

Ensembling Prioritized Hybrid Policies for Multi-agent Pathfinding

Huijie Tang^{*1}, Federico Berto^{*1}, Jinkyoo Park¹²

Abstract—Multi-Agent Reinforcement Learning (MARL) based Multi-Agent Path Finding (MAPF) has recently gained attention due to its efficiency and scalability. Several MARL-MAPF methods choose to use communication to enrich the information one agent can perceive. However, existing works still struggle in structured environments with high obstacle density and a high number of agents. To further improve the performance of the communication-based MARL-MAPF solvers, we propose a new method, Ensembling Prioritized Hybrid Policies (EPH). We first propose a selective communication block to gather richer information for better agent coordination within multi-agent environments and train the model with a Q-learning-based algorithm. We further introduce three advanced inference strategies aimed at bolstering performance during the execution phase. First, we hybridize the neural policy with single-agent expert guidance for navigating conflict-free zones. Secondly, we propose Q value-based methods for prioritized resolution of conflicts as well as deadlock situations. Finally, we introduce a robust ensemble method that can efficiently collect the best out of multiple possible solutions. We empirically evaluate EPH in complex multi-agent environments and demonstrate competitive performance against state-of-the-art neural methods for MAPF. We open-source our code at <https://github.com/ai4co/eph-mapf>.

I. INTRODUCTION

Multi-Agent Pathfinding (MAPF) involves finding collision-free paths for a group of agents while also aiming to minimize makespan or sum of costs [1]. MAPF has many practical applications, including warehouse robotics [2], aviation [3], and digital gaming [4] scenarios. However, the challenge intensifies with the realization that optimal solutions for MAPF are NP-hard, involving solving large-scale constraint satisfaction problems in combinatorial spaces [5], [6]. Classical algorithms to solve the MAPF problems such as heuristic centralized solvers typically invoke search-based algorithms, e.g., Conflict-Based Search (CBS) [7], as well as its improved versions ECBS [8] and EECBS [9]. However, as the number of agents increases, many heuristic solvers struggle to scale because of the high complexity of considering all agents in a centralized manner at once. In addition, in real-world applications, new tasks frequently occur during the execution of the initially planned paths [10]. Such centralized heuristics solvers have to re-plan paths for all the agents after seeing new tasks, making it burdensome to apply to real-world applications.

Multi-Agent Reinforcement Learning (MARL) based approaches [11], [12], [13], [14] offer another way to solve the

MAPF problem. Instead of treating MAPF as a centralized problem, they solve MAPF by regarding it as a sequential decision-making problem. One of the seminal MARL-based approaches to MAPF is PRIMAL [11]. In PRIMAL, the authors propose to learn a fully decentralized policy with partial agent-wise observation; the model is trained with MARL and imitation learning (IL). However, PRIMAL struggles to scale to high-obstacle-density environments with many agents. Also, PRIMAL assumes that agents move sequentially in each time step, which differs from real-world scenarios where agents move simultaneously in each time step. The improved version of PRIMAL, PRIMAL2 [12], also trains with MARL and IL. PRIMAL2 assumes the *disappear at target* [1] scenario. This assumption makes the problem easier because it means lower obstacle density when agents reach their goals and disappear, whereas in the *stay at target* [1] setting, agents constitute obstacles since they do not disappear after reaching their goals. Thus, the problem is much harder, especially when many agents exist. We adopt *stay at target* in this work.

Apart from the abovementioned MARL-based solvers [12], [11] that consider environments that do not allow communication, another work direction is MARL-based solvers with communication between agents when communication is possible, such as DHC [14] and DCC [13]. DHC is trained with Deep Q-learning [15] and adopts a similar setting as PRIMAL, where each agent has a partial observation. However, DHC is different in that it introduces *Graph Convolutional Communication*, which enables agents to communicate with other nearing agents for a better understanding of the environment, thus improving performance. Its improved version, DCC, improves communication by not learning broadcast but selective communication, thus reducing communication overhead. Recent MARL-based MAPF works like SCRIMP [16] and SACHA [17] also use communication mechanisms. For example, in SCRIMP, global communication is used to gather information from all agents in the environment, but this might be inefficient and burdensome in real-world applications.

Contributions. In this paper, we propose EPH (Ensembling Prioritized Hybrid Policies), a Q-learning-based MARL-MAPF solver with communication. Our communication scheme entails an enhanced selective communication block with improvements inspired by the latest Transformer variant to enable richer information extraction. Additionally, we introduce several strategies to be used during the inference phase to further bolster performance. Firstly, we propose Q value-based priority decisions, in which the Prioritized Conflict Resolution

^{*}Equal contributions

¹Department of Industrial and Systems Engineering, KAIST, South Korea

²OMELET

[‡]Authors are members of the AI4CO open research community.

Emails: {raylan.tang; fberto; jinkyoo.park}@kaist.ac.kr

decides the priority of agents involved in conflicts, and Advanced Escape Policy is responsible for breaking deadlocks. Secondly, we use hybrid expert guidance for agents that have no other *live agents* nearby. Lastly, we utilize ensembling to sample the best solutions from multiple solvers that run in parallel to leverage the different strengths of each strategy proposed. By improving communication capability and utilizing proposed strategies to help inference, we achieve competitive performance against state-of-the-art neural MARL-MAPF solvers.

II. RELATED WORKS

Search-based MAPF. Traditional centralized MAPF solvers often invoke search-based algorithms, and they can be categorized by their optimality. Optimal traditional heuristics solvers include Conflict-Based Search (CBS) [7] and its improved version [18], [19]. However, CBS relies on a constraint tree to find optimal solutions and can incur heavy computational overhead when the number of nodes in the constraint tree is high with many agents. The bounded-suboptimal solvers include CBS variants [9], [8]. These solvers can scale to environments with a larger number of agents while also guaranteeing the (sub-)optimality of generated solutions. However, their scalability is still hindered by their exponential time complexity [20]. Besides, when a new task or agent is added to the environment, such heuristics solvers have to re-plan the whole paths for all agents.

