

TAPE: Leveraging Agent Topology for Cooperative Multi-Agent Policy Gradient

Xingzhou Lou^{1,2*}, Junge Zhang^{1,2†}, Timothy J. Norman³, Kaiqi Huang^{1,2}, Yali Du⁴

¹School of Artificial Intelligence, University of Chinese Academy of Sciences

²Institute of Automation, Chinese Academy of Sciences

³University of Southampton

⁴King’s College London

louxingzhou2020@ia.ac.cn, jgzhang@nlpr.ia.ac.cn, t.j.norman@soton.ac.uk

kqhuang@nlpr.ia.ac.cn, yali.du@kcl.ac.uk

Abstract

Multi-Agent Policy Gradient (MAPG) has made significant progress in recent years. However, centralized critics in state-of-the-art MAPG methods still face the centralized-decentralized mismatch (CDM) issue, which means sub-optimal actions by some agents will affect other agent’s policy learning. While using individual critics for policy updates can avoid this issue, they severely limit cooperation among agents. To address this issue, we propose an agent topology framework, which decides whether other agents should be considered in policy gradient and achieves compromise between facilitating cooperation and alleviating the CDM issue. The agent topology allows agents to use coalition utility as learning objective instead of global utility by centralized critics or local utility by individual critics. To constitute the agent topology, various models are studied. We propose Topology-based multi-Agent Policy gradiEnt (TAPE) for both stochastic and deterministic MAPG methods. We prove the policy improvement theorem for stochastic TAPE and give a theoretical explanation for the improved cooperation among agents. Experiment results on several benchmarks show the agent topology is able to facilitate agent cooperation and alleviate CDM issue respectively to improve performance of TAPE. Finally, multiple ablation studies and a heuristic graph search algorithm are devised to show the efficacy of the agent topology.

1 Introduction

Recent years has witnessed dramatic progress of reinforcement learning (RL) and multi-agent reinforcement learning (MARL) in real life applications, such as unmanned vehicles (Liu et al. 2022), traffic signal control (Noaen et al. 2022) and on-demand delivery (Wang et al. 2023). Taking advantage of the centralized training decentralized execution (CTDE) (Oliehoek, Spaan, and Vlassis 2008; Kraemer and Banerjee 2016) paradigm, current cooperative MARL methods (Du et al. 2023; Wang et al. 2020a,b; Peng et al. 2021; Zhang et al. 2021; Zhou, Lan, and Aggarwal 2022) adopt value function factorization or a centralized critic to provide centralized learning signals to promote cooperation and achieve implicit or explicit credit assignment. Multi-agent

policy gradient (MAPG) (Lowe et al. 2017; Foerster et al. 2018; Zhou et al. 2020; Zhang et al. 2021; Zhou, Lan, and Aggarwal 2022; Du et al. 2019) applies RL policy gradient techniques (Sutton and Barto 2018; Silver et al. 2014; Lillcrap et al. 2015) to the multi-agent context. In CTDE, MAPG methods adopt centralized critics or value-mixing networks (Rashid et al. 2020b,a; Wang et al. 2020a) for individual critics so that agents can directly update their policies to maximize the global Q value Q_{tot}^π in their policy gradient. As a result, agents cooperate more effectively and obtain better expected team rewards.

The centralized critic approach has an inherent problem known as centralized-decentralized mismatch (CDM) (Wang et al. 2020c; Chen et al. 2022). The CDM issue refers to sub-optimal, or explorative actions of some agents negatively affecting policy learning of others, causing catastrophic miscoordination. The CDM issue arises because sub-optimal or explorative actions may lead to a small or negative centralized global Q value Q_{tot}^π , even if other agents take good or optimal actions. In turn, the small Q_{tot}^π will make the other agents mistake their good actions as bad ones and interrupt their policy learning. The Decomposed Off-Policy (DOP) approach (Wang et al. 2020c) deals with sub-optimal actions of other agents by linearly decomposed individual critics, which ignore the other agents’ actions in the policy gradient. But the use of individual critics severely limits agent cooperation.

We give an example to illustrate the issue of learning with centralised critics and individual critics respectively. Consider an one-step matrix game with two agents A, B where each agent has two actions a_0, a_1 . Reward $R(a_0, a_0) = 2, R(a_0, a_1) = -4, R(a_1, a_0) = -1$ and $R(a_1, a_1) = 0$. Assume agent A has a near-optimal policy with probability ϵ choosing non-optimal action a_1 and is using the COMA centralized critic (Foerster et al. 2018) for policy learning. If agent A takes optimal action a_0 and B takes the non-optimal action a_1 , agent A ’s counterfactual advantage $Adv_A(a_0, a_1) = Q_{tot}^\pi(a_0, a_1) - [(1 - \epsilon) Q_{tot}^\pi(a_0, a_1) + \epsilon Q_{tot}^\pi(a_1, a_1)] = -4\epsilon < 0$, which means agent A will mistakenly think a_0 as a bad action. Consequently, the sub-optimal action of agent B causes agent A to decrease the probability of taking optimal action a_0 and deviate from the optimal policy. Similar problems will occur with other centralized critics. If we employ individual critics, however, cooperation will be limited. Assume

*Work done while visiting King’s College London

†Correspondence

Copyright © 2024, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

both agents’ policies are initialized as random policies and learning with individual critics. For agent A , $Q_A(a_0) = \mathbb{E}_{a_B \sim \pi_B} [Q_{tot}^\pi(a_0, a_B)] = 0.5 \times 2 - 0.5 \times 4 = -1$. Similarly, we can get $Q_A(a_1) = -0.5$, $Q_B(a_0) = 0.5$, $Q_B(a_1) = -2$. The post-update joint-policy will be (a_1, a_0) and receive reward -1 , which is clearly sub-optimal.

In this paper, we aim to alleviate the CDM issue without hindering agent’s cooperation capacity by proposing an agent topology framework to describe the relationships between agents’ policy updates. Under the agent topology framework, agents connected in the topology consider and maximize each other’s utilities. Thus, the shared objective makes each individual agent form a coalition with its connected neighbors. Agents only consider the utilities of agents in the same coalition, facilitating in-coalition cooperation and avoiding influence of out-of-coalition agents. Based on the agent topology, we propose **Topology-based multi-Agent Policy gradient (TAPE)** for both stochastic and deterministic MAPG, where the agent topology can alleviate the bad influence of other agents’ sub-optimality without hindering cooperation among agents. Theoretically, we prove the policy improvement theorem for stochastic TAPE and give a theoretical explanation for improved cooperation by exploiting agent topology from the perspective of parameter-space exploration.

