

F2A2: Flexible Fully-decentralized Approximate Actor-critic for Cooperative Multi-agent Reinforcement Learning

Wenhao Li

LIWENHAO@CUHK.EDU.CN

*School of Data Science
The Chinese University of Hong Kong, Shenzhen
Shenzhen Institute of Artificial Intelligence and Robotics for Society
Shenzhen 518172, China*

Bo Jin

BJIN@TONGJI.EDU.CN

*Corresponding author
School of Software Engineering
Shanghai Research Institute for Intelligent Autonomous Systems
Tongji University
Shanghai 201804, China*

Xiangfeng Wang

XFWANG@CS.ECNU.EDU.CN

*Corresponding author
School of Computer Science and Technology
Key Laboratory of Mathematics and Engineering Applications, Ministry of Education
East China Normal University
Shanghai 200062, China*

Junchi Yan

YANJUNCHI@SJTU.EDU.CN

*Department of Computer Science and Engineering
Key Laboratory of Artificial Intelligence, Ministry of Education
Shanghai Jiao Tong University
Shanghai 200240, China*

Hongyuan Zha

ZHAHY@CUHK.EDU.CN

*School of Data Science
The Chinese University of Hong Kong, Shenzhen
Shenzhen Institute of Artificial Intelligence and Robotics for Society
Shenzhen 518172, China*

Editor: Ambuj Tewari

Abstract

Traditional centralized multi-agent reinforcement learning (MARL) algorithms are sometimes unpractical in complicated applications due to non-interactivity between agents, the curse of dimensionality, and computation complexity. Hence, several decentralized MARL algorithms are motivated. However, existing decentralized methods only handle the fully cooperative setting where massive information needs to be transmitted in training. The block coordinate gradient descent scheme they used for successive independent actor and critic steps can simplify the calculation, but it causes serious bias. This paper proposes a flexible fully decentralized actor-critic MARL framework, which can combine most of the actor-critic methods and handle large-scale general cooperative multi-agent settings. A primal-dual hybrid gradient descent type algorithm framework is designed to learn indi-

vidual agents separately for decentralization. From the perspective of each agent, policy improvement and value evaluation are jointly optimized, which can stabilize multi-agent policy learning. Furthermore, the proposed framework can achieve scalability and stability for the large-scale environment. This framework also reduces information transmission by the parameter sharing mechanism and novel modeling-other-agents methods based on theory-of-mind and online supervised learning. Sufficient experiments in cooperative Multi-agent Particle Environment and StarCraft II show that the proposed decentralized MARL instantiation algorithms perform competitively against conventional centralized and decentralized methods.

Keywords: cooperative MARL, decentralized, actor-critic, primal-dual method

1. Introduction

Multi-agent reinforcement learning (MARL, Zhang et al. (2019)) has shown remarkable performance in interactive and complicated cooperative multi-agent environments, e.g. multi-robot controlling (Matignon et al., 2012) and multi-player games (Peng et al., 2017). MARL algorithms generally model a cooperative multi-agent learning system as a Markov game (Littman, 1994a) (or a stochastic game, (Guillermo, 1982)), where the joint actions of multiple agents influence a shared environment. In particular, each agent can access the full observation of the environment and takes action according to its current policy. These actions together determine the successive states of the environment (Lowe et al., 2017; Foerster et al., 2018a; Rashid et al., 2018).

However, the global assumption that each agent can fully observe the environment is usually difficult to satisfy in many practical applications, such as intelligent connected vehicle (Adler and Blue, 2002). Hence, a more reasonable solution is to model the problem as a more general formulation, i.e., cooperative partially observable stochastic game (POSG) (Hansen et al., 2004). Various MARL methods have been proposed, including value-based (Jiang et al., 2018; Rashid et al., 2018), actor-critic-based (Foerster et al., 2018a; Wei et al., 2018; Iqbal and Sha, 2019). However, these methods are designed to solve a fully cooperative POSG with another global assumption, i.e., all agents share a global cost function. Therefore, some works try to solve the more general cooperative POSG problem, where the cost functions of the agents might correspond to different tasks, and are only known to the corresponding agent. The collective goal of the agents is to minimize the globally summed return (Jiang and Lu, 2018; Yang et al., 2018b; Li et al., 2019). This paper is focused on this general cooperative setting.

Existing MARL algorithms for cooperative POSG mostly follow two frameworks (Figure 1): Centralized Training Decentralized Execution (CTDE) (Oliehoek et al., 2008) and Decentralized Training Decentralized Execution (DTDE). CTDE assumes that a powerful central controller can receive and process all agents' information. The left of Figure 1 shows the CTDE framework of actor-critic-based MARL methods, where centralized training explicitly takes into account the observations, policies, and costs of all agents in the learning phase, thereby effectively solving the non-stationary environment problem in multi-agent reinforcement learning. However, CTDE suffers from some limitations. First, as the number of agents increases, the amount of information for the centralized controller to process will increase exponentially, eventually leading to the *curse of dimensionality*. This will bring heavy space-time overhead to the entire system. Second, the centralized assumption for

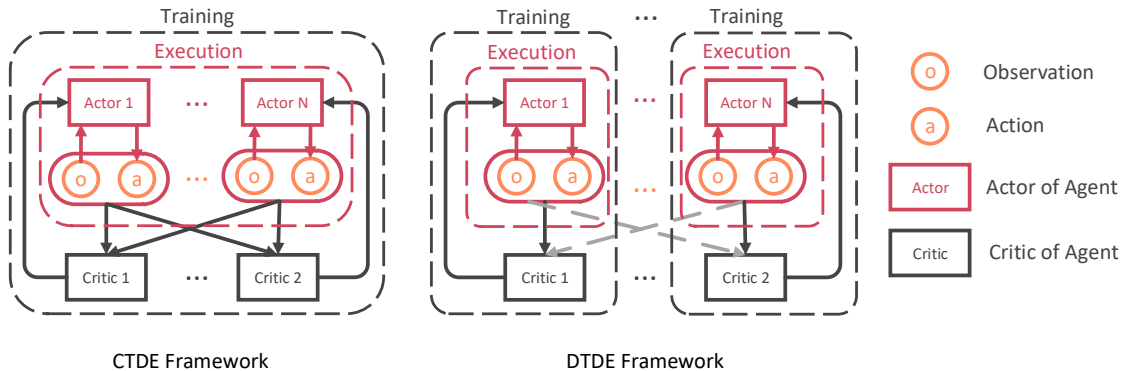


Figure 1: Centralized training decentralized execution (CTDE) and decentralized training decentralized execution (DTDE) actor-critic frameworks. The gray dotted arrow indicates that the relevant information of the agents is no longer scheduled by a centralized controller but is passing between agents by communication.

training is unrealistic in many real-world scenarios. Besides, the presence of a centralized controller also degrades the system’s capability to resist malicious attacks.

Many recent works try to follow the DTDE framework (Figure 1) to decentralize cooperative MARL (Zhang et al., 2018c; Doan et al., 2019a,b; Suttle et al., 2019). Unfortunately, existing decentralized cooperative MARL methods either only focus on the policy evaluation stage in fully observable small-scale multi-agent problems or only could handle the fully cooperative POSG setting where massive information needs to be transmitted in decentralized training. Therefore, our motivation is to design a more general DTDE framework for general cooperative POSG problems. To achieve a fully decentralized MARL algorithm, the main challenge need to be tackled is: How to effectively utilize partial observation of each agent to make the global decision in a fully decentralized scheme? In other words, it needs to drive the performance of decentralized MARL equal or close to the centralized one.

In this paper, we propose a Flexible Fully-decentralized Approximate Actor-critic (F2A2) algorithm for DTDE cooperative MARL under a novel and general joint actor-critic framework in Figure 2. A fully decentralized mechanism is designed based on consensus constraints and primal-dual optimization, whose benefit is that it can introduce a parameter sharing mechanism to increase the efficiency and adaptability for different settings. Moreover, to reduce the effect of the information loss caused by decentralized settings, a novel modeling-other-agents (MOA) technique is adopted based on theory-of-mind (TOM) (Rabinowitz et al., 2018) and online supervised learning, which enables the agent to estimate the information of other agents while making decisions based on local information, and improves the robustness and performance of the F2A2. Several decentralized versions of typical MARL algorithms are devised in the proposed framework, i.e., F2A2-COMA, F2A2-DDPG, F2A2-TD3, and F2A2-SAC. Extensive experiments on Cooperative Multi-agent Particle Environment and StarCraft II show that the proposed fully decentralized algorithms can

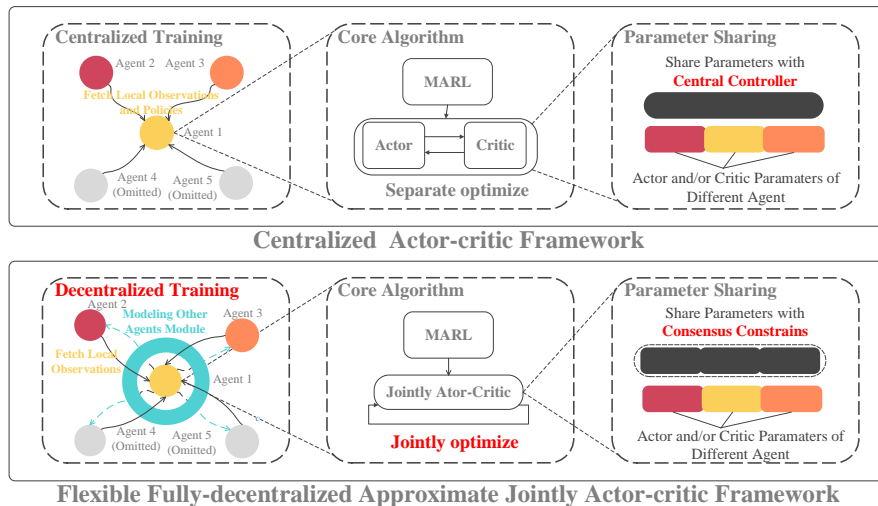


Figure 2: The F2A2 architecture. Compared with the upper traditional centralized training methods, F2A2 solves general cooperative POSG by a novel additive joint policy optimization objective and decentralizes via a separable primal-dual algorithm. The flexibility and superiority of the joint actor-critic framework allow it to combine any on/off-policy actor-critic algorithms with the advantages of the trust region. F2A2 contains a novel agent modeling scheme in decentralized training to improve the communication efficiency. Parameter sharing mechanism and consensus constraints enhance generality and scalability.

obtain competitive performance against conventional centralized and decentralized methods. Overall, the main contributions of this paper are:

- *Fully decentralized framework.* We reformulate actor-critic for general cooperative POSG problems in an additive joint form. Its separable characteristics lead to a novel fully decentralized actor-critic framework cooperating with separable primal-dual optimization. This is the first fully decentralized MARL to our knowledge.
- *Flexibility.* The proposed novel actor-critic framework is compatible with various actor-critic algorithms, and the general decentralizing mechanism can transform them into fully decentralized versions. Besides, the proposed framework has a flexible parameter-sharing and regularization mechanism, which makes F2A2 suitable for different kinds of settings, including small- and large-scale cooperative scenarios, on- and off-policy training schemes.
- *Performance.* A sufficient comparison is made between existing centralized and decentralized algorithms with F2A2. The proposed decentralized solutions have achieved remarkable performance with a more general setting even compared with its centralized version.

The following is the roadmap of this paper. Section 2 provides a brief but complete introduction to related works in centralized MARL and decentralized MARL and provide the

background material on cooperative POSG and actor-critic algorithm in Section 3. Section 4 describes the proposed model. The fully decentralized actor-critic MARL algorithms and related analysis from various aspects are presented in Section 5. Extensive experiments are presented in Section 6, and we conclude this paper in Section 7.

2. Related Work

The cooperative multi-agent problem exists widely in real world, such as multi-robot control (Matignon et al., 2012), multi-player games (Peng et al., 2017) and etc. In this section, existing cooperative multi-agent reinforcement learning methods is divided into *centralized* and *decentralized* methods based on their training mechanism. In this paper, the decentralized algorithm means it satisfies decentralization in both *training* and *execution* stages.