MARL-based MAPF. Recently, various approaches have emerged to address the MAPF problem using multi-agent reinforcement learning (MARL) techniques, building upon the foundation laid by PRIMAL [11]. These approaches exhibit diversity in their settings and strategies. In terms of the problem settings, some works consider scenarios with partial observation, where agents do not have complete knowledge of the environment [14], [13], [11], [12], while others operate under full observation, assuming agents possess complete information about the environment [21], [22]. Furthermore, the choice of RL algorithm varies across different works. Some employ Actor-Critic approaches [11], [12], [16], whereas others opt for value-based Q-Learning methods [14], [13], [22], [23]. In addition, certain approaches consider communication mechanisms among agents to enhance coordination [14], [13], [24], [25], [16]. These communication protocols allow agents to share information and collaborate effectively. Some recent MARL-based MAPF solvers also propose the use of techniques to help pathfinding during the inference phase. Gao et al. [26] propose an escape policy for structured environments that helps neural solvers escape from deadlock scenarios, which forces agents to take random actions attempting to break the deadlock. Such inference techniques can improve the performance of neural solvers under certain predefined conditions [27].

III. PROBLEM FORMULATION

A. MAPF Definition

The MAPF problem involves finding a set of collision-free paths for multiple agents, each with its own distinct starting location and a unique goal location, within a map containing obstacles. The map M is an undirected graph $M = (V, E)$, with some vertices being inaccessible obstacles $V_o \subset V$. An obstacle collision occurs when agent i reaches the location of an obstacle. Given an agent set N , each agent $i \in N$ has a starting vertex $v_i^0 \in V$ and a goal vertex $v_i^g \in V$. We define the observation space O_i and action space A_i for agent i , which is, at each time step $t = 0, \dots, t_{max}$, each agent i located in $v_i^t \in V$ has an observation of the map $o_i^t \in O_i$, decides on a single action $a_i^t \in A_i$, and move to v_i^{t+1} . t_{max} is the maximal allowed time to finish the MAPF problem. A vertex collision occurs when agent i and agent j reach the same vertex v at the same time, step t ; an edge conflict occurs when agent i and agent j traverse through the same edge (u, v) in the opposite direction. A MAPF solution exists if and only if for any agent $i \in N$, there exists $0 \leq t \leq t_{max}$ such that $a_i^t(\dots a_i^1(a_i^0(v_i^0))) = v_i^{t+1} = v_i^g$, where $a_i^0(v_i^0) = v_i^1, \dots, a_i^t(v_i^t) = v_i^{t+1} = v_i^g$; and for any agent $i \in N$, there are no collisions.

B. Environment Settings

Following the conventions of MAPF, we use a 2D grid world to represent the environment. Each cell in the $n \times n$ grid world can be either an empty cell or an obstacle that stops agents from passing¹. Each agent occupies an empty cell, and the task for each agent is moving from its starting cell to the goal cell without collisions. At the beginning of each episode, m starting cells and m goal cells are randomly selected among the empty cells for m agents, and we make sure that there is no overlap among $2m$ selected cells. At each timestep, the agents choose to move either up, down, left, right, or stay still. We consider four types of conflicts: *edge conflict* [1], *vertex conflict* [1], conflict with static obstacles, and out-of-bound conflict. These four types of conflicts are not allowed during pathfinding.

We use a partially observable environment. At each time step, each agent has a limited Field of View (FOV) that consists of six channels: four channels are the heuristic channels described in [14], the remaining two channels are used to represent the locations of all agents within the FOV, and the location of obstacles within the FOV, respectively.

IV. EPH: ENSEMBLING PRIORITIZED HYBRID POLICIES

We propose EPH and describe it in two sections, each with several subsections. **Section IV-A** describes how we formulate MAPF problem as a Q-Learning based MARL problem with communication and how we train such a model. **Section IV-B** gives an introduction to the advanced inference strategies that we propose to improve performance during execution. **Figure 1** shows the overall architecture of EPH.

¹We consider value 1 on map M as an obstacle and 0 as available.