Empirically, we use three prevalent random graph models (Erdős, Rényi et al. 1960; Watts and Strogatz 1998; Albert and Barabási 2002) to constitute the agent topology. Results show that the Erdős–Rényi (ER) model (Erdős, Rényi et al. 1960) is able to generate the most diverse topologies. With diverse coalitions, agents are able to explore different cooperation patterns and achieve strong cooperation performance. Evaluation results on a matrix game, Level-based foraging (Papoudakis et al. 2021) and SMAC (Samvelyan et al. 2019) show that TAPE outperforms all baselines and the agent topology is able to improve base methods’ performance by both facilitating cooperation among agents and alleviating the CDM issue. Moreover, to show the efficacy of the agent topology, we conduct multiple studies and devise a heuristic graph search algorithm.

Contributions of this paper are three-fold: Firstly, We propose an agent topology framework and Topology-based multi-Agent Policy gradient (TAPE) to achieve compromise between facilitating cooperation and alleviating CDM issue; Secondly, we theoretically establish policy improvement theorem for stochastic TAPE and elaborate the cause for improved cooperation by agent topology; Finally, empirical results demonstrate that the agent topology is able to alleviate the CDM issue without hindering cooperation among agents, resulting in strong performance of TAPE.

2 Preliminaries

The cooperative multi-agent task in this paper is modelled as **Decentralized Partially Observable Markov Decision Process (Dec-POMDP)** (Oliehoek and Amato 2016). A Dec-POMDP is a tuple $G = \langle I, S, \mathcal{A}, P, r, \mathcal{O}, O, n, \gamma \rangle$, where $I = \{1, \dots, n\}$ is a finite set of n agents, S is the state space, \mathcal{A} is the agent action space and γ is a discount factor. At each timestep, every agent $i \in I$ picks an action $a_i \in \mathcal{A}$ to form the joint-action $\mathbf{a} \in \mathbf{A} = \mathcal{A}^n$ to interact with the

environment. Then a state transition will occur according to a state transition function $P(s'|s, \mathbf{a}) : S \times \mathbf{A} \times S \rightarrow [0, 1]$. All agents will receive a shared reward by the reward function $r(s, \mathbf{a}) : S \times \mathbf{A} \rightarrow \mathbb{R}$. During execution, every agent draws a local observation $o \in \mathcal{O}$ by an observation function $O(s, a) : S \times \mathcal{A} \rightarrow \mathcal{O}$. Every agent stores an observation-action history $\tau^a \in T = (\mathcal{O} \times \mathcal{A})$, based on which agent i derives a policy $\pi_i(a_i|\tau_i)$. The joint policy $\pi = \{\pi_1, \dots, \pi_n\}$ consists of policies of all agents. The global Q value function $Q_{tot}^\pi(s, \mathbf{a}) = \mathbb{E}_\pi[\sum_{i=0} \gamma^i r_{t+i}|s_t = s, \mathbf{a}_t = \mathbf{a}]$ is the expectation of discounted future reward summed over the joint-policy π .

The policy gradient in stochastic MAPG method DOP is: $g = \mathbb{E}_\pi \left[\sum_i k_i(s) \nabla_{\theta_i} \log \pi_i(a_i|\tau_i) Q_i^{\phi_i}(s, a_i) \right]$, where $k_i \geq 0$ is the positive coefficient provided by the mixing network, and the policy gradient in deterministic MAPG methods is $g = \mathbb{E}_D \left[\sum_i \nabla_{\theta_i} \pi_i(\tau_i) \nabla_{a_i} Q_{tot}^\pi(s, \mathbf{a})|_{a_i = \pi_i(\tau_i)} \right]$, where Q_{tot}^π is the centralized critic and π_i is the policy of agent i parameterized θ_i .

3 Related Work

Multi-Agent Policy Gradient The policy gradient in **stochastic MAPG** methods has the form $\mathbb{E}_\pi \left[\sum_i \nabla_{\theta_i} \log \pi_i(a_i|\tau_i) G_i \right]$ (Foerster et al. 2018; Wang et al. 2020c; Lou et al. 2023b; Chen et al. 2022), where objective G_i varies across different methods, such as counterfactual advantage (Foerster et al. 2018) and polarized joint-action value (Chen et al. 2022). The objective in DOP is individual aristocratic utility (Wolpert and Tumer 2001), which ignores other agents’ utilities to avoid the CDM issue, but the cooperation is also limited by this objective. It is worth noting that polarized joint-action value (Chen et al. 2022) also aims to address the CDM issue, but it only applies to stochastic MAPG methods, and the polarized global Q value can be very unstable. **Deterministic MAPG** methods use gradient ascent to directly maximize the centralized global Q value Q_{tot}^π . Lowe *et al.* (Lowe et al. 2017) model the global Q value with a centralized critic. Current deterministic MAPG methods (Zhang et al. 2021; Peng et al. 2021; Zhou, Lan, and Aggarwal 2022) adopt value factorization to mix individual Q values to get Q_{tot}^π . As the global Q value is determined by the centralized critic for all agents, sub-optimal actions of one agent will easily influence all others.

Topology in Reinforcement Learning Adjodah *et al.* (Adjodah et al. 2019) discuss the communication topology issue in parallel-running RL algorithms such as A3C (Mnih et al. 2016). Results show that the centralized learner implicitly yields a fully-connected communication topology among parallel workers, which will harm their performance. In MARL with decentralized training, communication topology is adopted to enable inter-agent communication among networked agents (Zhang et al. 2018; Wang et al. 2019; Konan, Seraj, and Gombolay 2022; Du et al. 2021). The communication topology allows agent to share local information with each other during both training and execution and even achieve local consensus, which further leads to better cooperation performance. In MARL with centralized training, deep

coordination graph (DCG) (Böhmer, Kurin, and Whiteson 2020) factorizes the joint value function according to a coordination graph to achieve a trade-off between representational capacity and generalization. Deep implicit coordination graph (Li et al. 2020) allows to infer the coordination graph dynamically by agent interactions instead of domain expertise in DCG. Ruan *et al.* (Ruan et al. 2022) learn an action coordination graph to represent agents’ decision dependency, which further coordinates the dependent behaviors among agents.