2.1 Centralized Cooperative MARL.

Most of the recent centralized cooperative MARL algorithms are follow the centralized training decentralized execution (CTDE) mechanism. The rise of the CTDE framework mainly has two-fold reasons. At early stages, the cooperative multi-agent problem is often modeled as a sizeable single-agent problem via joint action learning (JAL) (Claus and Boutilier, 1998), by learning a centralized policy based on global state or joint observation and then executing the joint action. It has to assume that the agents can access the global state, or there is a communication channel to integrate the individual information of all agents, no matter for training or execution. Even if the above problems can be solved by technical means, when the number of agents grows, the joint action space’s size will increase exponentially, leading to the centralized policy learning infeasible. By contrast, independent learning (IL) (Tan, 1993) can learn decentralized policies, but it results in nonstationarity as each agent treats the other agents as part of its environment. To solve the infeasible policy learning and the non-stationary problem, Oliehoek et al. (2008) introduces the CTDE mechanism. Specifically, in the training stage, the agents can freely share individual observations, parameters, and/or gradients, similar to JAL. However, in the execution stage, the policies are decentralized carried out with local observations, just similar in IL.

The algorithms that follow the CTDE mechanism can be divided into two categories: value-based methods and actor-critic-based methods. For value-based methods, some works (Rashid et al., 2018; Son et al., 2019) obtain a decomposable joint Q-value function through centralized training and only use the local Q-value function for each agent in the execution stage. The graph convolutional method (Jiang et al., 2018) introduces the graph neural networks (GNN) (Scarselli et al., 2008) to factorize the joint Q-value function without decomposable assumption. Yang et al. (2018a) and Wang et al. (2020) introduce the mean-field approximation to mitigate the curse of dimensionality of joint Q-value function. The actor-critic-based algorithms are more suitable for the CTDE training mechanism than the value-based methods because of their unique structure composed of the actor and the critic. Actor-critic-based multi-agent reinforcement learning algorithms generally train a central-critic (based on individual observations and policies of all agents) centrally and use the actor, which is only based on the local observation during decentralized execution for each agent. COMA (Foerster et al., 2018a) learns a shared centralized counterfactual baseline for all agents, which addresses the credit assignment problem in cooperative multi-agent

learning. Wei et al. (2018) extends the Soft Q-Learning (Haarnoja et al., 2017) to multi-agent to solve relative over-generalization problem. A line of works (Lowe et al., 2017; Li et al., 2019; Reddy et al., 2019) extend the DDPG (Lillicrap et al., 2016) algorithm to multi-agent scenarios. Unlike previous methods, MADDPG (Lowe et al., 2017) centrally trains a central critic for each agent. M3DDPG (Li et al., 2019) introduces the minimax method to maximize the expected returns when the other agents have the lowest returns. RC-MADDPG (Reddy et al., 2019), in addition to maximizing the expected return of the agent, additionally introduces the variation of return (VOR) as an optimization goal, making the policy performance of the agent more stable. Ryu et al. (2018) introduced an additional generative cooperative policy network, which is used to generate the more diverse samples via maximizing the expected return of other agents, to encourage the exploration based on Lowe et al. (2017). MAAC (Iqbal and Sha, 2019), which is also based on MADDPG, introduces the attention mechanism (Bahdanau et al., 2015) to model centralized non-shared critic in actor-critic framework, enabling the agent to self-identify information quality.

In recent years there has been another main line of work to solve the cooperative POSG problem under CTDE framework, i.e. *learning to communicate* (Foerster et al., 2016; Sukhbaatar et al., 2016; Peng et al., 2017; Jiang and Lu, 2018; Sheng et al., 2020) and (Chu et al., 2020). The cooperation between agents is accomplished by passing information to each other, and the information is generated by a shared message generator which is training centrally.

2.2 Decentralized Cooperative MARL.

For the current mainstream CTDE framework, the existence of an unavoidable central controller during training makes the framework have many performance limitations, such as single points of failure, high communication requirements, massive computing burden, and limited flexibility and scalability. In order to solve the problems caused by the CTDE framework, there are also some works on fully decentralized frameworks in recent years. From the perspective of optimization goals, these methods are divided into two categories: *policy evaluation* and *optimal policy learning*. For methods that belong to policy evaluation, their purpose is to learn the optimal value function corresponding to a fixed multi-agent policy (the fixed policy does not have to be optimal). For the methods that belong to the optimal policy learning, the goal is the same as all the above-mentioned centralized methods, to learn an optimal control policy for all agents. Therefore, *the policy evaluation method can be regarded as a component of the optimal policy learning method*.

For policy evaluation methods, the objective of all agents is to jointly minimize the mean square projected Bellman error (MSPBE). Macua et al. (2015) and Stanković and Stanković (2016) are fully decentralized multi-agent extensions of gradient temporal difference (GTD-2) (Sutton et al., 2009) and linear temporal-difference with gradient correction (TDC) (Sutton et al., 2009). Lee et al. (2018) also develops a fully decentralized multi-agent extension of GTD-2 and using the ordinary differential equation (ODE) method to establish the asymptotic convergence. Wai et al. (2018) proposes a double averaging scheme that combines the dynamic consensus (Qu and Li, 2017) and the SAG algorithm (Schmidt et al., 2017) to solve the distributed saddle-point version of MSPBE and achieving the linear convergence rate. More recently, standard TD learning (Tesauro, 1995), instead of gradient-

TD, has been generalized to this MARL setting, with particular focuses on finite-sample analyses (Doan et al., 2019a).

For optimal policy learning methods, some early value-based multi-agent reinforcement learning algorithms include Team Q-learning (Littman, 2001), Distributed Q-learning (Lauer and Riedmiller, 2000), FMQ (Kapetanakis and Kudenko, 2002), Hyper-Q learning (Tesauro, 2003), OAB (Wang and Sandholm, 2002), Hysteretic Q-learning (Matignon et al., 2007) and Lenient Q-learning (Panait et al., 2008). All these works are based on tabular Q-learning. Littman (2001); Lauer and Riedmiller (2000); Kapetanakis and Kudenko (2002); Wang and Sandholm (2002); Matignon et al. (2007); Panait et al. (2008) only consider the fully cooperative POSG setting and works Lauer and Riedmiller (2000); Matignon et al. (2007); Panait et al. (2008) have the limitation that they only work in deterministic environments.

In recent years, Kar et al. (2013) combines the idea of consensus and innovation to the standard Q-learning algorithm, proposes the QD -learning algorithm. Arslan and Yüksel (2017) proposed a decentralized Q-learning algorithm for fully observable stochastic games with weakly acyclic and verified its effectiveness in some tabular cases. The algorithm is fully decentralized in that each decision-maker has access to only its local information, the global state information, and the local cost function. Zhang et al. (2018a) proposed two decentralized Q-learning algorithms which extended the fitted-Q iteration algorithm for single-agent RL into a multi-agent tabular-based scenario with both the cooperative and the competitive settings. Each agent maintains a local estimation of the global average value function by exchanges local information over a communication network to keep the decentralization. Yongacoglu et al. (2019) proposed a variant of Q-learning algorithm for fully observable and fully cooperative stochastic games, which is conduct policy evaluation and policy improvement at two-timescale, and verified its team optimality in some tabular cases under the same decentralized settings as in Arslan and Yüksel (2017). Cassano et al. (2019) also proposed a decentralized tabular Q-learning algorithm, in which each agent maintains a local estimation of the global average value function for fully observable stochastic games under the same communicable settings as in Zhang et al. (2018a,c). Zeng et al. (2020) proposed a similar MARL algorithm with the value function consensus constraint as Zhang et al. (2018a,c), expect that it is a Q-learning algorithm and is to solve the multi-task fully observable stochastic game.

Compared with the value-based MARL methods, the actor-critic-based algorithms are also more applied to decentralized works. Some works follow a local actor and a consensus critic update scheme. Cassano et al. (2018) and Suttle et al. (2019) propose a fully decentralized policy gradient and actor-critic method to solve fully cooperative POSG problems separately. Zhang et al. (2018c) derives two multi-agent actor-critic algorithms for policy optimization, which employs a distributed evaluation strategy by combining diffusion learning (Sayed, 2014) and TD, to solve the fully observable stochastic games. This method introduces each agent’s local estimation of the counterfactual baseline and uses a consensus constraint to make the local estimation as accurate as possible compared to the centralized counterfactual baseline. Although the above methods are decentralized for fully cooperative POSG (except for Zhang et al. (2018c), which is proposed for general cooperative problem), they still assume that all agents share value functions. The massive information exchange makes the above method unable to extend to large-scale multi-agent scenarios. Zhang and Zavlanos (2019) is not the same as Zhang et al. (2018c) to learn the local estimate of the

centralized optimal critic and then learn the optimal policy independently. Instead, it learns the local estimate of the optimal centralized policy and then learns the optimal critic independently. Later in Zhang et al. (2018b), the same idea as Zhang et al. (2018c) is extended to the continuous control setting. Zhang et al. (2018b) develop an on-policy actor-critic algorithm, using the recent development of the expected policy gradient method (Ciosek and Whiteson, 2018). Qu et al. (2019) is also working under the problem setting of Zhang et al. (2018c), which extends Dai et al. (2018) to the multi-agent scenario. To solve a MARL problem with safety constraints, Lu et al. (2021) formulate the problem as a distributed constrained Markov decision process with networked agents and proposes a decentralized policy gradient method based on Zhang et al. (2018c). Qu et al. (2022) extends Zhang et al. (2018c) to large-scale multi-agent scenarios by introducing truncated Q-functions and Hu et al. (2021) further introduces reward machines on the basis of Qu et al. (2022) to solve complex cooperation problems.

Remark 1. Although these actor-critic-based methods are decentralized, they still assume that all agents share value functions or policies. The massive information exchange makes them unable to extend to large-scale multi-agent scenarios. Meanwhile, the above actor-critic-based methods optimize the actor and critic *separately* via block coordinate gradient descent (BCGD) type methods according to their related problem formulations. This formulation is composed of successive independent actor and critic steps and ignores the influence between each other, which can introduce bias. The biggest difference between these methods and the proposed algorithm framework is that a primal-dual hybrid gradient (PDHG) type method is introduced to *jointly* optimize the actor and critic, thereby avoiding error accumulation and the unstable problem in practical deployment of methods above. As will be shown, the proposed method not only solves the more general cooperative POSG problem, but also avoids direct sharing of policies between agents via a novel MOA module.

3. Preliminaries

3.1 Markovian and Non-Markovian Policies and Values

Under a few conditions (Puterman, 2014), for any given MDP, there exists an optimal policy which is deterministic and Markovian (that maps agent’s immediate or current observation of the environment to actions). In contrast, when dealing with partially observed MDPs (POMDPs), we might be interested in the class of non-Markovian and history-dependent policies (that map the state-action trajectory history to distributions over actions). In this paper, we do not employ non-Markovian policies for several reasons.

Firstly, tackling non-Markovian and history dependent policies can be computationally undecidable (Madani et al., 1999) for the infinite horizon, or PSPACE-Complete (Papadimitriou and Tsitsiklis, 1987) in the finite horizon POMDPs. When extended to the POSG, this problem will become more serious. To avoid undesirable computational burdens, we focus on Markovian policies.

Secondly, indeed, many prior works study Markovian policies for POMDPs (Baxter and Bartlett, 2001; Littman, 1994b; Baxter et al., 2000; Azizzadenesheli et al., 2016) where, in general, the optimal policies are stochastic (Littman, 1994b; Singh et al., 1994; Montufar et al., 2015). Acknowledging the computation complexity of Markovian policies (Vlassis

et al., 2012), this line of work highlights the broad interest, importance, and applicability of Markovian policies.

Thirdly, one could alternatively consider history-dependent policies, which is a richer function class than Markovian policies. However, utilizing history-dependent policies requires turning a given POMDP problem to a potentially non-stationary MDP with states as concatenations of the historical data. As such, a thorough theoretical treatment of history-dependent policies for POMDPs would likely require leveraging theoretical analyses for non-stationary MDPs. When extended to the POSG, the theoretical analysis of non-stationarity of POSG is still in a very preliminary stage. So much discussion in this area is beyond the scope of this paper, and we will delve into this in future work.

Fourthly, and most importantly, the expressiveness of general policy classes is mainly entangled with the definition of the observation. Under some regularity conditions, for any given class of history-dependent policies on a POMDP, there exists a class of Markovian policies on a new POMDP (Littman, 1994b; Castro et al., 2009; Hausknecht and Stone, 2015; Azizzadenesheli et al., 2018) such that: (1) the observations of the new POMDP are the (typically) discrete concatenations of the historical data in the former POMDP; and (2) the two policy classes are equivalent, i.e., the respective policies result in the same behavior (e.g., action sequence). In other words, instead of making the policies non-Markovian and history-dependent, we can keep them Markovian and instead enrich the observation space. Similarly, for any limited-memory policy class (depending on a fixed window of history instead of the whole history, e.g., the policies in MDPs of order more than one), there is an enrichment of the observation that results in an equivalent class of Markovian policies. This viewpoint is known as the emergentism approach (O’Connor, 2021).