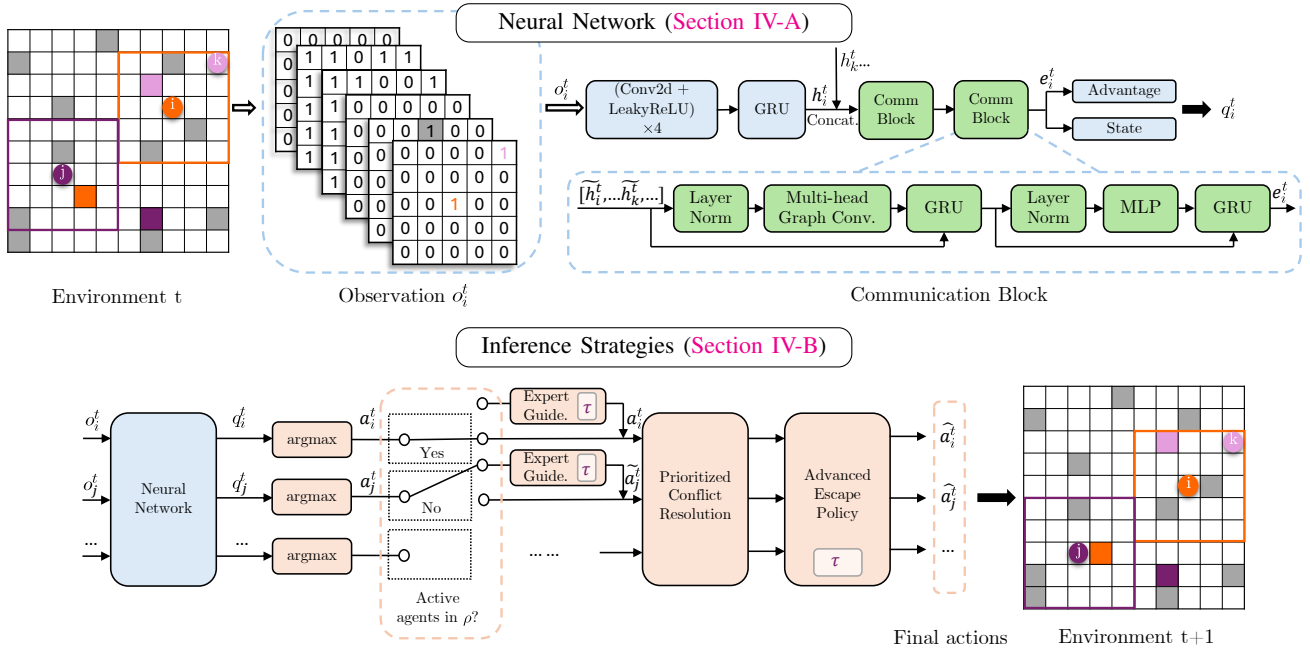


Fig. 1: Overview of a single inference step of EPH. The upper part is the neural network structure of EPH; the lower part is the illustration of how observations of all agents are transformed into actions by first feeding into the neural network, and then going through the proposed inference strategies. The ρ and τ in the lower part are the hyperparameters defined in Section IV-B.

A. Q-Learning based MARL with Communication for MAPF

In this section, we first introduce our model with the new communication block, and then we introduce the model training via Double Dueling Deep Q Networks (D3QN).

1) *Graph Convolution based Communication*: We show our communication architecture in Figure 1. Our communication structure is adapted from selective communication block in DCC [13]. Compared with ordinary communication, selective communication can effectively reduce communication time compared with non-selective counterparts, it can also selectively gather more informative and important hidden information from other agents within FOV, which is shown to be beneficial for pathfinding [13]. However, in the previous work, the selective graph convolution communication block did not include position-wise processing or normalization, and it is shown that including them helps stabilize training and gives better scalability [28]. Inspired by *GTrXL*-styled transformer in [28], we re-design the communication block to enable better coordination among agents by including an extra multi-layer perceptron layer, GRU block, and layer normalization, which is shown to be beneficial in many works across different fields [29], [16], [28]. The generation method of communication mask for selective communication, as well as the request-reply communication scheme that consists of two cascaded communication blocks, are kept the same as shown in DCC to reduce communication overhead and enrich the information gathered.

2) *Training of Q-learning based MARL with Communication*: We train the model with D3QN (Double Dueling Deep Q Network), which incorporates Deep Q-Learning Network (DQN) with both the Dueling DQN architecture [30], sep-

arating parameters for advantage and value estimation, and Double DQN [31], which uses a target network to estimate the value to ameliorate the effects of Q value overestimation [32]. The Q value for agent i is obtained via:

$$Q_{s,a}^i = Val_s(e_i^t) + Adv(e_i^t)_a - \frac{1}{|\mathcal{A}|} \sum_{a'} Adv(e_i^t)_{a'} \quad (1)$$

where $Val(\cdot)$ and $Adv(\cdot)$ represents state advantage and action advantage function, respectively. e_i^t is the final output of communication for agent i containing information from other agents whom i chooses to communicate with. \mathcal{A} is the action space. The loss for training EPH is expressed as:

$$\mathcal{L}(\theta) = \text{MSE}(R_t^i - Q_{s_t, a_t}^i(\theta)) \quad (2)$$

where MSE is the mean square error. We have $R_t^i = r_t^i + \gamma r_{t+1}^i + \dots + \gamma^n Q_{s_{t+n}, a_{t+n}}^i(\theta)$, where r_t^i is the reward received by agent i at time t , and θ is the parameter for target network which is updated to align with online parameter θ at a predefined interval. The reward structure for D3QN is taken from DHC [14] as shown in the Table I. Negative rewards are given to agents who don't reach the goal to expedite goal reaching in the shortest distance.

TABLE I: Reward Structure

Actions	Rewards
➡ Move and Stay off Goal	-0.075
✓ Stay on Goal	0
✘ Collision	-0.5
★ Reach Goal	3

In MAPF, the roles of each agent are identical: every agent has a unique start location and a goal location, and each of

them is expected to move in a collision-free manner and reach the goal eventually. Given this observation, instead of training multiple policies for multiple agents, it's more natural to adopt parameter sharing and train a single policy from a single agent's perspective. Although we adopt parameter sharing, it should be emphasized that the trained policy can be used for multi-agent pathfinding. This is because each agent perceives unique information from its own partial observations and inter-agent communication outcome and, consequently, takes different actions.

B. Advanced Inferences Strategies for EPH

We propose three inference techniques to further bolster pathfinding performance during inference. In Section IV-B.1, we hybridize the policy with single-agent optimal A^* paths to give guidance to agents that do not have any other *live agents* nearby. In Section IV-B.2, we introduce two Q value-based priority decision-making strategies, namely *Prioritized Conflict Resolution* for efficient conflict resolution and *Advanced Escape Policy* for priority-based avoidance of deadlocks. Finally, Section IV-B.3 introduces the ensembling, which runs multiple policies in parallel and samples the best possible solutions in the solution space.

1) *Hybrid Expert Guidance*: When agent communication is sparse, such as in scenarios where no other agents are located within the agent's field of view, it is beneficial to incorporate low-cost single-agent expert path to guide decision-making, thus hybridizing (EPH stands for Hybrid). We hereby define a *live agent* as an agent that is currently off its goal, i.e., $v_i^t \neq v_i^g$. We propose leveraging the optimal A^* path as expert guidance during inference for an agent that has no other *live agents* within a square centered on the agent itself; the length of the square is $2\rho + 1$ where ρ is the visibility radius for *live agents*. We adapt the efficient A^* , a well-established low-cost algorithm for efficient single-agent pathfinding, to multi-agent settings as follows:

$$a_i^{*t}, a_i^{*t+1}, \dots, a_i^{*T} = A_\tau^*(v_i^t, v_i^g, M, V^t, V^g) \quad (3)$$

where $a_i^{*t}, a_i^{*t+1}, \dots, a_i^{*T}$ is the sequence of actions that leads agent i to the goal and $\tau \in \{0, 1, 2\}$ is the A^* type. $V^t = \{v_1^t, \dots, v_i^t, \dots, v_m^t\}$ and $V^g = \{v_1^g, \dots, v_i^g, \dots, v_m^g\}$ is the current position set and goal set, respectively. We formulate each type τ as follows:

- i) A_0^* : is the classic A^* that takes the current map M for avoiding obstacles ensuring finding the optimal single-agent path.
- ii) A_1^* : treats *all agents as obstacles*, except the current agent i ; i.e., $M(v_j^t) \leftarrow 1 \forall j, \text{ and } j \neq i$. All other agents are considered temporary obstacles within the A^* calculation. This ensures collision avoidance.
- iii) A_2^* : this version considers *only inactive agents as obstacles*, i.e., $M(v_j^t) \leftarrow 1 \forall j \text{ s.t. } v_j^t = v_j^g$. This version tries to avoid agents already at goal only. This can generally obtain better paths compared to A_1^* , since agents that are still moving do not interfere with optimal path calculations.




Fig. 2: Types of A_τ^* we consider. $\tau \in \{0, 1, 2\}$. Changing the map representation improves the single-agent path based on each situation.

The illustration of three types of A^* is shown in Figure 2. We find that using different A^* types can be beneficial for pathfinding in different scenarios.

2) *Value-based Priority Decisions*: We further introduce two techniques for resolving conflict situations and deadlock scenarios based on priorities (EPH stands for Prioritized) provided by the Q values obtained by the trained neural model. This is inspired by priority-based planning methods [33], [34], [35] and, in particular, Priority-Based Search (PBS) [36], which first decides the priority for each agent and then plans individual paths for each agent in the order of priority. In PBS, the agent with lower priority is not allowed to collide with the agent with higher priority. In our Q value-based priority decisions, an agent with a lower Q value will be assigned a lower priority and will have to give way to another agent with a higher priority. Intuitively, higher Q-values correspond to higher expected rewards.

a) *Prioritized Conflict Resolution*: During pathfinding, agents may still take invalid actions, which leads to conflicts. Compared with previous works [14], [13] that recursively recover the states of agents that are involved in conflicts, we use several strategies to reduce collisions. When we first sample actions from the Q value matrix generated by our network, we filter out actions that make agents collide with static obstacles by masking the corresponding value in the Q value matrix. After we get static-obstacle-collision-free action for each agent, we let all agents execute their actions. If collisions between agents happen, we decide the priority of agents involved in agent collision by giving the highest priority to the agent with the highest Q value, and re-choosing actions for other agents with smaller Q values with their previous actions masked. The illustration of Prioritized Conflict Resolution is shown in Figure 3. Since Q value represents the goodness and preference of an agent choosing the corresponding action, it is natural to decide the priority of agents based on Q value.

b) *Advanced Escape Policy*: Gao et al. [26] first proposed an escape policy to avoid deadlock situations, a common case in structured environments. A deadlock can occur when an agent is off its goal and oscillates back and forth between two adjacent cells (Algorithm 1, lines 9-10). However, in the original escape policy in [26], agents choose to break deadlocks just by taking random actions when deadlocks are detected, which may be suboptimal and lead to

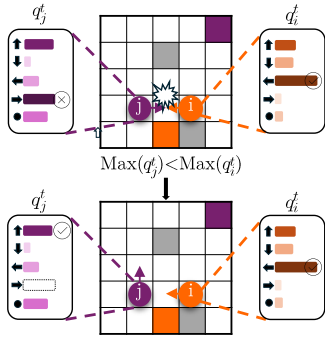


Fig. 3: Prioritized Conflict Resolution. Agents with higher Q values are prioritized when a conflict happens, which leads to shorter paths.

unnecessary randomness². In our Advanced Escape Policy, we use a two-stage approach to break the deadlocks. Firstly, we assign a priority to each agent by sorting the Q values in descending order. Then, for each agent in such order, we utilize A_τ^* mentioned in Section IV-B.1 to obtain the action. If A_τ^* fails to find a path because of other agents involved in the deadlock, we choose the action from the highest *valid* Q value, where *invalid* Q values (corresponding to actions that lead to a collision) are masked, i.e. set to $-\infty$. We also update a copy of the map such that the current next action is masked, ensuring there will be no conflict. Algorithm 1 shows our proposed Advanced Escape Policy.

3) *Ensembling*: In the inference phase, $\underline{\text{EPH}}$ (where $\underline{\text{E}}$ stands for the ensembling) employs an ensemble technique that operates by sampling best solutions from solvers with different settings that run in parallel, without engaging in adaptive learning. Each of the solvers has its own A^* type, either A_0^* , A_1^* or A_2^* , with different ρ . Some of the solvers have the value-based priority decision mechanism enabled, while others don't. This approach is designed to leverage the complementary strengths of each strategy proposed to navigate complex multi-agent pathfinding. We can divide the ensembling technique into two stages:

- i) *Parallel Solver Execution*: For every episode, $\underline{\text{EPH}}$ executes multiple solvers with different settings in parallel, each of them has different techniques enabled with different parameters. We treat each solution from one solver in the parallel execution as one independent pathfinding solution. Parallel execution allows for a comprehensive exploration of potential solutions, maximizing the likelihood of identifying an optimal path.
- ii) *Best Solution Sampling*: Upon completion of all parallel solvers, $\underline{\text{EPH}}$ evaluates each independent solution from one solver based on its *makespan* [1], i.e., time steps required to finish the single test episode. For each episode, we take as a final solution the one with the lowest *makespan* among all parallel solvers.

This ensemble technique enhances $\underline{\text{EPH}}$'s capability to solve MAPF problems by providing a robust framework for selecting the most effective pathfinding solver from a set of pre-defined options. By systematically sampling the

²With time $\rightarrow \infty$, a series of random actions should also guarantee a solution at the expense of solving time.