4 Topology-based Multi-Agent Policy Gradient

In this section, we propose Topology-based multi-Agent Policy gradient (TAPE), which exploits the agent topology for both stochastic and deterministic MAPG. This use of the agent topology provides a compromise between facilitating cooperation and alleviating CDM. The primary purpose of the agent topology is to indicate relationships between agents’ policy updates, so we focus on policy gradients of TAPE here and cover the remainder in supplementary material. First, we will define the agent topology.

The agent topology describes how agents should consider others’ utility during policy updates. Each agent is a vertex $v \in \mathcal{V}$ and \mathcal{E} is the set of edges. For a given topology, $(\mathcal{V}, \mathcal{E})$, if $e_{ij} \in \mathcal{E}$, the source agent i should consider the utility of the destination agent j in its policy gradient. The only constraint we place on a topology is that $\forall i, e_{ii} \in \mathcal{E}$, because agents should at least consider their own utility in the policy gradient. The topology captures the relationships between agents’ policy updates, not their communication network at test time (Foerster et al. 2016; Das et al. 2019; Wang et al. 2019; Ding, Huang, and Lu 2020). Connected agents consider and maximize each other’s utilities together. Thus, the shared objective makes each individual agent form a coalition with the connected neighbors. We use the adjacency matrix E to refer the agent topology in what follows.

In our agent topology framework, DOP (Wang et al. 2020c) (policy gradient given in section 2) and other independent learning algorithms’ has an edgeless agent topology. The adjacency matrix is the identity matrix and no edge exists in the topology. With no coalition, DOP agent will only maximize its own individual utility Q_i , and hence is poor at cooperation. Although DOP adopts a mixing network for the individual utilities to enhance cooperation, an agent’s ability to cooperate is still limited, which we will empirically show in the matrix game experiments. Methods with centralized critic such as COMA (Foerster et al. 2018), FACMAC (Peng et al. 2021) and PAC (Zhou, Lan, and Aggarwal 2022) yields the fully-connected agent topology. In these methods, there is only one coalition with all of the agents in it (all edges exist in the topology), and all agents update their policies based on the centralized critic. Consequently, they suffer from the CDM issue severely, because the influence of an agent’s sub-optimal behavior will spread to the entire multi-agent system.

4.1 Stochastic TAPE

Instead of global centralized critic (Foerster et al. 2018), we use the agent topology to aggregate individual utilities and

critics to facilitate cooperation among agents for stochastic MAPG (Wang et al. 2020c). To this end, a new learning objective *Coalition Utility* for the policy gradient is defined as below.

Definition 1 (Coalition Utility). *Coalition Utility U_i for agent i is the summation of individual utility U_j of connected agent j in agent topology E , i.e. $U_i = \sum_{j=1}^n E_{ij} U_j$, where $U_j(s, a_j) = Q_{tot}^\phi(s, \mathbf{a}) - \sum_{a'_j} \pi_j(a'_j | \tau_j) Q_{tot}^\phi(s, (a'_j, \mathbf{a}_{-j}))$.*

U_j is the aristocrat utility from (Wang et al. 2020c; Wolpert and Tumer 2001). $E_{ij} = 1$ only if agent j is connected to agent i in E and Q_{tot}^ϕ is the global Q value function. Coalition utility only depends on in-coalition agents because if agent j is not in agent i ’s coalition, $E_{ij} = 0$. With the coalition utility, we propose **stochastic TAPE** with the policy gradient given by

$$\begin{aligned} \nabla J_1(\theta) &= \mathbb{E}_\pi \left[\sum_i \nabla_{\theta_i} \log \pi_i(a_i | \tau_i) U_i \right] \quad (1) \\ &= \mathbb{E}_\pi \left[\sum_{i,j} E_{ij} k_j(s) \nabla_{\theta_i} \log \pi_i(a_i | \tau_i) Q_j^{\phi_j}(s, a_j) \right], \quad (2) \end{aligned}$$

where $k_j \geq 0$ is the weight for agent j ’s local Q value $Q_j^{\phi_j}$ provided by the mixing network. The policy gradient derivation from Eq. 1 to Eq. 2 is provided in the appendix A. Since the local utility of other in-coalition agents is maximized by the policy updates, cooperation among agents is facilitated. Pseudo-code and more details of stochastic TAPE are provided in the appendix E.1.

4.2 Deterministic TAPE

Current deterministic MAPG methods (Peng et al. 2021; Zhang et al. 2021; Zhou, Lan, and Aggarwal 2022) yield fully-connected agent topology, which makes agents vulnerable to bad influence of other agents’ sub-optimal actions. A mixing network f_{mix} is adopted to mix local Q value functions $Q_i^{\pi_i}$. Each agent uses deterministic policy gradient to update parameters and directly maximize global Q value $Q_{tot}^\pi = f_{mix}(s, Q_1^{\pi_1}, \dots, Q_n^{\pi_n})$. We use the agent topology to drop out utilities of out-of-coalition agents, so that influence of their sub-optimal actions will not spread to in-coalition agents. To this end, *Coalition Q* is defined as below.

Definition 2 (Coalition Q). *Coalition Q Q_{co}^i for agent i is the mixture of its in-coalition agents’ local Q values with mixing network f_{mix} , i.e.*

$$Q_{co}^i(s, \mathbf{a}) = f_{mix}(s, \mathbb{1}[E_{i1}]Q_1^{\pi_1}, \dots, \mathbb{1}[E_{i,n}]Q_n^{\pi_n}), \quad (3)$$

where $\mathbb{1}[E_{ij}]$ is the indicator function and $\mathbb{1}[E_{ij}] = 1$ only when edge E_{ij} exists in the topology.