Similar to the policy class, in this paper, we also focus on the Markovian value function. Recent work (Mao et al., 2020) in cooperative POSG has shown that, with appropriate enrichment of the observation space, the optimal state-value function (V) or state-action-value function (Q) obtained under the newly constructed POSG can approximate the optimal one obtained under the original POSG. However, Mao et al. (2020) adopts a value-based approach. Fortunately, after enriching the observation space, we can naturally extend the policy gradient theorem under MDP to POMDP based on Mao et al. (2020) as the basis for the subsequent theoretical derivation of this paper¹.

The above reasoning implies that, by considering the class of Markovian policies and value functions in episodic POMDPs, we often will not restrict the generality of the results in this paper and similar conclusions can also be obtained by extending to the POSG. Therefore, despite focusing on Markovian policies and value functions, our results also hold for both classes of limited-memory as well as history-dependent policies and value functions through representing histories as observations. So in this paper, both o and \mathbf{O} refer to the enriched observation spaces unless specified.

3.2 Problematic Instances of POMDPs

Certain aberrant instances of worst-case scenarios in POMDP can potentially render the proposed F2A2 computationally infeasible. Specifically, in the realm of POMDPs, the identification of the system dynamics is hindered by pathological instances. Such cases are

1. For the specific derivation process, please refer to the Appendix. C

characterized by observations that offer no valuable information for discerning the system’s behavior. The arduousness of the situation was proven by Krishnamurthy et al. (2016), who demonstrated that, in the worst-case scenario, the discovery of a near-optimal policy for a POMDP necessitates an exponential number of samples proportional to the horizon. This result is formally expressed through the following proposition, where \mathcal{A} represents the joint action space with a cardinality of $|\mathcal{A}| = A$.

Proposition 1 (Krishnamurthy et al. 2016; Jin et al. 2020) *There exists a class of 2-states T -horizon POMDPs whose observations reveal no information about the underlying states up to the end, such that any algorithm requires at least $A^{\Omega(T)}$ samples to learn an $\mathcal{O}(1)$ -optimal policy with a probability of $1/2$ or higher.*

A running example can be found in Liu et al. (2022, §3.1 Hard instances of POMDPs). For POMDPs with more than two states, the proposition presented earlier can be easily extended to situations where two mixtures of latent states with disjoint support exist, causing observations to fail to differentiate between the two mixtures. A mixture of states is represented by a probability vector $\nu \in \Delta_S$, where ν_1 and ν_2 are considered to have disjoint support if $\text{supp}(\nu_1) \cap \text{supp}(\nu_2) = \emptyset$.

To avoid the aforementioned pathological instances, a simple approach is to assume that any two latent state mixtures ν_1 and ν_2 with disjoint support yield distinct distributions over observations, meaning that $\mathbb{O}_t \nu_1 \neq \mathbb{O}_t \nu_2$ for all $t \in [T]$, where \mathbb{O}_h is the $O \times S$ emission matrix at step h . Here, \mathcal{S} and \mathcal{O} denote the state and joint observation spaces, respectively, with cardinalities of $|\mathcal{S}| = S$ and $|\mathcal{O}| = O$. It can be shown using a linear algebraic argument that this condition is equivalent to the rank of the emission matrix \mathbb{O}_t being S .

Proposition 2 (Jin et al. 2020) *The emission matrix \mathbb{O}_h is rank S if and only if the induced distributions over observations are distinct for any two mixtures of latent states with disjoint support.*

Building upon the proposition discussed earlier, Jin et al. (2020) suggested a more robust assumption to guarantee that observations possess adequate information to differentiate any two-state mixtures, given a sufficiently large number of samples.

Assumption 1 (α -weakly revealing condition, Liu et al. 2022) *There exists $\alpha > 0$, such that $\min_t \sigma_S(\mathbb{O}_t) \geq \alpha$.*

Liu et al. (2022) refer to Assumption 1 as the “weakly” revealing condition, in contrast to the rich observation or block MDP setup in the literature. However, it should be noted that Assumption 1 implicitly assumes that $S \leq O$, as \mathbb{O}_t is a matrix of size $O \times S$. Thus, it only holds in the *undercomplete* setting. To ensure the generality of F2A2, we adopt the *overcomplete* version assumption proposed by (Liu et al., 2022) where $S > O$.

Observations in a single step are insufficient to distinguish between any two mixtures of latent states due to information-theoretic limitations. Instead, we must examine the distribution of observations for a sequence of m consecutive steps corresponding to the enriched observation space in F2A2. It is worth noting that the number of possible joint observable sequences of length m , $(\mathbf{o}_1, \mathbf{a}_1, \dots, \mathbf{a}_{m-1}, \mathbf{o}_m)$, is $O^m A^{m-1}$, which is greater

than S when $m \geq \Omega(\log S)$. To formalize this, we define the m -step emission-action matrices $\left\{ \mathbb{M}_t \in \mathbb{R}^{(A^{m-1}O^m) \times S} \right\}_{t \in [T-m+1]}$, and introduce the following assumption.

Assumption 2 (m -step α -weakly revealing condition, Liu et al. 2022) *There exists $m \in \mathbb{N}, \alpha > 0$ such that $\min_{t \in [T-m+1]} \sigma_S(\mathbb{M}_t) \geq \alpha$.*

Assumption 2 guarantees that the observable sequence in the next m consecutive steps contains enough information to distinguish any two mixtures of states when a sufficiently large number of enriched observations is available. This assumption encompasses Assumption 1 as a specific case with $m = 1$.

3.3 Cooperative Partially Observable Stochastic Games

Partially observable stochastic game (POSG) (Hansen et al., 2004) is denoted as a tuple based on Markov Game as follows:

$$\langle \mathcal{X}, \mathcal{S}, \{\mathcal{A}^i\}_{i=1}^n, \{\mathcal{O}^i\}_{i=1}^n, \mathcal{P}, \mathcal{E}, \{\mathcal{C}^i\}_{i=1}^n \rangle,$$

where n is the total number of agents, \mathcal{X} represents the agent space, \mathcal{S} is a finite set of states, \mathcal{A}^i is a finite action set of agent i , $\mathcal{A} = \mathcal{A}^1 \times \mathcal{A}^2 \times \dots \times \mathcal{A}^n$ is the finite set of joint actions, $\mathcal{P}(s'|s, \mathbf{a})$ is the Markovian state transition probability function, \mathcal{O}^i is a finite observation set of agent i , $\mathcal{O} = \mathcal{O}^1 \times \mathcal{O}^2 \times \dots \times \mathcal{O}^n$ is the finite set of joint observations, $\mathcal{E}(\mathbf{o}|s)$ is the Markovian observation emission probability function, and $\mathcal{C}^i : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow \mathbb{R}$ is the cost function of agent i .

The game in POSG unfolds over a finite or infinite sequence of stages (or timesteps), where the number of stages is called *horizon*. This paper considers the episodic infinite horizon problem. The objective for each agent is to minimize the expected discounted cumulative cost received during the game. For a cooperative POSG, the definition in (Song et al., 2020) is quoted,

$$\forall x \in \mathcal{X}, \forall x' \in \mathcal{X} \setminus \{x\}, \forall \pi_x \in \Pi_x, \forall \pi_{x'} \in \Pi_{x'}, \frac{\partial \mathcal{C}^{x'}}{\partial \mathcal{C}^x} \geq 0,$$

where x and x' are a pair of agents in agent space \mathcal{X} ; π_x and $\pi_{x'}$ are the corresponding policies in the policy space Π_x and $\Pi_{x'}$ separately. Intuitively, this definition means that there is no conflict of interest for any pair of agents. The most common example of cooperative POSG is the fully cooperative POSG (also called decentralized partially observable Markov decision process, Dec-POMDP), that all the agents share the same global cost at each stage, and $\mathcal{C}^1 = \mathcal{C}^2 = \dots = \mathcal{C}^n$.

This paper aims to solve the general cooperative POSG. Each agent completes a common task based on their local observations, cost, and learning process. Without loss of generality, the optimization goal of the general cooperative POSG problem is defined as follows

$$\min_{\Psi} \sum_{i=1}^n \sum_{t=0}^{\infty} \mathbf{E}_{s_0 \sim p_0, \mathbf{o} \sim \mathcal{E}, \mathbf{a} \sim \pi_{\Psi}} [\gamma^t \mathcal{C}_{t+1}^i], \quad (3.1)$$

where $\Psi := \{\psi^i\}_{i=1}^n$ denotes parameters of the approximated policy $\pi_{\psi^i}^i : \mathcal{O}^i \rightarrow \mathcal{A}^i$ of all agents and $\pi_{\Psi} := \prod_{i=1}^n \pi_{\psi^i}^i$ represents the factorizable joint policy of all agents. γ is the

discount factor. Note p_0 is the distribution of the initial state s_0 . c_{t+1}^i represents the reward received by the agent i at timestep $t + 1$ after executes action a_t^i in local observation o_t^i .

3.4 Single-agent Actor-critic-type Algorithms for POMDP

Single-agent actor-critic methods optimize actor $\pi(\cdot|o; \psi)$ directly by minimizing the expected discounted accumulated cost:

$$\sum_{t=0}^{\infty} \mathbf{E}_{s_0 \sim p_0, o \sim \mathcal{E}, a \sim \pi_\psi} [\gamma^t c_{t+1}], \quad (3.2)$$

where ψ is the parameters of the approximated policy, p_0 is the distribution of the initial state s_0 . However, this may lead to a high estimation variance with policy gradient methods. Instead, a critic $V_\phi(o)$ is introduced to estimate the expected accumulate cost. The optimization formulation of the vanilla actor-critic algorithm (Sutton et al., 2000) can be reformalized into a bi-level problem as follows (Yang et al., 2018c, 2019; Hong et al., 2020):

$$\begin{aligned} \psi^* &= \arg \min_{\psi} J_{\text{actor}}(\psi, \phi^*(\psi)) := \mathbf{E}_{s_0 \sim p_0, o_0 \sim \mathcal{E}} [V_{\phi^*}^\pi(o_0)], \\ \text{where } \phi^*(\psi) &= \arg \min_{\phi} J_{\text{critic}}(\psi, \phi) := \mathbf{E}_{s \sim d_\psi, o \sim \mathcal{E}} [(V_\phi^\pi(o) - V_{tg}^\pi(o))^2], \end{aligned} \quad (3.3)$$

where $V_{tg}^\pi(o) := \mathbf{E}_{a \sim \pi_\psi, s' \sim \mathcal{P}, o' \sim \mathcal{E}} (\mathcal{C}_{s,a}^{s'} + \gamma V_\phi^\pi(o'))$ and $\mathcal{C}_{s,a}^{s'}$ is equivalent to $\mathcal{C}^i(s, a, s')$ defined in previous subsection. d_ψ is the distribution of the state-occupancy measure of policy π_ψ .

Note the second term of the right side of the Eq. (3.3) is the *bellman error* and the $V_{tg}(\cdot)$ stands for the temporal difference target in it. It can be seen that the first term of the right side of the Eq. (3.3) is equivalent to Eq. (3.2). The critic $V_\phi(\cdot)$ is introduced to estimate the expected accumulated cost to reduce the variance. Moreover, correspondingly the second term, bellman error term, is used to fit the introduced critic $V_{tg}(\cdot)$. Traditional actor-critic algorithms (Lillicrap et al., 2016; Schulman et al., 2015, 2017; Haarnoja et al., 2018; Fujimoto et al., 2018), minimizes Eq. (3.3) where updates the actor and the critic parameters alternatively as follows,

$$\min_{\psi} J_{\text{actor}}(\psi, \phi) := \mathbf{E}_{s_0 \sim p_0, o_0 \sim \mathcal{E}} [V_\phi^\pi(o_0)], \quad (3.4a)$$

$$\min_{\phi} J_{\text{critic}}(\psi, \phi) := \mathbf{E}_{s \sim d_\psi, o \sim \mathcal{E}} [(V_\phi^\pi(o) - V_{tg}^\pi(o))^2]. \quad (3.4b)$$

The algorithm scheme (3.4) can be considered as applying the block coordinate gradient descent (BCGD) type algorithm on (3.3).

It is worth noting that when solving (3.4b), the bootstrapping method is generally used (Sutton and Barto, 2018). Bootstrapping methods are not in fact instances of true gradient descent (Barnard, 1993). They take into account the effect of changing the parameter ϕ on the estimated value, but ignore its effect on the target. They include only a part of the gradient and, accordingly, are called semi-gradient methods. Although semi-gradient (bootstrapping) methods do not converge as robustly as gradient methods, they do converge reliably in important cases such as the linear case (Sutton and Barto, 2018). Moreover, they offer important advantages that make them offer clearly preferred. One reason for this is

that they typically enable significantly faster learning. Another is that they enable learning to be continual and online, without waiting for the end of an episode. This enables them to be used on continuing problems and provides computational advantages.