Algorithm 1: Advanced Escape Policy

Input: Goals set V^g ; agents' current positions set V^t ; position set for each agent for the past five time steps $V_i^p = \{v_i^t, v_i^{t-1}, v_i^{t-2}, v_i^{t-3}, v_i^{t-4}\}$; Q value matrix $Q = (q_1^t \cdots q_i^t \cdots q_m^t)$, $\forall i = 1, \dots, m$; initial map M where 1 is an obstacle and 0 is available; A^* type τ .

Output: Actions $A = \{a_1^t, \dots, a_i^t, \dots, a_m^t\}$

- 1 **Function** MaskQValues (q_i^t, M, v_i^t):
 - // mask Q values of obstacles
 - 2 **for** a **in** \mathcal{A} **do**
 - 3 **if** $M[a(v_i^t)] = 1$ **then**
 - 4 $q_i^t(a) \leftarrow -\infty$;
 - 5 **return** q_i^t ;
 - // Init actions based on max Q values
 - 6 $A^{\text{init}} = \{a_1^t, \dots, a_i^t, \dots, a_m^t\} \leftarrow \arg \max Q$;
 - 7 $Idx \leftarrow \text{argsort}(Q_{A^{\text{init}}}, \text{descending}=\text{True})$; // sort indexes by Q values (first is highest)
 - 8 **for** i **in** Idx **do**
 - // check if active and has deadlock
 - 9 **if** $v_i^t \neq v_i^g$ **then**
 - 10 **if** $v_i^{t-1} = v_i^{t-3}$ & $v_i^{t-2} = v_i^{t-4}$ **then**
 - // if A_τ^* path found, use A_τ^*
 - 11 **if** $\exists A_\tau^*(v_i^t, v_i^g, M, V^t, V^g)$ **then**
 - 12 $a_i^t \leftarrow A_\tau^*(v_i^t, v_i^g, M, V^t, V^g)[0]$
 - 13 **else**
 - 14 $q_i^t \leftarrow \text{MaskQValues}(q_i^t, M, v_i^t)$;
 - 15 $a_i^t \leftarrow \arg \max q_i^t$
 - 16 $M[a_i^t(v_i^t)] \leftarrow 1$ // set next pos. as obstacle

best solutions from multiple solvers running in parallel, $\underline{\text{EPH}}$ efficiently navigates the complexities of multi-agent environments, ensuring high performance and reliability in the inference phase.

V. EXPERIMENTS

A. Training and Testing Settings

We use D3QN with a prioritized experience replay buffer [37] for training $\underline{\text{EPH}}$. Curriculum learning [38], which is shown to be effective in previous works [14], [17], is also adopted. The curriculum learning starts from one agent in 10×10 map, and ends on 40×40 map with 16 agents. Every time the success rate of the model under training reaches 0.9 for the current training environment, we increase the difficulty of the map. Thus, experience from more complex environments will be added to the buffer, and the model learns from more and more complex scenarios gradually. The obstacle density of maps in training is sampled from a triangular distribution between 0 and 0.5 with a peak of 0.33. The FOV size in our work is 9×9 to align with previous works [14], [13]. The maximum episode length for training episodes is 256. We save model checkpoints every 1000 training steps and take the best one as the final model based on performance on a random-map validation set. It takes 20 hours to perform 150k training steps on a server with a single NVIDIA RTX A6000 and AMD EPYC 7542 (128)






Fig. 4: Comparative Analysis of Success Rate (left) and Average Episode Length (right) on random maps.

TABLE II: Solution quality for different MAPF solvers. We report episode length (EL, lower is better \downarrow) and success rate (SR, higher is better \uparrow).

		Heuristics Solvers								Neural Solvers										
		CBS		ODrM*		wPBS		PRIMAL		DHC		DCC		SACHA		SCRIMP		EPH		
Map	m	EL	SR	EL	SR	EL	SR	EL	SR	EL	SR	EL	SR	EL	SR	EL	SR	EL	SR	
	den312d	4	51.74	100%	51.76	100%	69.32	92%	196.54	40%	86.56	100%	82.99	100%	81.43	100%	82.34	100%	79.05	100%
	8	55.50	100%	78.74	88%	116.32	70%	245.02	8%	100.70	100%	97.95	99%	89.73	100%	99.58	100%	91.42	100%	
	16	118.97	68%	186.44	34%	208.28	24%	256.00	0%	109.24	100%	108.29	97%	96.74	100%	105.78	100%	104.87	100%	
	64	251.86	2%	256.00	0%	248.06	4%	256.00	0%	124.38	98%	119.15	97%	104.30	98%	115.39	100%	110.78	100%	
	warehouse	4	77.79	100%	77.79	100%	104.41	94%	355.80	42%	146.12	99%	135.89	99%	134.59	99%	197.79	87%	134.56	100%
	8	83.48	100%	100.37	96%	170.46	80%	451.82	18%	198.82	91%	169.50	96%	166.72	93%	304.22	66%	151.94	100%	
	16	81.64	100%	133.59	88%	340.18	40%	492.04	8%	281.37	74%	208.72	90%	198.72	76%	366.98	48%	164.05	100%	
	64	262.15	58%	417.22	22%	512.00	0%	505.58	4%	432.28	28%	335.81	58%	354.33	48%	451.40	21%	176.35	100%	
		64	494.93	4%	512.00	0%	512.00	0%	512.00	0%	512.00	1%	473.92	14%	437.29	28%	504.26	4%	189.58	100%

@2.9GHz. We also retrain DCC with the same setting for fair comparison.