During policy update, out-of-coalition agents’ Q values are always masked out, so agent i ’s policy learning will not be affected by out-of-coalition agents. Based on Coalition Q , we propose **deterministic TAPE**, whose policy gradient is




Figure 1: (a) gives the proposed three matrix games of different levels. We use different colors for different levels of game. Blue represents Easy, green represents Medium and red represents Hard. (b), (c) and (d) give evaluation results. Stochastic TAPE has the best performance because the agents directly maximize joint utility to achieve strong cooperation. The only difference between TAPE and DOP is that TAPE adopts the agent topology. Although COMA is seen as a weak baseline on SMAC, it achieves much better performance than DOP. QMIX fails to perform well in these games as they are not monotonic games.

given by

$$\nabla J_2(\theta) = \mathbb{E}_{\mathcal{D}} \left[\sum_i \nabla_{\theta_i} \pi_i(\tau_i) \nabla_{a_i} \hat{Q}_{co}^i(s, \mathbf{a}) \Big|_{a_i = \pi_i(\tau_i)} \right] \quad (4)$$

where $\hat{Q}_{co}^i(s, \mathbf{a}) = f_{\text{mix}} \left(s, \mathbb{1}[E_{i1}] \hat{Q}_1^{\phi_1}, \dots, \mathbb{1}[E_{i,n}] \hat{Q}_n^{\phi_n} \right)$ and $\hat{Q}_i^{\phi_i}(\tau_i, a_i, m_i) = Q_i^{\phi_i}(\tau_i, a_i, m_i) - \alpha \log \pi_i(a_i | \tau_i)$ is the local soft Q value (Zhang et al. 2021) augmented with assistive information m_i which contains information to assist policy learning towards the optimal policy as in (Zhou, Lan, and Aggarwal 2022). After dropping out agents not in the coalition, the bad influence of out-of-coalition sub-optimal actions will not affect in-coalition agents. More details and pseudo-code are provided in the appendix E.2.

5 Analysis

5.1 Agent Topology

Although the agent topology can be any arbitrary topology, a proper agent topology should be able to explore diverse cooperation pattern, which is essential for robust cooperation (Li et al. 2021; Strouse et al. 2021; Lou et al. 2023a; Yan et al. 2023). We studied three prevalent random graph model: Barabási–Albert (BA) model (Albert and Barabási 2002), Watts–Strogatz (WS) model (Watts and Strogatz 1998) and Erdős–Rényi (ER) model (Erdős, Rényi et al. 1960). BA model is a scale-free network commonly used for citation and signaling biological networks (Barabási and Albert 1999). WS model is known as the small-world network where each nodes can be reached through a small number of nodes, resulting in the six degrees of separation (Travers and Milgram 1977). While in ER model, each edge between any two nodes has an independent probability of being present. Formally, the adjacency matrix E of ER agent topology $(\mathcal{V}, \mathcal{E})$ for n agents is defined as $\forall i \in \{1, \dots, n\}, E_{ii} = 1; \forall i, j \in \{1, \dots, n\}, i \neq j, E_{ij} = 1$ with probability p otherwise 0.

In research question 1 of section 6.3, we found that ER model is able to generate the most diverse topologies, which in turn help the agents explore diverse cooperation pattern and achieve strongest performance. Thus, we use ER model to constitute the agent topology in the experiments.

5.2 Theoretical Results

We now establish policy improvement theorem of stochastic TAPE, and prove a theorem for the cooperation improvement from the perspective of exploring the parameter space, which is a common motivation in RL research (Schulman, Chen, and Abbeel 2017; Haarnoja et al. 2018; Zhang et al. 2021; Adjodah et al. 2019). We assume the policy to have tabular expressions.

The following theorem states that stochastic TAPE updates can monotonically improve the objective function $J(\pi) = \mathbb{E}_{\pi} [\sum_t \gamma^t r_t]$.

Theorem 1. [stochastic TAPE policy improvement theorem] *With tabular expressions for policies, for any pre-update policy π and updated policy $\hat{\pi}$ by policy gradient in Eq. 2 that satisfy for any agent i , $\hat{\pi}_i(a_i | \tau_i) = \pi_i(a_i | \tau_i) + \beta_{a_i, s} \delta$, where δ is a sufficiently small number, we have $J(\hat{\pi}) \geq J(\pi)$, i.e. the joint policy is improved by the update.*

Please refer to Appendix B for the proof of Theorem 1. Although this policy improvement theorem is established for policies with tabular expressions, we provide conditions in the proof, under which policy improvement is guaranteed even with function approximators.

Next, we provide a theoretical insight that compared to using individual critics, stochastic TAPE can better explore the parameter space to find more effective cooperation pattern. One heuristic for measuring such capacity is the diversity of parameter updates during each iteration (Adjodah et al. 2019), which is measured by the variance of parameter updates.

Given state s and action a_i , let $\xi_{a_i, s}^{\text{TAPE}}$ and $\xi_{a_i, s}^{\text{DOP}}$ denote the stochastic TAPE and DOP parameter updates respectively. The following theorem states that stochastic TAPE policy update is more diverse so that it can explore the parameter space more effectively.

Theorem 2. *For any agent i and $\forall s, a_i$, the stochastic TAPE policy update $\xi_{a_i, s}^{\text{TAPE}}$ and DOP policy update $\xi_{a_i, s}^{\text{DOP}}$ satisfy that $\text{Var} [\xi_{a_i, s}^{\text{TAPE}}] \geq \text{Var} [\xi_{a_i, s}^{\text{DOP}}]$, and $\Delta = \text{Var} [\xi_{a_i, s}^{\text{TAPE}}] - \text{Var} [\xi_{a_i, s}^{\text{DOP}}]$ is in proportion to p^2 , where p is the probability of edges being present in the Erdős–Rényi model, i.e. $\Delta \propto p^2$.*

Theorem 2 shows that compared to solely using individual critics, our agent topology provides larger diversity in policy updates to find better cooperation pattern. More details and

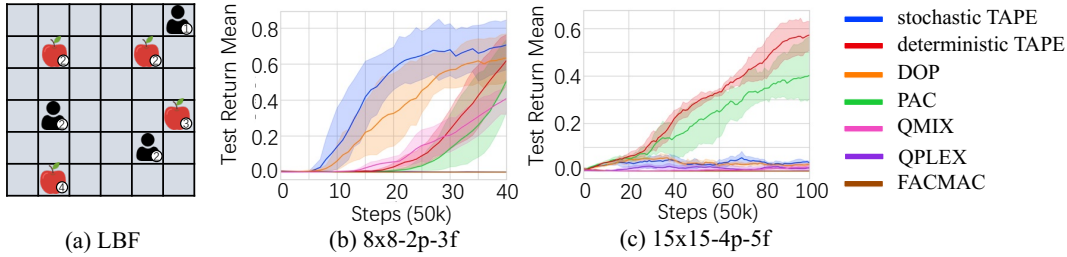


Figure 2: (a) gives a scenario 6x6-3p-4f in LBF. 6x6-3p-4f stands for 6x6 grid-world with 3 players and 4 fruits. (b) In 8x8-2p-3f, stochastic TAPE achieve best performance. While in the more difficult task 15x15-4p-5f (c), deterministic TAPE outperform its base method and all other baselines. See stochastic TAPE against DOP, and deterministic TAPE against PAC for comparison.

proof are provided in the appendix C. It is worth noting that although a large hyperparameter p in the agent topology means larger diversity in parameter updates, the CDM issue will also become severer because the connections among agents become denser. Thus, p must be set properly to achieve compromise between facilitating cooperation and avoiding CDM issue, which we will show later in the experiments.