4. Joint Actor-Critic MARL Framework

Actor-critic type algorithms are popular for MARL, which combines the advantages of both policy gradient and value-based methods, often being more tractable and efficient in either high-dimensional or continuous action space. Therefore, many existing MARL methods (Foerster et al., 2018a; Wei et al., 2018; Iqbal and Sha, 2019) are based on the actor-critic framework.

For general cooperative POSG, this paper first draw on the work of Yang et al. (2018c), Yang et al. (2019) and Hong et al. (2020) to look at the actor-critic algorithm from a bi-level perspective and extend it to multi-agent scenarios,

$$\min_{\{\psi^i\}} \sum_{i=1}^n \mathbf{E} \left[V_{\phi^{i,*}}^{\pi,i}(\mathbf{o}_0) \right] \quad s.t. \quad \{\phi^{i,*}\} = \arg \min_{\{\phi^i\}} \sum_{i=1}^n \mathbf{E} \left[\left(V_{\phi^i}^{\pi,i}(\mathbf{o}) - V_{tg}^{\pi,i}(\mathbf{o}) \right)^2 \right], \quad (4.1)$$

where the expectations are taken on $s_0 \sim d_0, s \sim d_\Psi, \mathbf{o}_0, \mathbf{o} \sim \mathcal{E}$; $\{\phi^i\}$ and $\{\psi^i\}$ denotes the critic parameters and the actor parameters of all agents respectively, which are alternatively optimized; $V_{tg}^{\pi,i}(\mathbf{o})$ is defined as

$$\mathbf{E}_{a \sim \pi_\Psi, s' \sim \mathcal{P}, \mathbf{o}' \sim \mathcal{E}} \left(C_{s,a}^{s',i} + \gamma V_{\phi^i}^{\pi,i}(\mathbf{o}') \right).$$

Then, the above bi-level form can be simplified to a single-level problem with a penalty term and obtain the standard optimization formulation for multi-agent cases:

$$\min_{\{\psi^i\}, \{\phi^i\}} \left\{ \mathcal{J}_{ac}(\{\psi^i\}, \{\phi^i\}) := \alpha_1 \sum_{i=1}^n \mathbf{E} \left[V_{\phi^i}^{\pi,i}(\mathbf{o}_0) \right] + \alpha_2 \sum_{i=1}^n \mathbf{E} \left[\left(V_{\phi^i}^{\pi,i}(\mathbf{o}) - V_{tg}^{\pi,i}(\mathbf{o}) \right)^2 \right] \right\}. \quad (4.2)$$

For each agent i , the critic $V_{\phi^i}^{\pi,i}(\mathbf{o})$ is determined by all $\boldsymbol{\pi} = \prod_{i=1}^n \pi_{\psi^i}^i$ and the specific ϕ^i based on joint observation \mathbf{o} ; π^i represents the policy of each agent i . It is worth noting that each agent's critic $V_{\phi^i}^{\pi,i}(\mathbf{o})$ defined in Equation (4.2) are only based on their critic parameters ϕ^i but on all agents' actor parameters $\{\psi^i\}$. The reason is that in the general cooperative POSG problem, each agent has an independent cost function C^i mentioned in Section 3.1 but this cost function is based on the policy of all agents. Further, the objective function for each agent i in actor and critic parameter updating phases can be defined as follows respectively:

$$\begin{aligned} J_{\text{actor}}^i(\{\psi^i\}, \phi^i) &:= \mathbf{E} \left[V_{\phi^i}^{\pi,i}(\mathbf{o}_0) \right], \\ J_{\text{critic}}^i(\{\psi^i\}, \phi^i) &:= \mathbf{E} \left[\left(V_{\phi^i}^{\pi,i}(\mathbf{o}) - V_{tg}^{\pi,i}(\mathbf{o}) \right)^2 \right]. \end{aligned} \quad (4.3)$$

Therefore, the optimization problem (4.2) can be reformulated into

$$\min_{\mathbf{w}} \mathcal{J}(\mathbf{w}) := \underbrace{\alpha_1 \sum_{i=1}^n J_{\text{actor}}^i(\{\psi^i\}, \phi^i)}_{\mathcal{J}_{\text{actor}}(\mathbf{w})} + \underbrace{\alpha_2 \sum_{i=1}^n J_{\text{critic}}^i(\{\psi^i\}, \phi^i)}_{\mathcal{J}_{\text{critic}}(\mathbf{w})} + \mathcal{R}(\mathbf{w}), \quad (4.4)$$

where $\mathbf{w} := (\{\psi^i\}, \{\phi^i\})$ denotes all the parameters; α_1 and α_2 denote penalty factors in actor and critic, respectively; $\mathcal{R}(\cdot)$ denotes an *optional* regularizer to prevent over-fitting or for sparsity and etc., and it is assumed with separable structure in the proposed framework. Some works (Farebrother et al., 2018; Liu et al., 2019) have shown that model regularization techniques have a significant impact on the performance of single-agent reinforcement learning models. The scale of multi-agent reinforcement learning models is generally much larger than that of single-agent models. Thus the impact of regularization should not be ignored. Therefore, the regularization term $\mathcal{R}(\cdot)$ is usually added to the multi-agent actor-critic objective function to make the proposed framework more generalizable.

When solving the general cooperative POSG problem, in order to learn the decentralized policies to achieve global collaboration, the following three types of techniques are mainly used (Heider and Simmel, 1944; Rasouli et al., 2017; Schroeder de Witt et al., 2019). (1) *common knowledge* based algorithms (Brafman and Tennenholtz, 2003; Schroeder de Witt et al., 2019) use the common knowledge protocol to achieve global collaboration by establishing common knowledge about other agents’ actions or observations, based on the global common knowledge generator; (2) *explicit communication* based algorithms (Peng et al., 2017; Jiang and Lu, 2018) achieve certain consensus by explicit exchanging information in the decision phase. But the information exchanged need to be able to be understood by all agents, so the message generation module also needs to be shared between agents; (3) *implicit communication* algorithms (Foerster et al., 2018a; Iqbal and Sha, 2019) directly share individual action and/or observation between agents. At the same time, similar to the explicit communication-based algorithms, the model of each agent to process this global information must also be consistent with maintaining the same understanding of the environment for the achievement of global collaboration.

In order to propose an effective framework to be compatible with these collaboration skills, and considering that the core modules in these methods can be accessed globally, the flexible *parameter sharing* mechanism is introduced to the proposed framework. In addition, the parameter sharing mechanism also can be used for better algorithm scalability (Yang et al., 2018b). Recent work (Terry et al., 2020; Christianos et al., 2021; Grupen et al., 2021) also shows that parameter sharing plays a crucial role in improving algorithm performance. To achieve highly integration of the parameter sharing mechanism and the general cooperative POSG optimization objective, without loss of generalization, the parameters are reformulated and divided into shared and non-shared parts, denoted as \mathbf{w}_{in} and \mathbf{w}_{sh} respectively (“*in*” and “*sh*” denote “individual” and “shared” parameters, respectively). Then, the problem (4.4) can be reformulated into a more general form, i.e.,

$$\min_{\mathbf{w}_{in}, \mathbf{w}_{sh}} \mathcal{J}(\mathbf{w}_{in}, \mathbf{w}_{sh}) = \alpha_1 \mathcal{J}_{\text{actor}}(\mathbf{w}_{in}, \mathbf{w}_{sh}) + \alpha_2 \mathcal{J}_{\text{critic}}(\mathbf{w}_{in}, \mathbf{w}_{sh}) + \mathcal{R}(\mathbf{w}_{in}, \mathbf{w}_{sh}), \quad (4.5)$$

Algorithm	\mathbf{w}_{in}		\mathbf{w}_{sh}		$\mathcal{J}_{actor}, \mathcal{J}_{critic}$	
	ψ_{in}^i	ϕ_{in}^i	ψ_{sh}	ϕ_{sh}	$\mathcal{J}_{actor}^i(\{\psi_{in}^i\}, \phi_{in}^i, \psi_{sh}, \phi_{sh})$	$\mathcal{J}_{critic}^i(\{\psi_{in}^i\}, \phi_{in}^i, \psi_{sh}, \phi_{sh})$
MADDPG	$\pi^i(o^i)$	\emptyset	\emptyset	$Q(o, \mathbf{a})$	$\mathbb{E}[Q_{\phi_{sh}}^{\pi, i}(o, \mathbf{a})]$	$\mathbb{E}[(Q_{\phi_{sh}}^{\pi, i}(o, \mathbf{a}) - Q_{ig}^{\pi, i})^2]$
COMA	$\pi^i(o^i)$ (GRUs ²)	\emptyset	\emptyset	$\frac{Q(o, \mathbf{a})}{V(o)}$ (MLP ³)	$\mathbb{E}[Q_{\phi_{sh}}^{\pi, i}(o, \mathbf{a}) - \mathcal{B}(o, \mathbf{a}^i)]$	$\mathbb{E}[(Q_{\phi_{sh}}^{\pi, i}(o, \mathbf{a}) - Q_{ig}^{\pi, i})^2]$
MAAC	$\pi^i(o^i)$ (MLP)	ζ^i, Γ^i	\emptyset	$\{V_h, W_h^{key}, W_h^{que}\}_{h=1}^H$	$\mathbb{E}[Q_{\phi_{in}, \phi_{sh}}^{\pi, i}(o, \mathbf{a}) + \alpha \mathcal{H}(\cdot \pi_{\psi_{in}^i}^i(o^i)) - \mathcal{B}(o, \mathbf{a}^i)]$	$\mathbb{E}[(Q_{\phi_{in}, \phi_{sh}}^{\pi, i}(o, \mathbf{a}) - Q_{ig}^{\pi, i})^2]$

Table 1: Summary of MADDPG, COMA and MAAC algorithms.

where the general form of \mathcal{J}_{actor} , \mathcal{J}_{critic} and $\mathcal{R}(\mathbf{w}_{in}, \mathbf{w}_{sh})$ are as follows:

$$\begin{aligned}
 \mathcal{J}_{actor}(\mathbf{w}_{in}, \mathbf{w}_{sh}) &= \sum_{i=1}^n \mathcal{J}_{actor}^i(\{\psi_{in}^i\}, \phi_{in}^i, \psi_{sh}, \phi_{sh}), \\
 \mathcal{J}_{critic}(\mathbf{w}_{in}, \mathbf{w}_{sh}) &= \sum_{i=1}^n \mathcal{J}_{critic}^i(\{\psi_{in}^i\}, \phi_{in}^i, \psi_{sh}, \phi_{sh}), \\
 \mathcal{R}(\mathbf{w}_{in}, \mathbf{w}_{sh}) &= \sum_{i=1}^n \underbrace{(r_{\psi}^{se}(\psi_{in}^i) + r_{\phi}^{se}(\phi_{in}^i))}_{\mathcal{R}_{in}(\mathbf{w}_{in})} + \underbrace{(r_{\psi}^{sh}(\psi_{sh}) + r_{\phi}^{sh}(\phi_{sh}))}_{\mathcal{R}_{sh}(\mathbf{w}_{sh})},
 \end{aligned}$$

with $\mathbf{w}_{in} = (\{\psi_{in}^i\}, \{\phi_{in}^i\})$ and $\mathbf{w}_{sh} = (\psi_{sh}, \phi_{sh})$; the regularizer $\mathcal{R}(\cdot)$ is separable with $r_{\psi}^{se}, r_{\phi}^{se}, r_{\psi}^{sh}$ and r_{ϕ}^{sh} be related regularization functions. Many existing MARL algorithms are equivalent to solve this general optimization formulation (4.5). Some state-of-the-art algorithms, i.e. MADDPG (Lowe et al., 2017), COMA (Foerster et al., 2018a), MAAC (Iqbal and Sha, 2019) and etc are summarized in the following Table 1, which include all elements \mathbf{w}_{in} , \mathbf{w}_{sh} , \mathcal{J}_{actor} , \mathcal{J}_{critic} in framework (4.5) for them. More detailed derivation can be found in supplemental material.

