V. Markakis, and P. Keskinocak, "Agent coordination with regret clearing," in *AAAI*, 2008. 5
- [33] L. Cohen, T. Uras, and S. Koenig, "Feasibility study: Using highways for bounded-suboptimal multi-agent path finding," in *Eighth Annual Symposium on Combinatorial Search*, 2015. 5
- [34] H. Ma, D. Harabor, P. J. Stuckey, J. Li, and S. Koenig, "Searching with consistent prioritization for multi-agent path finding," in *AAAI*, vol. 33, no. 01, 2019, pp. 7643–7650. 5
- [35] L. Cohen, G. Wagner, D. Chan, H. Choset, N. Sturtevant, S. Koenig, and T. S. Kumar, "Rapid randomized restarts for multi-agent path finding solvers," in *Eleventh Annual Symposium on Combinatorial Search*, 2018. 6
- [36] M. Phillips and M. Likhachev, "Sipp: Safe interval path planning for dynamic environments," in *IEEE ICRA*, 2011, pp. 5628–5635. 5