6 Experiment

In this section, we first demonstrate that by ignoring other agents in the policy gradient to avoid bad influence of their sub-optimal actions, cooperation among agents is severely harmed. To this end, three one-step matrix games that require strong cooperation are proposed. Then, we evaluate the efficacy of the proposed methods on (a) Level-Based Foraging (LBF) (Papoudakis et al. 2021); (b) Starcraft II Multi-Agent Challenge (SMAC) (Samvelyan et al. 2019), and answer several research questions via various ablations and a heuristic graph search technique. Our code is available here¹.

6.1 Matrix Game

We propose 3 one-step matrix games, which are harder versions of the example in introduction. The matrix games are given in Fig. 1(a). We use different colors to show rewards in different games (blue for Easy, green for Medium and red for Hard). The optimal joint policy is for both agents to take action a_0 . But agent A_0 lacks motivation to choose a_0 because it is very likely to receive a large penalty (-8 or -16). Thus, this game requires strong cooperation among agents. In the Medium game, we further increase the penalty for agent 0 to choose a_0 . In the Hard game, we keep the large penalty and add a local optimal reward at (a_1, a_1) . Note that these matrix games are not monotonic games (Rashid et al. 2020b) as the optimal action for each agent depends on other agents. The evaluation results are given in Fig. 1.

With the agent topology to encourage cooperation, stochastic TAPE outperforms other methods by a large margin and is able to learn optimal joint policy even in the Hard game. DOP agents optimize individual utilities, ignoring utilities of other agents to avoid the influence of their sub-optimal actions, which result in severe miscoordination in these games. But since DOP agents adopt stochastic policy, they may receive

some large reward after enough exploration. But the learning efficiency is much lower than stochastic TAPE. COMA is a weak baseline on complex tasks (Samvelyan et al. 2019) (0% win rate in all maps in section 6.3). But since COMA agents optimize global Q value (expected team reward sum) instead of individual utility in DOP, it can achieve better results on these tasks requiring strong cooperation. These matrix games demonstrate the importance of considering the utility of other agents in cooperative tasks. With the agent topology, stochastic TAPE can facilitate cooperation among agents and alleviate CDM issue simultaneously.

6.2 Level-Based Foraging

In Level-Based Foraging (LBF (Papoudakis et al. 2021)), agents navigate a grid-world and collect randomly-scattered food items. Agents and food items are assigned with levels. A food item is only allowed to be collected when near-by agents' level sum is larger than the food level. Reward is only given when a food item is collected, assigning the environment with sparse-reward property. Test return is 1 when all food items are collected. Compared baselines include both value-based methods: QMIX (Rashid et al. 2020b) and QPLEX (Wang et al. 2020a), and policy-based methods: DOP (Wang et al. 2020c), FACMAC (Peng et al. 2021) and PAC (Zhou, Lan, and Aggarwal 2022). Scenario illustration and results are given in Fig. 2.

To make 8x8-2p-3f more difficult, food items can only be collected when all agents participate. In this simple and sparse-reward task, with the stochastic policy and enhanced cooperation, stochastic TAPE outperforms all other methods on convergence speed and performance. While in 15x15-4p-5f, only state-of-the-art method PAC and deterministic TAPE learn to collect food items. With the agent topology to keep out bad influence of other agents' sub-optimal actions, deterministic TAPE achieves best performance.

6.3 StarCraft Multi-Agent Challenge

StarCraft Multi-Agent Challenge (SMAC) (Samvelyan et al. 2019) is a challenging benchmark built on StarCraft II, where agents must cooperate with each other to defeat enemy teams controlled by built-in AI. We evaluate the proposed methods and baselines with the recommended evaluation protocol and metric in six maps including three hard maps (3s_vs_4z,

¹github.com/LxzGordon/TAPE

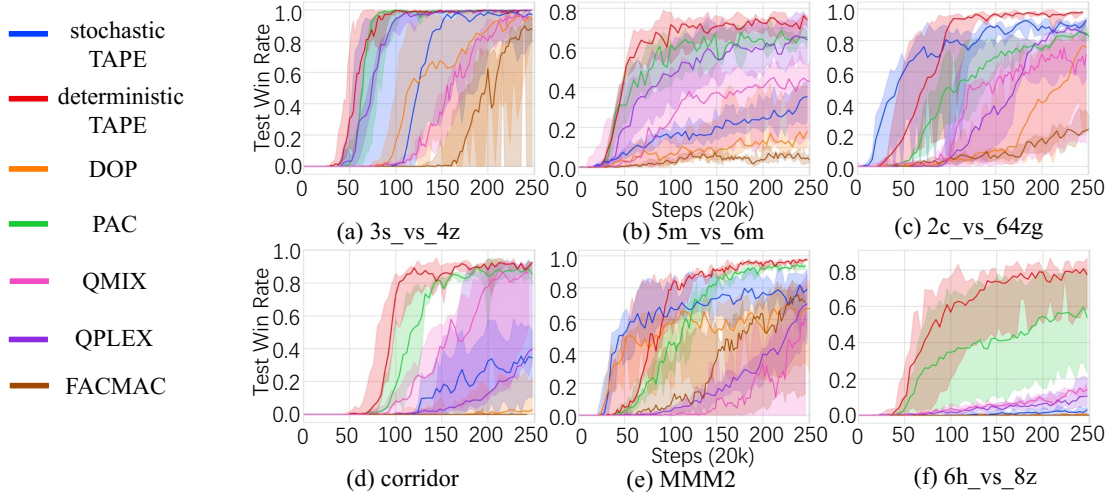


Figure 3: Experiment results on SMAC. (a-c) give the results in hard maps, and (d-f) are results in super-hard maps. After adopting our agent topology to facilitate cooperation and alleviate CDM issue, stochastic TAPE and deterministic TAPE outperforms their base methods respectively. See stochastic TAPE against DOP, and deterministic TAPE against PAC for comparison.

5m_vs_6m and 2c_vs_64zg) and three super hard maps (corridor, MMM2 and 6h_vs_8z). All algorithms are run for four times with different random seeds. Each run lasts for 5×10^6 environmental steps. During training, each algorithm has four parallel environment to collect training data.