Most MARL methods solve the actor-critic framework based on the block coordinate gradient descent (BCGD) type techniques, whose standard procedure is composed of successive independent actor and critic steps, i.e.,

$$\begin{cases}
 \text{Critic-step:} & \begin{pmatrix} \phi_{in}^i \\ \phi_{sh} \end{pmatrix} \leftarrow \begin{pmatrix} \phi_{in}^i \\ \phi_{sh} \end{pmatrix} - \gamma \begin{pmatrix} \nabla_{\phi_{in}^i} \mathcal{J}_{critic}(\mathbf{w}_{in}, \mathbf{w}_{sh}) \\ \nabla_{\phi_{sh}} \mathcal{J}_{critic}(\mathbf{w}_{in}, \mathbf{w}_{sh}) \end{pmatrix}; \\
 \text{Actor-step:} & \begin{pmatrix} \psi_{in}^i \\ \psi_{sh} \end{pmatrix} \leftarrow \begin{pmatrix} \psi_{in}^i \\ \psi_{sh} \end{pmatrix} - \gamma \begin{pmatrix} \nabla_{\psi_{in}^i} \mathcal{J}_{actor}(\mathbf{w}_{in}, \mathbf{w}_{sh}) \\ \nabla_{\psi_{sh}} \mathcal{J}_{actor}(\mathbf{w}_{in}, \mathbf{w}_{sh}) \end{pmatrix},
 \end{cases} \quad (4.6)$$

while the actor parameters $(\{\psi_{in}^i\}, \psi_{sh})$ and the critic parameters $(\{\phi_{in}^i\}, \phi_{sh})$ are fixed in critic and actor steps respectively. By the way, the three state-of-the-art MARL algorithms mentioned above follow this iterative scheme (4.6), and the details are presented in supplemental material. In this case, together with the existence of sharing parameters, most typical MARL algorithms are included in the CTDE algorithm framework. These algorithms have to maintain a globally accessible shared module in order to handle the parameter sharing. More important, when optimizing actor and critic parts separately, they ignored the influence between each other. Although this separation optimization scheme simplifies the solution calculation, it also introduces serious bias.

5. Flexible Fully-decentralized Approximate Actor-critic Framework

Current decentralized multi-agent actor-critic algorithms (Zhang et al., 2018b,c; Zhang and Zavlanos, 2019; Suttle et al., 2019) also use BCGD-type procedure to optimize actor and critic of each agent. For example, Zhang et al. (2018c) can be regarded as a distributed version of the COMA. The core idea of Zhang and Zavlanos (2019) is the opposite of Zhang et al. (2018c), and it also uses the same procedure as the former. In centralized training, the concentrative information collection and process might guarantee the global convergence to a certain degree (Konda and Tsitsiklis, 2000). However, in decentralized training, the hysteretic information exchange would cause error accumulation and an unstable problem in many practical deployments. The standard procedure needs to be modified to a more rational and flexible form to achieve full decentralization, satisfying the demand for the mutual observation, reward assignment, and policy interaction between agents, and synchronize the optimization. At the same time, the above algorithms generally assume that the local observations and policies of other agents are known when performing decentralized optimization. This constraint requires a large amount of inter-agent communication so that the above-mentioned decentralized algorithms cannot be extended to a large-scale multi-agent environment.

Based on the analysis above, the following issues need to be considered to propose a fully decentralized multi-agent actor-critic algorithm. First, the algorithm can solve error accumulation and instability problems caused by the BCGD-type procedure of the current multi-agent actor-critic algorithms under decentralized training. Second, it can work with fewer constraints. It does not require the precise policies of all other agents, thereby avoiding the high communication costs. Third, it can be flexibly combined with most actor-critic algorithms and compatible with on-policy and off-policy techniques.

Therefore, the fully decentralized algorithm framework is proposed to solve the general formulation above (4.5). Firstly, the consensus variables $\{\tilde{\mathbf{w}}_{sh}^i\}$ are introduced to help achieve the fully decentralized structure. Recall the definition of $\mathcal{J}_{\text{actor}}(\mathbf{w}_{in}, \mathbf{w}_{sh})$, $\mathcal{J}_{\text{critic}}(\mathbf{w}_{in}, \mathbf{w}_{sh})$ and the regularizer $\mathcal{R}(\mathbf{w}_{in}, \mathbf{w}_{sh})$, the comprehensive formulation can be obtained as follows:

$$\begin{aligned} \min_{\mathbf{w}_{in}, \{\tilde{\mathbf{w}}_{sh}^i\}, \mathbf{w}_{sh}} \quad & \alpha_1 \sum_{i=1}^n \mathcal{J}_{\text{actor}}^i(\mathbf{w}_{in}, \tilde{\mathbf{w}}_{sh}^i) + \alpha_2 \sum_{i=1}^n \mathcal{J}_{\text{critic}}^i(\mathbf{w}_{in}, \tilde{\mathbf{w}}_{sh}^i) + \mathcal{R}(\mathbf{w}_{in}, \mathbf{w}_{sh}), \\ \text{s.t.} \quad & \mathbf{w}_{sh} = \tilde{\mathbf{w}}_{sh}^i, \quad i = 1, \dots, n. \end{aligned} \tag{5.1}$$

This can be considered as an equivalent reformulation of the general formulation (4.5). (5.1) is a typical linear constrained optimization problem while the consensus constraints only relate to the shared parameters for all n agents. $\{\tilde{\mathbf{w}}_{sh}^i\}$ are introduced to help communicating the shared parameters and each $\tilde{\mathbf{w}}_{sh}^i$ belongs to agent i respectively. As a result, the primal-dual hybrid gradient (PDHG) type method (or inexact alternating direction method of multipliers (ADMM) type method) is utilized to solve (5.1) because the PDHG-type method naturally has decentralized computing architecture. The augmented Lagrangian

function is defined as follows,

$$\mathcal{L}(\mathbf{w}_{in}, \{\tilde{\mathbf{w}}_{sh}^i\}, \mathbf{w}_{sh}, \{\boldsymbol{\lambda}_i\}) = \underbrace{\alpha_1 \sum_{i=1}^n J_{\text{actor}}^i(\mathbf{w}_{in}, \tilde{\mathbf{w}}_{sh}^i) + \alpha_2 \sum_{i=1}^n J_{\text{critic}}^i(\mathbf{w}_{in}, \tilde{\mathbf{w}}_{sh}^i)}_{\sum_{i=1}^n J_{ac}^i(\mathbf{w}_{in}, \tilde{\mathbf{w}}_{sh}^i)} + \mathcal{R}(\mathbf{w}_{in}, \mathbf{w}_{sh}) - \sum_{i=1}^n \langle \boldsymbol{\lambda}_i, \mathbf{w}_{sh} - \tilde{\mathbf{w}}_{sh}^i \rangle + \frac{\beta}{2} \sum_{i=1}^n \|\mathbf{w}_{sh} - \tilde{\mathbf{w}}_{sh}^i\|^2, \quad (5.2)$$

where $\{\boldsymbol{\lambda}_i\}$ denote the Lagrangian dual variables concerning the consensus linear constraints with a unified penalty parameter β (in order to express more clearly, a unified penalty parameter is used, and it can also be modified into separate and different parameters). Motivated by the popular ADMM, the PDHG algorithm framework can be designed by alternatively calculate the primal variables \mathbf{w}_{in} , $\{\tilde{\mathbf{w}}_{sh}^i\}$, \mathbf{w}_{sh} and the dual variables $\{\boldsymbol{\lambda}_i\}$. Similar to the definition $\mathbf{w}_{in} = (\{\psi_{in}^i\}, \{\phi_{in}^i\})$ and $\mathbf{w}_{sh} = (\psi_{sh}, \phi_{sh})$, we have $\tilde{\mathbf{w}}_{sh}^i = (\tilde{\psi}_{sh}^i, \tilde{\phi}_{sh}^i)$. The augmented Lagrangian $\mathcal{L}(\mathbf{w}_{in}, \{\tilde{\mathbf{w}}_{sh}^i\}, \mathbf{w}_{sh}, \{\boldsymbol{\lambda}_i\})$ can be denoted equivalently as

$$\mathcal{L}\left(\left\{\left(\phi_{in}^i, \tilde{\phi}_{sh}^i\right)\right\}, \left\{\left(\psi_{in}^i, \tilde{\psi}_{sh}^i\right)\right\}, (\psi_{sh}, \phi_{sh}), \{\boldsymbol{\lambda}_i\}\right). \quad (5.3)$$

The classical ADMM framework works on this augmented Lagrangian function \mathcal{L} , and in each iteration it minimizes variable blocks $\left\{\left(\phi_{in}^i, \tilde{\phi}_{sh}^i\right)\right\}$, $\left\{\left(\psi_{in}^i, \tilde{\psi}_{sh}^i\right)\right\}$ and (ψ_{sh}, ϕ_{sh}) based on Gauss-Seidel scheme and further updates the Lagrangian multiplier $\{\boldsymbol{\lambda}_i\}$. Instead of minimizing the Lagrangian function directly, the gradient descent technique is employed to approximately updating primal variable blocks progressively, and as a result, the brief primal-dual hybrid gradient algorithm framework is proposed in the following calculation scheme (Boyd et al., 2011; Chambolle and Pock, 2011)

$$\begin{bmatrix} \phi_{in}^i \\ \tilde{\phi}_{sh}^i \end{bmatrix} \leftarrow \begin{bmatrix} \phi_{in}^i \\ \tilde{\phi}_{sh}^i \end{bmatrix} - \beta_{\phi} \begin{bmatrix} \nabla_{\phi_{in}^i} \mathcal{L}(\mathbf{w}_{in}, \{\tilde{\mathbf{w}}_{sh}^i\}, \mathbf{w}_{sh}, \{\boldsymbol{\lambda}_i\}) \\ \nabla_{\tilde{\phi}_{sh}^i} \mathcal{L}(\mathbf{w}_{in}, \{\tilde{\mathbf{w}}_{sh}^i\}, \mathbf{w}_{sh}, \{\boldsymbol{\lambda}_i\}) \end{bmatrix}, \quad i = 1, \dots, n; \quad (5.4a)$$

$$\begin{bmatrix} \psi_{in}^i \\ \tilde{\psi}_{sh}^i \end{bmatrix} \leftarrow \begin{bmatrix} \psi_{in}^i \\ \tilde{\psi}_{sh}^i \end{bmatrix} - \beta_{\psi} \begin{bmatrix} \nabla_{\psi_{in}^i} \mathcal{L}(\mathbf{w}_{in}, \{\tilde{\mathbf{w}}_{sh}^i\}, \mathbf{w}_{sh}, \{\boldsymbol{\lambda}_i\}) \\ \nabla_{\tilde{\psi}_{sh}^i} \mathcal{L}(\mathbf{w}_{in}, \{\tilde{\mathbf{w}}_{sh}^i\}, \mathbf{w}_{sh}, \{\boldsymbol{\lambda}_i\}) \end{bmatrix}, \quad i = 1, \dots, n; \quad (5.4b)$$

$$\begin{bmatrix} \psi_{sh} \\ \phi_{sh} \end{bmatrix} \leftarrow \begin{bmatrix} (1/n) \sum_{i=1}^n \tilde{\psi}_{sh}^i \\ (1/n) \sum_{i=1}^n \tilde{\phi}_{sh}^i \end{bmatrix}; \quad (5.4c)$$

$$\boldsymbol{\lambda}_i \leftarrow \boldsymbol{\lambda}_i - \beta (\mathbf{w}_{sh} - \tilde{\mathbf{w}}_{sh}^i), \quad i = 1, \dots, n. \quad (5.4d)$$

The update step (5.4a) calculates the critic parameters ϕ_{in}^i together with the splitting shared critic parameters $\tilde{\phi}_{sh}^i$, and (5.4b) updates the actor parameters ψ_{in}^i with the splitting shared actor parameters $\tilde{\psi}_{sh}^i$. (5.4c) aims to update the shared parameters (ϕ_{sh}, ψ_{sh}) by averaging all the splitting shared critic and actor respectively. It is evident that each agent i computes its actor and critic parameters and its splitting shared parameters. The overall shared parameters are updated through (5.4c) and broadcast to all agents. (5.4) can be decentralized implemented and full details can be found in Algorithm 1.

Noting that if no shared parameters is used in formulation (4.5) and (5.1), then the whole problem is degenerated to a simplified version, i.e.,

$$\min_{\mathbf{w}_{in}} \mathcal{J}(\mathbf{w}_{in}) = \alpha_1 \mathcal{J}_{\text{actor}}(\mathbf{w}_{in}) + \alpha_2 \mathcal{J}_{\text{critic}}(\mathbf{w}_{in}) + \mathcal{R}_{in}(\mathbf{w}_{in}). \quad (5.5)$$

Algorithm 1 The F2A2 Algorithm Framework.