We evaluate EPH on random maps and structured environments. For random maps, we test on 40×40 and 80×80 maps with 0.3 obstacle density; the maximum allowed time step for them is 256 and 386, respectively. The set of parameters τ and ρ used by ensemble for random maps are selected from a Cartesian product of the spaces $\{0, 1, 2\} \times \{2, 3, 4, 5\}$. For structured environments, we use the benchmark from [1] for the test. Specifically, we choose the Dragon Age Origins map den312d (65×81) and warehouse map (161×63) from the benchmark for the test to showcase EPH’s performance in complex structured environments. The maximum allowed time step for den312d and warehouse is 256 and 512, respectively. The set of parameters τ and ρ used by ensemble for structured maps are selected from a Cartesian product of the spaces $\{0, 1, 2\} \times \{3, 4\}$. For random maps, we report the success rate and episode length averaged across 100 test instances for each case. For structured maps, we average across 300 to align with the setting in one of our baseline SACHA [17]. The success rate is the percentage of test instances fully solved, and the episode length is the *makespan* [1] of an instance.

B. Results on Random Maps

We choose DHC, DCC, and current state-of-the-art MARL-MAPF solver SCRIMP [16] as our baselines. DHC and DCC both utilize local communication to gather information from other agents nearby to give the model more information and understanding beyond the agent’s own FOV, and both use Q-learning. While DHC chooses to

communicate with all other agents nearby, DCC selectively communicates with them. SCRIMP uses Transformer-based global communication to even gather more information from all agents in the environment to enrich the information one agent can receive. Imitation learning is also utilized to help SCRIMP learn from expert demonstrations from multi-agent solvers during training alongside PPO [39].

Figure 4 (left) shows the success rate of EPH compared with baselines. On 40×40 map, EPH beats all the baselines in terms of success rate, reaching 100% success rate for all the cases; on the 80×80 map, EPH outperforms both DHC and DCC in all cases, while falling behind SCRIMP by 1% on 32 and 64 agents scenarios. Figure 4 (right) illustrates the average episode length, which indicates the solution quality, of all four models. EPH outperforms both DHC and DCC in terms of average episode length in all cases. When compared against SCRIMP, EPH can outperform SCRIMP on larger 80×80 map. On 40×40 map, EPH can achieve a shorter average episode length than SCRIMP except in the 64-agent case. It should be noted that SCRIMP uses global communication, which could give the model more information than we could. However, global communication is burdensome and inefficient in real life; also, global communication may be limited by geographical constraints. By contrast, EPH, which achieves better performance than baselines in most cases, only uses improved local selective communication. EPH equipped with local selective communication is information-efficient, less burdensome, and practical in real-life scenarios.

C. Results on Structured Maps

We additionally test on structured maps that were never seen during training to showcase the performance of baselines and EPH in terms of scalability and generalization to unseen scenarios. For structured maps, apart from the abovementioned baselines, we further include three additional heuristic solvers: CBS [7], wPBS [10] and ODrM* [40]; as well as two neural solvers: PRIMAL [11] and the recent SACHA [17] as baselines. SACHA is a neural MAPF solver trained with Soft Actor-Critic and employs optional global communication blocks. We choose to compare with SACHA equipped with global communication block for fair comparison. For the heuristics solvers, the runtime limit is 120 seconds for CBS and wPBS, and 20 seconds for ODrM*, same as the settings used in SACHA. EPH takes a slightly higher runtime compared to DCC. At the same time, conflict checks and running A* have a negligible effect on runtime³.

Table II shows the performance of different solvers in two structured environments. In den312d map, EPH beats all neural baselines in terms of success rate. For the average episode length, we outperform all neural baselines except SACHA in den312d map with 8 to 32 agents. It is noticeable that though SACHA achieves shorter episode length in these cases, it fails to achieve a higher success rate when the number of agents increases. By contrast, EPH has better scalability and achieves 100% success rate in den312d map with 64 agents. In the warehouse map, which is a common setup in real-world scenarios, EPH excels all neural baselines in all cases in terms of both metrics, showcasing the practicality of our method in real-world applications. It is noticeable that heuristic solvers generally offer better solutions in cases with few agents. However, the scalability issue, which is common for heuristics solvers, prevents them from reaching better results when the agent number increases.

D. Ablation Studies

TABLE III: Effectiveness of EPH components with a large number of agents $m = 64$. *Base* is the model that only has improved communication.

Method	den312d		warehouse	
	EL	SR	EL	SR
Base	140.70	97%	479.42	15%
+ Prioritized Decisions	129.84	100%	454.16	22%
+ Hybrid Guidance	130.16	100%	255.31	88%
+ Ensemble	121.66	100%	189.58	100%

In Table III, we showcase the effectiveness of the Q value-based priority decision-making strategy, hybrid expert guidance, and ensembling outlined in Section IV-B. *Base* is the model that only has the proposed improved communication block. We use A_2^* in the relative parts of value-based priority decision-making strategy and hybrid expert

³We note that we did not parallelize ensembling in the implementation; parallelizing ensembling would significantly speed up the algorithm. We leave this implementation for future work.

guidance. In den312d, it can be seen that all three inference strategies help achieve better performance compared with *Base*. In warehouse, the same trend is preserved. However, we observe that different strategies have different improvements in different scenarios. For example, hybrid guidance does not offer much improvement in maps that are not highly congested and structured, such as den312d; but in highly structured environments like warehouse, it offers significant improvements. This signals that ensembling to take advantage of different solvers would be the best choice to give the model the best performance. In addition, it is also noticeable that the *Base* outperforms DCC in both structured maps, illustrating the effectiveness of our improved communication scheme.

We also conduct an ablation study regarding the impact of A^* types of the hybrid expert guidance, i.e., the τ , on EPH’s performance. We observe that in den312d map with 64 agents, different τ only leads to 4% of difference in success rate. However in warehouse map with 64 agents, the difference jumps to 79%, and the model with hybrid expert guidance with A_0^* only achieves 18% success rate. This signifies again that different strategies have their own applicability and the effectiveness of ensembling.

VI. CONCLUSIONS

In this paper, we proposed EPH, a novel learning approach for Multi-Agent Path Finding (MAPF). Our approach employs an enhanced selective communication scheme and, at inference time, efficiently samples multiple solutions by ensembling solvers with configurations of neural value-based priorities for resolving conflicts, advanced escape policy for deadlock avoidance, and hybrid expert guidance. EPH demonstrated competitive performance against state-of-the-art neural MAPF baselines in complex environments.