Overall results The overall results in six maps are provided in Fig. 3. We can see deterministic TAPE outperforms all other methods in terms of performance and convergence speed. In 6h_vs_8z, one of the most difficult maps in SMAC, deterministic TAPE achieves noticeably better performance than its base method PAC and other baselines. It’s worth noting that after integrating agent topology, both stochastic TAPE and deterministic TAPE have better performance compared to the base methods. This demonstrates the efficacy of the proposed agent topology in facilitating cooperation for DOP and alleviating CDM issue for PAC. Especially, in 2c_vs_64zg, stochastic TAPE outperforms all of the baselines except for our deterministic TAPE while its base method DOP struggles to perform well.

Next, we answer three research questions by ablations and additional experiments. The research questions are: **Q1.** What is the proper model to constitute the agent topology? **Q2.** Is there indeed a compromise between facilitating cooperation and suffering from the CDM issue? **Q3.** Is the agent topology capable of compromising between facilitating cooperation and the CDM issue to achieve best performance?

Q1. We study three prevalent random graph models: Barabási–Albert (BA) model (Albert and Barabási 2002), Watts–Strogatz (WS) model (Watts and Strogatz 1998) and the Erdős–Rényi (ER) model (Erdős, Rényi et al. 1960) via visualization and ablation study. First, we generate 1000 topologies for 12 agents with each model and give the visualization result in Fig. 4(a), where x -axis is average degree and y -axis is connectivity (minimum number of edges required, by removing which the graph becomes two sub-graphs). Average degree and connectivity are two essential factors for

agent topology as they reflect the level of CDM issue and cooperation. Compared to the other two models, ER model generates much more diverse topologies, covering the area from edgeless topology to fully-connected topology. Then, we evaluate stochastic TAPE with each model on MMM2, a super hard map in SMAC. Results are given in Fig. 4(b). For the random graph models, the larger the graph diversity in Fig. 4(a), the stronger the performance is. Thus, we constitute the agent topology with ER model in other experiments. For fully-connected topology, the performance demonstrates very large variance, because once a sub-optimal action occurs, its bad influence will easily spread through the centralized critic to all other agents. It is worth nothing that the graphs can also be generated via Bayesian optimization, but this may also result in limited graph diversity, causing unstable or even worse performance. Thus, how to generate agent topology via optimization-based methods remains a challenge.

Q2. The compromise here means the more connection among agents to improve performance, the severer CDM issue becomes, and when it is too severe, it will in turn affect performance. To answer this research question, we devise a heuristic graph search technique. During policy training of agent i , we generate n topologies with the ER model in each step and use them to update the agent policy. After obtaining n updated policy $[\pi_1^i, \dots, \pi_n^i]$, we evaluate the post-update global Q value $Q_{tot}^{\pi^{-i}, \pi_j^i}$ and choose the policy with largest global Q value as the updated policy, i.e. $\pi^i = \arg \max_j Q_{tot}^{\pi^{-i}, \pi_j^i}$. The motivation of this heuristic graph search technique is that global Q value is the expected future reward sum, which shows the post-update performance. Using this technique, we can find the topology with better performance. Then, we respectively use the graph search technique when p is small or large and give the visualization of preferred topologies in Fig. 5. The results confirm that the compromise does exist, because (1) facilitating cooperation

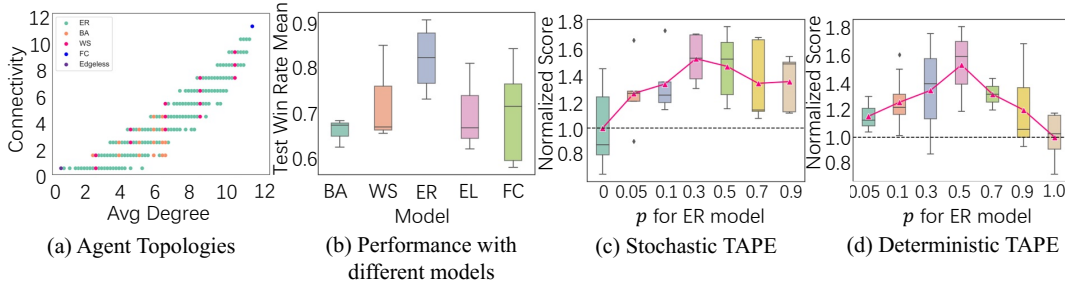


Figure 4: (a) and (b) show the results and performance of using different models to constitute agent topologies. BA is Barabási–Albert model, WS is Watts–Strogatz model, ER is Erdős–Rényi model, EL (Edgeless) and FC (Fully-Connected) are the topologies adopted in DOP and PAC respectively. ER has the most diverse topoloies and strongest performance. (c) and (d) show the performance of stochastic TAPE and deterministic TAPE in MMM2 with difference hyperparameter p for ER model. Evaluation metric is test win rate and scores are normalized by the base method. In base method DOP, $p = 0$ and base method PAC $p = 1$. The boxplot is obtained with four different random seeds, and the red lines show the mean performance.




Figure 5: The heatmaps show the difference between the frequency of edges being present and the probability p . Source and Destination represent starting node and destination node of an edge. During training, over 1 million agent topology is generated. According to the law of large numbers, the difference is always around 0 when the heuristic graph search technique is not used in (b). In (a) and (c), we adopt the heuristic graph search technique to choose the agent topology with strongest performance. When p is too small (0.01 in (a)), the connection among agents is too sparse, weakening cooperation among agents. Therefore, agent topologies with more edges can facilitate cooperation and are preferred by the graph search technique. As a results, the difference is always positive in (a). On the contrary, when the connection is too dense ($p = 0.3$ in (c)), topologies with less edges are preferred because they stop bad influence of sub-optimal actions from spreading and have better performance, resulting in negative differences in (c).

by building more agent connections when there is little CDM issue (Fig. 5(a)), and (2) removing connections to stop bad influence of sub-optimal actions from spreading when CDM issue is severe (Fig. 5(c)), can both improve performance.

Q3. We answer this research question by giving the performance with different hyperparameter p , as it controls the level of CDM issue and cooperation. The results are given in Fig. 4(c), (d). Large p stands for dense connections, where agents are easily affected by sub-optimal actions of other agents but cooperation is strongly encouraged. Small p means sparse connections, where sub-optimal actions’ influence will not easily spread but cooperation among agents is limited. (c) and (d) are drawn at the end of training and half of training to show the convergence performance and speed respectively. We can see the performances first increase when p is small and later decrease when p is too large. The best performance appears at the point where the cooperation is strong and CDM issue is acceptable. From the results, we can say our ER agent topology is able to compromise between cooperation and alleviating the CDM issue to achieve the best performance.