- 1: Initialize independent parameters \mathbf{w}_{in} , shared parameters $\tilde{\mathbf{w}}_{sh}^i$. Set consensus parameters $\tilde{\mathbf{w}}_s$, dual parameters $\{\boldsymbol{\lambda}_i\}$ and the unified penalty parameter β to zero. For each agent i , initialize all other policies estimation parameters $\psi_{j,in}^i$ for $j = 1, \dots, i-1, i+1, \dots, n$.
- 2: **for** each episode **do**
- 3: **for** $t = 1$ to the pre-defined max length of the episode **do**
- 4: Each agent observes initial observation o_i^t ;
- 5: For agent i , select action a_i^t by current policy π_i ;
- 6: Execute $\mathbf{a}^t = (a_1^t, \dots, a_N^t)$ and get the cost c_i^{t+1} and next observation o_i^{t+1} ;
- 7: Store $(\mathbf{o}^t, \mathbf{a}^t, \mathbf{c}^{t+1}, \mathbf{o}^{t+1})$ to replay buffer \mathcal{D} ;
- 8: **for** each agent i **do**
- 9: Sample a batch tuple $\{(\mathbf{o}^k, \mathbf{a}^k, \mathbf{c}^{k+1}, \mathbf{o}^{k+1})\}$ from replay buffer \mathcal{D} ;
- 10: For each tuple sample the next action \mathbf{a}^{k+1} according to the estimated policies $\tilde{\boldsymbol{\pi}}$ of all other agents;
- 11: **Critic-step:** Update independent critic parameters and consensus critic parameters $\{\phi_{in}^i, \tilde{\phi}_{sh}^i\}$:

$$\begin{bmatrix} \phi_{in}^i \\ \tilde{\phi}_{sh}^i \end{bmatrix} \leftarrow \begin{bmatrix} \phi_{in}^i \\ \tilde{\phi}_{sh}^i \end{bmatrix} - \beta_\phi \begin{bmatrix} \nabla_{\phi_{in}^i} \mathcal{L}(\mathbf{w}_{in}, \{\tilde{\mathbf{w}}_{sh}^i\}, \mathbf{w}_{sh}, \{\boldsymbol{\lambda}_i\}) \\ \nabla_{\tilde{\phi}_{sh}^i} \mathcal{L}(\mathbf{w}_{in}, \{\tilde{\mathbf{w}}_{sh}^i\}, \mathbf{w}_{sh}, \{\boldsymbol{\lambda}_i\}) \end{bmatrix},$$

and use the formula in Prop. 4 or its variants to calculate the sampled off-policy policy gradient and α is the step-size of gradient descent method.

- 12: **Actor-step:** Update independent actor parameters and consensus actor parameters $\{\psi_{in}^i, \tilde{\psi}_{sh}^i\}$ by minimizing the loss:

$$\begin{bmatrix} \psi_{in}^i \\ \tilde{\psi}_{sh}^i \end{bmatrix} \leftarrow \begin{bmatrix} \psi_{in}^i \\ \tilde{\psi}_{sh}^i \end{bmatrix} - \beta_\psi \begin{bmatrix} \nabla_{\psi_{in}^i} \mathcal{L}(\mathbf{w}_{in}, \{\tilde{\mathbf{w}}_{sh}^i\}, \mathbf{w}_{sh}, \{\boldsymbol{\lambda}_i\}) \\ \nabla_{\tilde{\psi}_{sh}^i} \mathcal{L}(\mathbf{w}_{in}, \{\tilde{\mathbf{w}}_{sh}^i\}, \mathbf{w}_{sh}, \{\boldsymbol{\lambda}_i\}) \end{bmatrix},$$

where use the recent obtained $\{\phi_{in}^i, \tilde{\phi}_{sh}^i\}$ into critic-step; and use the formula in Prop. 4 or its variants to calculate the sampled off-policy policy gradient; and α is the step-size of gradient descent method.

- 13: **Consensus-step:** Eq. 5.4 can be implemented in a decentralized manner. Specifically, agent i broadcast the recent obtained consensus parameter $\tilde{\phi}_{sh}^i, \tilde{\psi}_{sh}^i$ and receive all others recent obtained consensus parameter $\tilde{\phi}_{sh}^j, \tilde{\psi}_{sh}^j$ for all $j \neq i$. Then update shared parameter \mathbf{w}_{sh} :

$$\begin{bmatrix} \psi_{sh} \\ \phi_{sh} \end{bmatrix} \leftarrow \begin{bmatrix} (1/n) \sum_{i=1}^n \tilde{\psi}_{sh}^i \\ (1/n) \sum_{i=1}^n \tilde{\phi}_{sh}^i \end{bmatrix},$$

- 14: and update the dual multiplier parameters $\{\boldsymbol{\lambda}_i\}$: $\boldsymbol{\lambda}_i \leftarrow \boldsymbol{\lambda}_i - \beta(\mathbf{w}_{sh} - \tilde{\mathbf{w}}_{sh}^i)$
 - 15: **if** $t \bmod$ the update frequency of the policy estimation model $== 0$ **then**
 - 16: Sample a batch tuple $\{(\mathbf{o}^k, \mathbf{a}^k, \mathbf{c}^{k+1}, \mathbf{o}^{k+1})\}$ from replay buffer \mathcal{D} ;
 - 17: Update the estimated policies $\tilde{\boldsymbol{\pi}}$ of all others by supervised learning.
 - 18: **end if**
 - 19: **end for**
 - 20: **end for**
 - 21: **end for**
-

The scheme (5.4) will be simplified into

$$\phi_{in}^i \leftarrow \phi_{in}^i - \nabla_{\phi_{in}^i} \mathcal{J}(\mathbf{w}_{in}), \quad i = 1, \dots, n; \quad (5.6a)$$

$$\psi_{in}^i \leftarrow \psi_{in}^i - \nabla_{\psi_{in}^i} \mathcal{J}(\mathbf{w}_{in}), \quad i = 1, \dots, n. \quad (5.6b)$$

(5.6) can be considered to minimize (5.5) concerning the critic parameters $\{\phi_{in}^i\}$ and actor parameters $\{\psi_{in}^i\}$ alternatively. This basic scheme is different from traditional multi-agent actor-critic algorithms (e.g., 3.4), while the BCGD-type scheme is employed on the jointly MARL framework (5.5).

Remark 2. Although some previous value-based multi-agent reinforcement learning algorithms (Lauer and Riedmiller, 2000; Matignon et al., 2007; Panait et al., 2008; Arslan and Yüksel, 2017) do not pass any messages between agents during the learning process and belong to pure decentralization, it is still reasonable to name the proposed framework as “fully decentralized” similar with Zhang et al. (2018c). These messages transmit between agents in the learning procedure are not uniformly collected and distributed by a centralized controller, but each agent sends and receives them individually.

Remark 3. There are two key formulations in this paper, i.e., Equation 4.1 and 4.2. As discussed above, problem (4.1) is the core problem that actor-critic type methods aim to solve, and problem (4.2) can be considered as an approximation version of problem (4.1). Problem (4.2) has the separable structure which motivated us to design the proposed Algorithm 1. The optimal solutions of problem (4.1) and (4.2) seem to be different, however the optimal solution set of problem (4.1) is more difficult to guarantee because of its bi-level programming structure. Very few works have discussed the theoretical analysis of the RL algorithm to solve the bi-level formulation, while nearly all of them focused on the single-agent case. For instance, Yang et al. (2018c), Yang et al. (2019) and Hong et al. (2020) consider the single-agent actor-critic algorithm as a specific solution to the corresponding bi-level problem, and give some convergence results on linear function approximation case. As for the multi-agent case, Zhang et al. (2018c) proposes two decentralized actor-critic algorithms with function approximation, and convergence analyses of the algorithms are provided when the value functions are approximated within the class of linear functions. Zhang et al. (2018c) still could not guarantee the convergence to the stationary point of the bi-level programming problem (4.1). For the stability of the training, a regularization term is added to the problem (4.2) to make the problem into the more general problem (4.5), which can be equivalently converted into problem (5.1). In this paper, some convergence results of the proposed Algorithm 1 can be further established. Expressly, Algorithm 1 can be incorporated into the algorithm framework of Hong et al. (2016) for solving the problem (4) (5.1 in the modified version). If Algorithm 1 can satisfy (Hong et al., 2016, Assumption A) then theoretical results similar to can (Hong et al., 2016, Theorem 2.4) be obtained, i.e., *Any limit point $\{\{\phi_{in}^{i,*}\}, \{\psi_{in}^{i,*}\}, \{\phi_{sh}^*\}, \{\psi_{sh}^*\}, \{\boldsymbol{\lambda}^{i,*}\}\}$ of the sequence $\{\{\phi_{in}^{i,k}\}, \{\psi_{in}^{i,k}\}, \{\phi_{sh}^k\}, \{\psi_{sh}^k\}, \{\boldsymbol{\lambda}^{i,k}\}\}$ which is obtained from Algorithm 1 is a stationary point of problem (5.1). Because problem (5.1) is equivalent with problem (4.5), we have that $\{\{\phi_{in}^{i,*}\}, \{\psi_{in}^{i,*}\}, \{\phi_{sh}^*\}, \{\psi_{sh}^*\}, \{\boldsymbol{\lambda}^{i,*}\}\}$ is a stationary point of problem (4.5).* The policy or value approximation functions are not limited to the linear case, and the local

convergence of the obtained sequence for general cases in Algorithm 1 can be proved. To emphasize, the convergence to the stationary point of the bi-level programming formulation problem (4.1) could not be guaranteed.

5.1 Instantiation Algorithms

In the proposed PDHG algorithm framework, the gradient of the augmented Lagrangian function \mathcal{L} concerning the primal variables (see Eq. 5.4a and 5.4b) need to be calculated to optimize the actors and critics of each agent. The proposed optimization objective functions are different from traditional centralized and decentralized MARL algorithms. Therefore, the related results in these algorithms cannot be directly used. This section will give the detailed form of the gradient of the Lagrangian function \mathcal{L} . In Section 4, it can be seen that the proposed algorithm framework has good flexibility, so the proposed framework could introduce various single-agent reinforcement learning algorithms as the backbone.

Specifically, one on-policy algorithm, i.e. COMA (Foerster et al., 2018a), and three state-of-the-art single-agent off-policy actor-critic algorithms, i.e. DDPG (Lillicrap et al., 2016), TD3 (Fujimoto et al., 2018) and SAC (Haarnoja et al., 2018) are incorporated into the proposed decentralized framework.

5.1.1 ON-POLICY F2A2 INSTANTIATION ALGORITHMS

For the on-policy methods, the current state-of-the-art algorithm COMA is chosen as the backbone algorithm. The corresponding proposition for the COMA algorithm is proposed to calculate the on-policy joint gradient.

Proposition 3 (On-Policy COMA-Based Joint Gradient) d_π represents the distribution of the state-occupancy measure of policy π , and δ is the TD(0)-error. The counterfactual baseline $\mathcal{B}(\mathbf{o}, \mathbf{a}^i)$ also be introduced from COMA. So the gradient of $J_{ac}^i(\mathbf{w}_{in}, \tilde{\mathbf{w}}_{sh}^i)$ is

$$\begin{aligned} \nabla_{\psi_{in}^i} J_{ac}^i(\mathbf{w}_{in}, \tilde{\mathbf{w}}_{sh}^i) &= \\ \mathbf{E}_{s \sim d_\pi, \mathbf{o} \sim \mathcal{E}, \mathbf{a} \sim \pi} &\left[\alpha_1 \nabla_{\psi_{in}^i} \log \pi_{\psi_{in}^i}^i(a^i | o^i) \left(Q_{\tilde{\phi}_{sh}^i}^{\pi, i}(\mathbf{o}, \mathbf{a}) - \mathcal{B}(\mathbf{o}, \mathbf{a}^i) + \frac{\alpha_2}{\alpha_1} \delta^2 \right) \right], \\ \nabla_{\tilde{\phi}_{sh}^i} J_{ac}^i(\mathbf{w}_{in}, \tilde{\mathbf{w}}_{sh}^i) &= \mathbf{E}_{s \sim d_\pi, \mathbf{o} \sim \mathcal{E}, \mathbf{a} \sim \pi} \left[(\alpha_1 + 2\alpha_2\delta) \nabla_{\tilde{\phi}_{sh}^i} Q_{\tilde{\phi}_{sh}^i}^{\pi, i}(\mathbf{o}, \mathbf{a}) \right]. \end{aligned}$$

For other single-agent on-policy actor-critic algorithms need to be incorporated, we just need to replace the J_{actor}^i and J_{critic}^i part in $J_{ac}^i(\mathbf{w}_{in}, \tilde{\mathbf{w}}_{sh}^i)$ with the corresponding form and derive their corresponding on-policy joint gradients. The specific forms of F2A2-COMA and the proof of Proposition 3 are given in supplemental material. With the gradient calculated above, below, the F2A2-COMA algorithm is formally proposed.