We finally define some future works that may arise from EPH. Firstly, training with other RL algorithms such as on-policy algorithms [41], [39], [42], where priorities may be assessed via values from the critic network, could further boost the performance. Importantly, better hybridization techniques with existing inexpensive low-level solvers may help in particular by avoiding deadlock situations arising in highly-structured environments [12]. Ultimately, we believe hybridizing learned neural solvers with their classical counterparts, as has been suggested in other combinatorial domains for both obtaining better solutions [43], [44], [45], [46] and generating better heuristics [47], [48], is a promising research avenue for scalable and generalizable MAPF.

ACKNOWLEDGMENT

We thank Qiushi Lin for providing us with help for the performance results of baselines in structured maps.

REFERENCES

- [1] R. Stern, N. Sturtevant, A. Felner, S. Koenig, H. Ma, T. Walker, J. Li, D. Atzmon, L. Cohen, T. Kumar *et al.*, “Multi-agent pathfinding: Definitions, variants, and benchmarks,” in *Proceedings of the International Symposium on Combinatorial Search*, vol. 10, no. 1, 2019, pp. 151–158.

- [2] P. R. Wurman, R. D'Andrea, and M. Mountz, "Coordinating hundreds of cooperative, autonomous vehicles in warehouses," *AI magazine*, vol. 29, no. 1, pp. 9–9, 2008.
- [3] R. Morris, C. S. Pasareanu, K. S. Luckow, W. Malik, H. Ma, T. S. Kumar, and S. Koenig, "Planning, scheduling and monitoring for airport surface operations," in *AAAI Workshop: Planning for Hybrid Systems*, 2016, pp. 608–614.
- [4] H. Ma, J. Yang, L. Cohen, T. Kumar, and S. Koenig, "Feasibility study: Moving non-homogeneous teams in congested video game environments," in *AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*, 2017.
- [5] J. Yu and S. LaValle, "Structure and intractability of optimal multi-robot path planning on graphs," in *AAAI Conference on Artificial Intelligence*, 2013.
- [6] J. Banfi, N. Basilico, and F. Amigoni, "Intractability of time-optimal multirobot path planning on 2d grid graphs with holes," *IEEE Robotics and Automation Letters*, vol. 2, no. 4, pp. 1941–1947, 2017.
- [7] G. Sharon, R. Stern, A. Felner, and N. R. Sturtevant, "Conflict-based search for optimal multi-agent pathfinding," *Artificial Intelligence*, vol. 219, pp. 40–66, 2015.
- [8] M. Barer, G. Sharon, R. Stern, and A. Felner, "Suboptimal variants of the conflict-based search algorithm for the multi-agent pathfinding problem," in *Proceedings of the International Symposium on Combinatorial Search*, vol. 5, no. 1, 2014, pp. 19–27.
- [9] J. Li, W. Ruml, and S. Koenig, "Eecbs: A bounded-suboptimal search for multi-agent path finding," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 35, no. 14, 2021, pp. 12 353–12 362.
- [10] J. Li, A. Tinka, S. Kiesel, J. W. Durham, T. S. Kumar, and S. Koenig, "Lifelong multi-agent path finding in large-scale warehouses," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 35, no. 13, 2021, pp. 11 272–11 281.
- [11] G. Sartoretti, J. Kerr, Y. Shi, G. Wagner, T. S. Kumar, S. Koenig, and H. Choset, "Primal: Pathfinding via reinforcement and imitation multi-agent learning," *IEEE Robotics and Automation Letters*, 2019.
- [12] M. Damani, Z. Luo, E. Wenzel, and G. Sartoretti, "Primal 2: Pathfinding via reinforcement and imitation multi-agent learning-lifelong," *IEEE Robotics and Automation Letters*, 2021.
- [13] Z. Ma, Y. Luo, and J. Pan, "Learning selective communication for multi-agent path finding," *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 1455–1462, 2021.
- [14] Z. Ma, Y. Luo, and H. Ma, "Distributed heuristic multi-agent path finding with communication," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2021, pp. 8699–8705.
- [15] V. Mnih, K. Kavukcuoglu, D. Silver, A. Graves, I. Antonoglou, D. Wierstra, and M. Riedmiller, "Playing atari with deep reinforcement learning," *arXiv preprint arXiv:1312.5602*, 2013.
- [16] Y. Wang, B. Xiang, S. Huang, and G. Sartoretti, "SCRIMP: Scalable communication for reinforcement-and imitation-learning-based multi-agent pathfinding," *arXiv preprint arXiv:2303.00605*, 2023.
- [17] Q. Lin and H. Ma, "SACHA: Soft actor-critic with heuristic-based attention for partially observable multi-agent path finding," *IEEE Robotics and Automation Letters*, 2023.
- [18] G. Gange, D. Harabor, and P. J. Stuckey, "Lazy cbs: implicit conflict-based search using lazy clause generation," in *ICAPS*, 2019.
- [19] J. Li, G. Gange, D. Harabor, P. J. Stuckey, H. Ma, and S. Koenig, "New techniques for pairwise symmetry breaking in multi-agent path finding," in *Proceedings of the International Conference on Automated Planning and Scheduling*, vol. 30, 2020, pp. 193–201.
- [20] J. Chung, J. Fayyad, Y. A. Younes, and H. Najjaran, "Learning to team-based navigation: A review of deep reinforcement learning techniques for multi-agent pathfinding," *arXiv preprint arXiv:2308.05893*, 2023.
- [21] Z. He, L. Dong, C. Sun, and J. Wang, "Asynchronous multithreading reinforcement-learning-based path planning and tracking for unmanned underwater vehicle," *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, vol. 52, no. 5, pp. 2757–2769, 2021.