7 Conclusion and Future Work

In this paper, we propose an agent topology framework, which aims to alleviate the CDM issue without limiting agents’ cooperation capacity. Based on the agent topology, we propose TAPE for both stochastic and deterministic MAPG methods. Theoretically, we prove the policy improvement theorem for stochastic TAPE and give a theoretical explanation about the improved cooperation among agents. Empirically, we evaluate the proposed methods on several benchmarks. Experiment results show that the methods outperform their base methods and other baselines in terms of convergence speed and performance. A heuristic graph search algorithm is devised and various studies are conducted, which validate the efficacy of our proposed agent topology.

Limitation and Future Work In this work, we consider constructing agent topology with existing random graph models without learning-based methods. Our future work is to adaptively learn the agent topology that can simultaneously facilitate agent cooperation and alleviate the CDM issue.

Acknowledgments

This work is supported by Basic Cultivation Fund project, CAS (JCPYJJ-22017), Strategic Priority Research Program of Chinese Academy of Sciences (XDA27010300).

References

- Adjodah, D.; Calacci, D.; Dubey, A.; Goyal, A.; Krafft, P.; Moro, E.; and Pentland, A. 2019. Leveraging Communication Topologies Between Learning Agents in Deep Reinforcement Learning. *arXiv preprint arXiv:1902.06740*.
- Albert, R.; and Barabási, A.-L. 2002. Statistical mechanics of complex networks. *Reviews of modern physics*, 74(1): 47.
- Barabási, A.-L.; and Albert, R. 1999. Emergence of scaling in random networks. *science*, 286(5439): 509–512.
- Böhmer, W.; Kurin, V.; and Whiteson, S. 2020. Deep coordination graphs. In *International Conference on Machine Learning*, 980–991. PMLR.
- Chen, W.; Li, W.; Liu, X.; and Yang, S. 2022. Learning Credit Assignment for Cooperative Reinforcement Learning. *arXiv preprint arXiv:2210.05367*.
- Das, A.; Gervet, T.; Romoff, J.; Batra, D.; Parikh, D.; Rabbat, M.; and Pineau, J. 2019. Tarmac: Targeted multi-agent communication. In *International Conference on Machine Learning*, 1538–1546. PMLR.
- Ding, Z.; Huang, T.; and Lu, Z. 2020. Learning individually inferred communication for multi-agent cooperation. *Advances in Neural Information Processing Systems*, 33: 22069–22079.
- Du, Y.; Han, L.; Fang, M.; Dai, T.; Liu, J.; and Tao, D. 2019. LIIR: learning individual intrinsic reward in multi-agent reinforcement learning. In *Proceedings of the 33rd International Conference on Neural Information Processing Systems (NeurIPS)*, 4403–4414.
- Du, Y.; Leibo, J. Z.; Islam, U.; Willis, R.; and Sunehag, P. 2023. A Review of Cooperation in Multi-agent Learning. *arXiv preprint arXiv:2312.05162*.
- Du, Y.; Liu, B.; Moens, V.; Liu, Z.; Ren, Z.; Wang, J.; Chen, X.; and Zhang, H. 2021. Learning Correlated Communication Topology in Multi-Agent Reinforcement learning. In *International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, 456–464.
- Erdős, P.; Rényi, A.; et al. 1960. On the evolution of random graphs. *Publ. Math. Inst. Hung. Acad. Sci.*, 5(1): 17–60.
- Foerster, J.; Assael, I. A.; De Freitas, N.; and Whiteson, S. 2016. Learning to communicate with deep multi-agent reinforcement learning. *Advances in neural information processing systems*, 29.
- Foerster, J.; Farquhar, G.; Afouras, T.; Nardelli, N.; and Whiteson, S. 2018. Counterfactual multi-agent policy gradients. In *Proceedings of the AAAI conference on artificial intelligence*, volume 32.
- Haarnoja, T.; Zhou, A.; Hartikainen, K.; Tucker, G.; Ha, S.; Tan, J.; Kumar, V.; Zhu, H.; Gupta, A.; Abbeel, P.; et al. 2018. Soft actor-critic algorithms and applications. *arXiv preprint arXiv:1812.05905*.
- Konan, S.; Seraj, E.; and Gombolay, M. 2022. Iterated reasoning with mutual information in cooperative and byzantine decentralized teaming. *arXiv preprint arXiv:2201.08484*.
- Kraemer, L.; and Banerjee, B. 2016. Multi-agent reinforcement learning as a rehearsal for decentralized planning. *Neurocomputing*, 190: 82–94.
- Li, C.; Wang, T.; Wu, C.; Zhao, Q.; Yang, J.; and Zhang, C. 2021. Celebrating diversity in shared multi-agent reinforcement learning. *Advances in Neural Information Processing Systems*, 34: 3991–4002.
- Li, S.; Gupta, J. K.; Morales, P.; Allen, R.; and Kochenderfer, M. J. 2020. Deep implicit coordination graphs for multi-agent reinforcement learning. *arXiv preprint arXiv:2006.11438*.
- Lillicrap, T. P.; Hunt, J. J.; Pritzel, A.; Heess, N.; Erez, T.; Tassa, Y.; Silver, D.; and Wierstra, D. 2015. Continuous control with deep reinforcement learning. *arXiv preprint arXiv:1509.02971*.
- Liu, H.; Kiumarsi, B.; Kartal, Y.; Koru, A. T.; Modares, H.; and Lewis, F. L. 2022. Reinforcement learning applications in unmanned vehicle control: A comprehensive overview. *Unmanned Systems*, 1–10.
- Lou, X.; Guo, J.; Zhang, J.; Wang, J.; Huang, K.; and Du, Y. 2023a. PECAN: Leveraging Policy Ensemble for Context-Aware Zero-Shot Human-AI Coordination. *arXiv preprint arXiv:2301.06387*.
- Lou, X.; Zhang, J.; Du, Y.; Yu, C.; He, Z.; and Huang, K. 2023b. Leveraging Joint-action Embedding in Multi-agent Reinforcement Learning for Cooperative Games. *IEEE Transactions on Games*.
- Lowe, R.; Wu, Y. I.; Tamar, A.; Harb, J.; Pieter Abbeel, O.; and Mordatch, I. 2017. Multi-agent actor-critic for mixed cooperative-competitive environments. *Advances in neural information processing systems*, 30.
- Mnih, V.; Badia, A. P.; Mirza, M.; Graves, A.; Lillicrap, T.; Harley, T.; Silver, D.; and Kavukcuoglu, K. 2016. Asynchronous methods for deep reinforcement learning. In *International conference on machine learning*, 1928–1937. PMLR.
- Noaen, M.; Naik, A.; Goodman, L.; Crebo, J.; Abrar, T.; Abad, Z. S. H.; Bazzan, A. L.; and Far, B. 2022. Reinforcement learning in urban network traffic signal control: A systematic literature review. *Expert Systems with Applications*, 116830.
- Oliehoek, F. A.; and Amato, C. 2016. *A concise introduction to decentralized POMDPs*. Springer.
- Oliehoek, F. A.; Spaan, M. T.; and Vlassis, N. 2008. Optimal and approximate Q-value functions for decentralized POMDPs. *Journal of Artificial Intelligence Research*, 32: 289–353.
- Papoudakis, G.; Christianos, F.; Schäfer, L.; and Albrecht, S. V. 2021. Benchmarking Multi-Agent Deep Reinforcement Learning Algorithms in Cooperative Tasks. In *Proceedings of the Neural Information Processing Systems Track on Datasets and Benchmarks (NeurIPS)*.
- Peng, B.; Rashid, T.; Schroeder de Witt, C.; Kamienny, P.-A.; Torr, P.; Böhmer, W.; and Whiteson, S. 2021. Facmac:

- Factored multi-agent centralised policy gradients. *Advances in Neural Information Processing Systems*, 34: 12208–12221.
- Rashid, T.; Farquhar, G.; Peng, B.; and Whiteson, S. 2020a. Weighted qmix: Expanding monotonic value function factorisation for deep multi-agent reinforcement learning. *Advances in neural information processing systems*, 33: 10199–10210.
- Rashid, T.; Samvelyan, M.; De Witt, C. S.; Farquhar, G.; Foerster, J.; and Whiteson, S. 2020b. Monotonic value function factorisation for deep multi-agent reinforcement learning. *The Journal of Machine Learning Research*, 21(1): 7234–7284.
- Ruan, J.; Du, Y.; Xiong, X.; Xing, D.; Li, X.; Meng, L.; Zhang, H.; Wang, J.; and Xu, B. 2022. GCS: graph-based coordination strategy for multi-agent reinforcement learning. *arXiv preprint arXiv:2201.06257*.
- Samvelyan, M.; Rashid, T.; de Witt, C. S.; Farquhar, G.; Nardelli, N.; Rudner, T. G. J.; Hung, C.-M.; Torr, P. H. S.; Foerster, J.; and Whiteson, S. 2019. The StarCraft Multi-Agent Challenge. *CoRR*, abs/1902.04043.
- Schulman, J.; Chen, X.; and Abbeel, P. 2017. Equivalence between policy gradients and soft q-learning. *arXiv preprint arXiv:1704.06440*.
- Silver, D.; Lever, G.; Heess, N.; Degris, T.; Wierstra, D.; and Riedmiller, M. 2014. Deterministic policy gradient algorithms. In *International conference on machine learning*, 387–395. Pmlr.
- Strouse, D.; McKee, K.; Botvinick, M.; Hughes, E.; and Everett, R. 2021. Collaborating with humans without human data. *Advances in Neural Information Processing Systems*, 34: 14502–14515.
- Sutton, R. S.; and Barto, A. G. 2018. *Reinforcement learning: An introduction*. MIT press.
- Travers, J.; and Milgram, S. 1977. An experimental study of the small world problem. In *Social networks*, 179–197. Elsevier.
- Wang, J.; Ren, Z.; Liu, T.; Yu, Y.; and Zhang, C. 2020a. Qplex: Duplex dueling multi-agent q-learning. *arXiv preprint arXiv:2008.01062*.
- Wang, S.; Hu, S.; Guo, B.; and Wang, G. 2023. Cross-Region Courier Displacement for On-Demand Delivery With Multi-Agent Reinforcement Learning. *IEEE Transactions on Big Data*.
- Wang, T.; Gupta, T.; Mahajan, A.; Peng, B.; Whiteson, S.; and Zhang, C. 2020b. Rode: Learning roles to decompose multi-agent tasks. *arXiv preprint arXiv:2010.01523*.
- Wang, T.; Wang, J.; Zheng, C.; and Zhang, C. 2019. Learning nearly decomposable value functions via communication minimization. *arXiv preprint arXiv:1910.05366*.
- Wang, Y.; Han, B.; Wang, T.; Dong, H.; and Zhang, C. 2020c. Off-policy multi-agent decomposed policy gradients. *arXiv preprint arXiv:2007.12322*.
- Watts, D. J.; and Strogatz, S. H. 1998. Collective dynamics of ‘small-world’ networks. *nature*, 393(6684): 440–442.
- Wolpert, D. H.; and Tumer, K. 2001. Optimal payoff functions for members of collectives. *Advances in Complex Systems*, 4(02n03): 265–279.
- Yan, X.; Guo, J.; Lou, X.; Wang, J.; Zhang, H.; and Du, Y. 2023. An Efficient End-to-End Training Approach for Zero-Shot Human-AI Coordination. In *NeurIPS 2023*.
- Zhang, K.; Yang, Z.; Liu, H.; Zhang, T.; and Basar, T. 2018. Fully decentralized multi-agent reinforcement learning with networked agents. In *International Conference on Machine Learning*, 5872–5881. PMLR.
- Zhang, T.; Li, Y.; Wang, C.; Xie, G.; and Lu, Z. 2021. Fop: Factorizing optimal joint policy of maximum-entropy multi-agent reinforcement learning. In *International Conference on Machine Learning*, 12491–12500. PMLR.
- Zhou, H.; Lan, T.; and Aggarwal, V. 2022. PAC: Assisted Value Factorisation with Counterfactual Predictions in Multi-Agent Reinforcement Learning. *arXiv preprint arXiv:2206.11420*.
- Zhou, M.; Liu, Z.; Sui, P.; Li, Y.; and Chung, Y. Y. 2020. Learning implicit credit assignment for cooperative multi-agent reinforcement learning. *Advances in neural information processing systems*, 33: 11853–11864.