F2A2-COMA. COMA (Counterfactual Multi-Agent Policy Gradient) method learns a centralized critic with a *counterfactual baseline* which is inspired by *difference rewards* to solve the *multi-agent credit assignment* problem. The COMA algorithm is introduced into the F2A2 framework, and the F2A2-COMA algorithm is then proposed accordingly. F2A2-

COMA focuses on settings with discrete actions but can be easily extended to continuous action spaces by estimating counterfactual baseline with Monte Carlo samples or using functional forms that render it analytical, e.g., Gaussian policies and critic. In the F2A2-COMA algorithm, all agents share a centralized counterfactual baseline function, but the policy functions are independent of each other. All the other settings are the same as the COMA algorithm.

5.1.2 OFF-POLICY F2A2 INSTANTIATION ALGORITHMS

The following proposition for off-policy methods is firstly proposed since the above single-agent off-policy algorithms are either based on or related to the DDPG algorithm.

Proposition 4 (Off-Policy DDPG-Based Joint Gradient) π_0 represents the data collection policy sampled from experience replay buffer, d_0 represents the distribution of the state-occupancy measure of policy π_0 , and δ is the TD(0)-error. $\epsilon, \alpha_1, \alpha_2$ are hyperparameters. $\psi_i := \{\psi_{in}^i, \tilde{\psi}_{sh}^i\}$, $\phi_i := \{\phi_{in}^i, \tilde{\phi}_{sh}^i\}$, the gradients of $J_{ac}^i(\mathbf{w}_{in}, \tilde{\mathbf{w}}_{sh}^i)$ w.r.t. $\{\psi_{in}^i, \tilde{\psi}_{sh}^i\}$ and $\{\phi_{in}^i, \tilde{\phi}_{sh}^i\}$ are

$$\begin{aligned} & \nabla_{\psi_{in}^i} J_{ac}^i(\mathbf{w}_{in}, \tilde{\mathbf{w}}_{sh}^i) \\ &= \mathbf{E}_{s \sim d_0, o \sim \mathcal{E}, a \sim \pi} \left[(\alpha_1 + 2\alpha_2 \delta \left(\frac{\pi_0^i(a^i|o^i)}{\pi_{\psi_{in}^i}^i(a^i|o^i)} \right)) \nabla_{\psi_{in}^i} \pi_{\psi_{in}^i}^i(a^i|o^i) \nabla_{a^i} Q_{\tilde{\phi}_{sh}^i}^{\pi, i}(\mathbf{o}, \mathbf{a}) \right], \\ & \nabla_{\tilde{\phi}_{sh}^i} J_{ac}^i(\mathbf{w}_{in}, \tilde{\mathbf{w}}_{sh}^i) \\ &= \mathbf{E}_{s \sim d_0, o \sim \mathcal{E}, a \sim \pi_0} \left[\left(\alpha_1 \left(\frac{\pi_{\psi_{in}^i}^i(a^i|o^i)}{\pi_0^i(a^i|o^i)} \right) + 2\alpha_2 \delta \right) \nabla_{\tilde{\phi}_{sh}^i} Q_{\tilde{\phi}_{sh}^i}^{\pi, i}(\mathbf{o}, \mathbf{a}) \right] + \\ & \quad \mathbf{E}_{s \sim d_0, o \sim \mathcal{E}, a \sim \pi} \left[\left(\alpha_1 + 2\alpha_2 \delta \left(\frac{\pi_0^i(a^i|o^i)}{\pi_{\psi_{in}^i}^i(a^i|o^i)} \right) \right) \nabla_{\tilde{\phi}_{sh}^i} Q_{\tilde{\phi}_{sh}^i}^{\pi, i}(\mathbf{o}, \mathbf{a}) \right]. \end{aligned}$$

The J_{actor}^i and J_{critic}^i part in $J_{ac}^i(\mathbf{w}_{in}, \tilde{\mathbf{w}}_{sh}^i)$ can be replaced with the corresponding form of the single-agent off-policy actor-critic algorithm need to be incorporated, and their corresponding joint gradients can be derived accordingly. The specific forms of F2A2-DDPG, F2A2-TD3, and F2A2-SAC are shown in Figure 10 in the supplementary material. The proof of Proposition 4 and its variants (of F2A2-TD3 and F2A2-SAC) are given in supplemental material. With the gradient calculated above, below, three off-policy F2A2 instantiation algorithms are proposed formally.

F2A2-DDPG. The DDPG algorithm is introduced into the F2A2 framework, and the F2A2-DDPG algorithm is proposed accordingly. In F2A2-DDPG, all agents share a centralized value function, but the policy functions are independent. At the same time, considering that the simulated environments in this paper are all designed with discrete action space, DDPG cannot directly deal with the above situation. So this paper learns from the ideas of Lowe et al. (2017), rather than using policies that deterministically output an action,

policies that produce differentiable samples through a Gumbel-Softmax distribution (Jang et al., 2017) is then used.

F2A2-TD3. While DDPG can achieve excellent performance sometimes, it is frequently brittle concerning hyperparameters and other kinds of tuning. For example, a standard failure mode for DDPG is that the learned Q-function begins to dramatically overestimate Q-values, leading to policy breaking because it exploits the errors in the Q-function. TD3 (Twin Delayed DDPG) is an algorithm that addresses this issue by introducing three critical tricks, clipped double Q-learning, delayed policy updates, and target policy smoothing. In the F2A2-TD3 algorithm, all agent’s two centralized Q-value functions share each other, and the policy function is independent. Each agent has a different policy update frequency and policy smoothing noise for better exploration. Since TD3 cannot be applied to discrete action space, the same approach as F2A2-DDPG is adopted to modify the TD3 algorithm.

F2A2-SAC. SAC (Soft Actor Critic) is an algorithm that optimizes a stochastic policy in an off-policy way, forming a bridge between stochastic policy optimization and DDPG-style approaches. It is not a direct successor to TD3 (having been published roughly concurrently). However, it incorporates the clipped double-Q trick, and due to the inherent stochasticity of the policy in SAC, it also winds up benefiting from something like target policy smoothing. Same as F2A2-TD3, all agents have the same centralized Q-value function and independent policy function.

Remark 4. In the practical implementation, we use the truncated importance sampling ratio inspired by Munos et al. (2016) for all off-policy instantiation algorithms (F2A2-DDPG, F2A2-TD3, F2A2-SAC) to stabilize the training process, i.e.,

$$\left(\frac{\pi_{\psi_{in}^i}^i(a^i|o^i)}{\pi_0^i(a^i|o^i)} \right) \rightarrow \min \left(\epsilon, \frac{\pi_{\psi_{in}^i}^i(a^i|o^i)}{\pi_0^i(a^i|o^i)} \right), \text{ and } \left(\frac{\pi_0^i(a^i|o^i)}{\pi_{\psi_{in}^i}^i(a^i|o^i)} \right) \rightarrow \min \left(\epsilon, \frac{\pi_0^i(a^i|o^i)}{\pi_{\psi_{in}^i}^i(a^i|o^i)} \right),$$

and we set $\epsilon = 1$ in all experiments. Compared to original importance sampling ratio, it does not suffer from the variance explosion of the product of importance sampling ratios. Truncated importance sampling ratio has other theoretical advantages and interested readers can refer to the original paper (Munos et al., 2016).

5.2 Modeling Other Agents

Existing decentralized reinforcement learning works (Zhang et al., 2018c; Doan et al., 2019b; Suttle et al., 2019) assume that each agent can observe others’ actions. While such an assumption is too strict, and the more realistic assumption is that the agent has to model other agents’ policies and predict other agents’ actions based on their historical observations. Some further improvements are involved into the proposed algorithm framework to make the proposed framework still works better under such a more general assumption. Specifically, in this more general assumption the symbol π , in Proposition 4, 3 and their variants, are changed to $\tilde{\pi}$ and the symbol \mathbf{a} is changed to $\tilde{\mathbf{a}} := \tilde{\pi}(\mathbf{o})$. Another important issue is that the approximate policy $\tilde{\pi}$ and action $\tilde{\mathbf{a}}$ may lead to distorted gradients in Algorithm 1. However, gradients with noise may help optimization algorithms converge to global

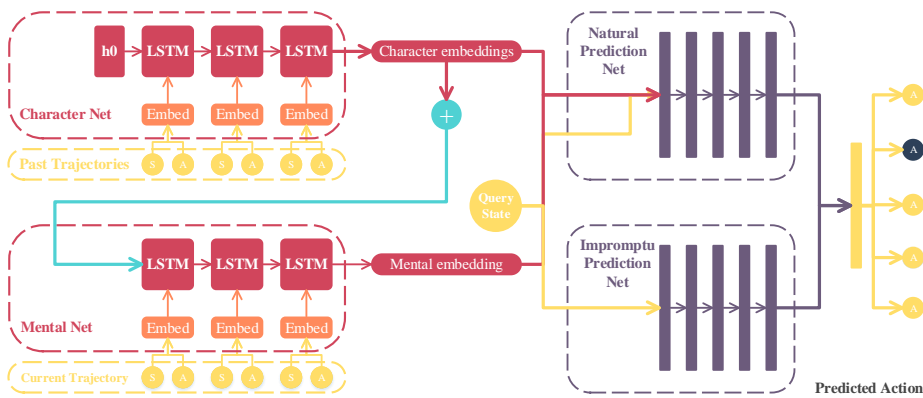


Figure 3: Modeling other agents’ policies. Character and Mental networks are used to encode the historical and instant information of the agent. Natural Prediction and Impromptu Prediction networks obtain the final prediction via online supervised learning with previous outputs and the current state.

optimal solution (Jastrzebski et al., 2017; Neelakantan et al., 2017; Smith and Topin, 2019). Although the estimated policy might bring bias, it could reduce the algorithm’s variance, which further improves the robustness of the system. This will be demonstrated later in the experiments.

To better estimate the policies of other agents in complex multi-agent environments, a novel modeling other agents (MOA) approach is devised based on theory-of-mind (TOM) inspired by Rabinowitz et al. (2018), and similar approach have also been found in other works (Jaques et al., 2019). The method is further connected with online supervised learning. The algorithm architecture is shown in Figure 3. Specifically, each agent first randomly samples a fixed number of trajectories in the environment before training and encode each trajectory using an LSTM. Then, the average of all the trajectories encoding is taken as the *character* of the agent. After training starts, for each current unfinished trajectory, another LSTM is used to encode its historical segment as the current *mental* of the agent. On the one hand, the current state, the character, and the current mental of the agent are used together as input to the *natural prediction network*; on the other hand, the *impromptu prediction network* is proposed to predict the action of agents only relying on the current state. The outputs of the two networks are combined to produce the final predictions. Note that the entire network’s training process is performed together with the reinforcement learning algorithm, and the training data is periodically collected online.

Finally, as detailed in Figure 4, we propose the fully decentralized algorithm framework, which follows the primal-dual hybrid gradient scheme and simultaneously splits joint tasks with the divide-and-conquer strategy. Algorithm 1 is the corresponding pseudocode.

5.3 Information Setting Comparison

Finally, to clearly show the differences among the proposed F2A2, the MARL algorithms in the CTDE framework, and existing decentralized MARL algorithms, the information settings of the above algorithms is summarized and the details are shown in Table 2.

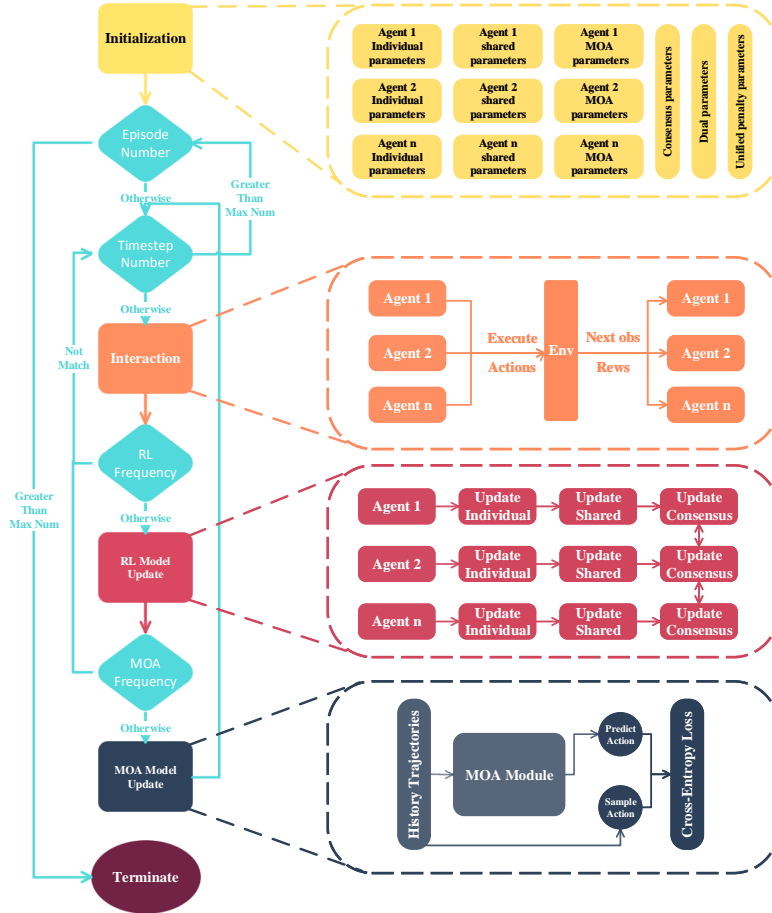


Figure 4: Flowchart corresponding to Algorithm 1.