- [22] D. Wang, H. Deng, and Z. Pan, "Mrcdrl: Multi-robot coordination with deep reinforcement learning," *Neurocomputing*, 2020.
- [23] L. Chen, Y. Wang, Y. Mo, Z. Miao, H. Wang, M. Feng, and S. Wang, "Multiagent path finding using deep reinforcement learning coupled with hot supervision contrastive loss," *IEEE Transactions on Industrial Electronics*, vol. 70, no. 7, pp. 7032–7040, 2022.
- [24] W. Li, H. Chen, B. Jin, W. Tan, H. Zha, and X. Wang, "Multi-agent path finding with prioritized communication learning," in *2022 International Conference on Robotics and Automation (ICRA)*. IEEE, 2022, pp. 10 695–10 701.
- [25] L. Chen, Y. Wang, Z. Miao, Y. Mo, M. Feng, and Z. Zhou, "Multi-agent path finding using imitation-reinforcement learning with transformer," in *2022 IEEE International Conference on Robotics and Biomimetics (ROBIO)*. IEEE, 2022, pp. 445–450.
- [26] J. Gao, Y. Li, X. Yang, and M. Tan, "Rde: A hybrid policy framework for multi-agent path finding problem," *arXiv preprint arXiv:2311.01728*, 2023.
- [27] H. Tang, F. Berto, Z. Ma, C. Hua, K. Ahn, and J. Park, "HiMAP: Learning heuristics-informed policies for large-scale multi-agent pathfinding," in *AAMAS*, 2024.
- [28] E. Parisotto, F. Song, J. Rae, R. Pascanu, C. Gulcehre, S. Jayakumar, M. Jaderberg, R. L. Kaufman, A. Clark, S. Noury *et al.*, "Stabilizing transformers for reinforcement learning," in *International conference on machine learning*. PMLR, 2020, pp. 7487–7498.
- [29] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, "Attention is all you need," *Advances in neural information processing systems*, vol. 30, 2017.
- [30] Z. Wang, T. Schaul, M. Hessel, H. Hasselt, M. Lanctot, and N. Freitas, "Dueling network architectures for deep reinforcement learning," in *International conference on machine learning*. PMLR, 2016.
- [31] H. Van Hasselt, A. Guez, and D. Silver, "Deep reinforcement learning with double q-learning," in *AAAI conference on artificial intelligence*, 2016.
- [32] H. Hasselt, "Double q-learning," *Advances in neural information processing systems*, vol. 23, 2010.
- [33] J.-C. Latombe, *Robot motion planning*. Springer Science & Business Media, 2012, vol. 124.
- [34] D. Silver, "Cooperative pathfinding," in *AAAI conference on artificial intelligence and interactive digital entertainment*, 2005.
- [35] K. Okumura, M. Machida, X. Défago, and Y. Tamura, "Priority inheritance with backtracking for iterative multi-agent path finding," *Artificial Intelligence*, vol. 310, p. 103752, 2022.
- [36] H. Ma, D. Harabor, P. J. Stuckey, J. Li, and S. Koenig, "Searching with consistent prioritization for multi-agent path finding," in *AAAI conference on artificial intelligence*, 2019.
- [37] T. Schaul, J. Quan, I. Antonoglou, and D. Silver, "Prioritized experience replay," *arXiv preprint arXiv:1511.05952*, 2015.
- [38] Y. Bengio, J. Louradour, R. Collobert, and J. Weston, "Curriculum learning," in *Proceedings of the 26th annual international conference on machine learning*, 2009, pp. 41–48.
- [39] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, "Proximal policy optimization algorithms," *arXiv preprint arXiv:1707.06347*, 2017.
- [40] C. Ferner, G. Wagner, and H. Choset, "Odrm* optimal multirobot path planning in low dimensional search spaces," in *ICRA*. IEEE, 2013.
- [41] R. J. Williams, "Simple statistical gradient-following algorithms for connectionist reinforcement learning," *Machine learning*, 1992.
- [42] F. Berto, C. Hua, J. Park, L. Luttmann, Y. Ma, F. Bu, J. Wang, H. Ye, M. Kim, S. Choi, N. G. Zepeda, A. Hottung, J. Zhou, J. Bi, Y. Hu, F. Liu, H. Kim, J. Son, H. Kim, D. Angioni, W. Kool, Z. Cao, J. Zhang, K. Shin, C. Wu, S. Ahn, G. Song, C. Kwon, L. Xie, and J. Park, "RL4CO: an Extensive Reinforcement Learning for Combinatorial Optimization Benchmark," *arXiv preprint arXiv:2306.17100*, 2024, <https://github.com/ai4co/rl4co>.
- [43] A. Hottung and K. Tierney, "Neural large neighborhood search for the capacitated vehicle routing problem," *arXiv preprint arXiv:1911.09539*, 2019.
- [44] W. Kool, L. Bliet, D. Numeroso, Y. Zhang, T. Catshoek, K. Tierney, T. Vidal, and J. Gromicho, "The euro meets neurips 2022 vehicle routing competition," in *NeurIPS 2022 Competition Track*. PMLR, 2022, pp. 35–49.
- [45] H. Ye, J. Wang, H. Liang, Z. Cao, Y. Li, and F. Li, "Glop: Learning global partition and local construction for solving large-scale routing problems in real-time," *AAAI 2024*, 2024.
- [46] H. Ye, J. Wang, Z. Cao, H. Liang, and Y. Li, "Deepaco: Neural-enhanced ant systems for combinatorial optimization," *Advances in Neural Information Processing Systems*, vol. 36, 2024.
- [47] F. Liu, X. Tong, M. Yuan, X. Lin, F. Luo, Z. Wang, Z. Lu, and Q. Zhang, "An example of evolutionary computation+ large language model beating human: Design of efficient guided local search," *arXiv preprint arXiv:2401.02051*, 2024.
- [48] H. Ye, J. Wang, Z. Cao, F. Berto, C. Hua, H. Kim, J. Park, and G. Song, "Large language models as hyper-heuristics for combinatorial optimization," *arXiv preprint arXiv:2402.01145*, 2024, <https://github.com/ai4co/LLM-as-HH>.