Schemes	Algorithms	Joint Observation	Joint Policy	Joint Cost	Shared Actor	Shared Critic
Centralized Algorithms	MADDPG (Lowe et al., 2017)	Accessible	Accessible	-	-	Accessible
	COMA (Foerster et al., 2018a)	Accessible	Accessible	Accessible	-	Accessible
	MAAC (Iqbal and Sha, 2019)	Accessible	Accessible	-	-	Accessible
Decentralized Algorithms	MA-AC (Zhang and Zavlanos, 2019)	Accessible	Communication	-	Communication	-
	MA-AC (Zhang et al., 2018c)	Accessible	Communication	-	-	Communication
	F2A2	Accessible	Estimation	-	Communication	Communication

Table 2: The information settings of typical centralized, decentralized algorithms, and F2A2. Dashes indicate that certain elements do not exist.

For centralized MARL algorithms (i.e., MADDPG, COMA, MAAC), each agent must use joint observations (and the joint actions taken by all agents under the current joint policy) as input when calculates the joint Q-value (and trains the joint Q-value function), due to the existence of the centralized critic. Therefore, the joint observations and joint policy must be accessible for each agent during the algorithm training process. For MADDPG and MAAC, because they solve the general cooperative POSG problem, each agent only needs to optimize its own expected cumulative costs. Therefore, each agent does not need

Environment	Scenario	Scale	Agent #
Cooperative MPE (Iqbal and Sha, 2019)	Cooperative Treasure Collection	Small-Scale	8
	Rover Tower	Small-Scale	8
StarCraft II Micromanagement (Samvelyan et al., 2019)	Map 2s3z	Small-Scale	5
	Map 3m	Small-Scale	3
	Map 8m	Small-Scale	8
MAgent (Zheng et al., 2018)	Battle	Large-Scale	256

Table 3: Attributes of the experiment environments.

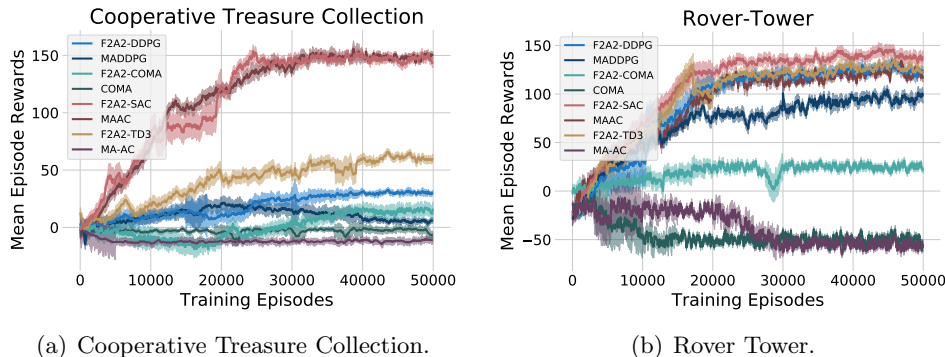
to access the joint cost. However, for the COMA algorithm, it is necessary to access the joint cost because it solves the fully cooperative MARL problem.

Existing decentralized algorithms achieve global access to some centralized modules (e.g., joint policy, centralized critic, centralized actor) by communicating with each other. However, for the partially observed multi-agent environment, to make the agent’s policy have a more robust representation ability, its policy function is often modeled by a recurrent neural network. Transmission of such a large amount of parameters through communication will make the algorithm less scalable. Therefore, F2A2 estimates the policies of the other agents by introducing the MOA module.

6. Experiments

The numerical experiments of the proposed framework are conducted from the following three aspects: effectiveness, performance, and scalability. Specifically, the effectiveness of F2A2 is verified in two general cooperative POSG environments designed in Iqbal and Sha (2019) first. Then, F2A2 is combined with the recurrent neural network to verify the performance on the more challenging cooperative Starcraft II unit micromanagement tasks (Samvelyan et al., 2019). Finally, the scalability is verified in a large-scale general cooperative POSG environment MAgent(Zheng et al., 2018). Table 3 contains the concise introduction for each environments.

In the above environments, the corresponding centralized baseline method MADDPG, MAAC, and COMA is set for the F2A2-framework instantiation F2A2-DDPG, F2A2-SAC, and F2A2-COMA. Like MAAC, the counterfactual baseline proposed by COMA is also introduced in F2A2-SAC. For fairness, the attentional critic is used as same as the MAAC algorithm in F2A2-SAC. In addition, as far as we know, there is no published centralized MARL algorithm using TD3 as the backbone. Therefore, the performance of the F2A2-TD3 is used as an indicator to measure the adaptability of the F2A2 framework to the single-agent algorithm instead of setting a corresponding baseline for the F2A2-TD3 algorithm. Specifically, because the TD3 algorithm is generally better than DDPG when the F2A2 framework adopts these single agent algorithms as the backbone, it can keep the relative order of performance after instantiation (that is, F2A2-TD3 should be better than F2A2-DDPG). For decentralized works, MA-AC (Zhang et al., 2018c) is chosen as the baseline.



(a) Cooperative Treasure Collection.

(b) Rover Tower.

Figure 5: Cooperative Multi-agent Particle Environment. The solid line represents the average under 10 seeds, and the shaded part represents the 95% confidence interval.

6.1 Cooperative Multi-agent Particle Environments

The *Multi-agent Particle Environments* (MPE) was first used in Lowe et al. (2017). However, the MPE includes cooperative tasks, competitive tasks, and cooperative-competitive-hybrid tasks, and the number of agents is small. To this end, the larger-scale collaborative environment based on the MPE environment proposed by Iqbal and Sha (2019), which is denoted as *Cooperative Multi-agent Particle Environments (Cooperative MPE)*, is used to more effectively verify the effectiveness of the algorithm. The two cooperative environments are introduced in **Appendix** in detail.

Figure 5(a) and 5(b) plot F2A2 instantiations’ performance relative to their centralized baselines and the decentralized baseline against the training episodes in Cooperative MPE. It can be seen that the algorithms instantiated by F2A2 can reach or exceed the performance of their corresponding centralized baselines, which shows that the F2A2 framework can effectively make global decisions based on local observations and also reflects the effectiveness of PDHG and MOA. At the same time, the performance of the F2A2-TD3 algorithm can exceed the F2A2-DDPG algorithm, which shows that the F2A2 framework can flexibly integrate the single-agent algorithms and guarantee their original advantages.

Specifically, from Figure 5(a) and 5(b), it can be seen that in the *Rover-Tower* environment, the performance gaps between the F2A2 instantiated algorithms and corresponding centralized baselines are greater than the *Cooperative Treasure Collection* environment. We believe this is due to the characteristics of the environments. Compared with the *Cooperative Treasure Collection* environment, the interaction between agents in the *Rover-Tower* environment is more sparse. Therefore, too much consideration of the behavior of other agents can easily cause the algorithm to overfit other agents and converge to a poor local optimal. The MOA module in F2A2 can play a role in regularization, allowing the algorithm to ignore other agents’ effects partially.

We mentioned earlier that MA-AC could be seen as a decentralized version of the centralized algorithm COMA. From the performance comparison of the two in the figure, it can be seen that if joint training and MOA are lacking, the performance of the decentralized algorithm cannot exceed the performance of the centralized algorithm. It should be noted here that, in theory, if the centralized algorithm can find the optimal global solution, then

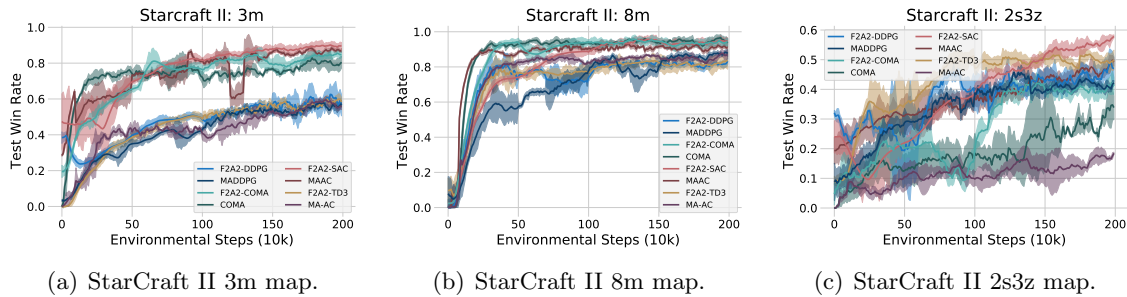


Figure 6: StarCraft II Environment. The solid line represents the average under 10 random seeds, and the shaded part represents the corresponding 95% confidence interval.


the decentralized algorithm cannot exceed the performance of the centralized algorithm. However, the current centralized MARL algorithms can only achieve local optimal.

6.2 StarCraft II Micromanagement

This section focuses on the decentralized micromanagement problem in StarCraft II. The combat scenarios where two groups of identical units are placed symmetrically on the map are considered. The units of the first group, allied, are controlled by the proposed algorithms. The enemy units are controlled by a built-in StarCraft II AI, which makes use of handcrafted heuristics. The initial placement of units within the groups varies across episodes. The difficulty of the computer AI controlling the enemy units is set to *medium*. The results on a set of maps where each unit group consists of 3 Marines (3m), 8 Marines (8m), and 2 Stalkers and 3 Zealots (2s3z) are compared with baselines.

Similar to the work of Foerster et al. (2018a) and Rashid et al. (2018), the action space of agents consists of the following set of discrete actions: *move[direction]*, *attack[enemy id]*, *stop*, and *noop*. Agents can only move in four directions: north, south, east, or west. A unit is allowed to perform the *attack[enemy id]* action only if the enemy is within its shooting range. This facilitates the decentralization of the problem. The introduction of the unit sight range achieves partial observability. Moreover, agents can only observe others if they are alive and cannot distinguish between units that are dead or out of range. At each time step, the agents receive a joint cost equal to the total negative damage dealt on the enemy units. In addition, agents receive a bonus of 10 points after killing each opponent and 200 points after killing all opponents. These costs are all normalized to ensure the maximum cumulative cost achievable in an episode is -20 .

Figure 6(a), 6(b) and 6(c) plot F2A2 instantiations’ win rate in test environment relative to their centralized baselines against the environmental steps in three StarCraft II micromanagement maps. It can be seen from the figure that even if the Starcraft II environment is more complicated than Cooperative MPE, the instantiation algorithms of the F2A2 framework can still reach or exceed its corresponding centralized baseline in performance. This demonstrates the robustness of the F2A2 framework. In addition, the performance comparison between MA-AC and COMA shows similar results to the Cooperative MPE.



(a) Negative costs of different numbers of agents over training epochs. (b) Convergence times as the number of agents changes.




Figure 7: MAgent training. (a-b) the negative cost (reward) and convergence performance of F2A2 for 12 – 256 agents, the solid line represents the average under 10 random seeds. (c-e), F2A2 is first trained in the environment with 256 agents and then transferred to a larger environment with 625 agents directly. Here some interesting patterns that the agents show after transfer are shown.

6.3 Large-scale MAgent Environment

The *Battle* scenario in the MAgent platform provided by Zheng et al. (2018) is chosen as the simulation environment. The Battle game is a general cooperative POSG scenario with two armies fighting against each other in a grid world, each empowered by a different RL algorithm. Each army consists of 12 – 256 homogeneous agents. The goal of each army is to get fewer costs by collaborating with teammates to destroy all the opponents. The agent can take action to either move to or attack nearby grids. Ideally, the agent should learn skills such as chasing to hunt after training. The default cost setting is adopted: 0.005 for every move, -0.2 for attacking an enemy, -5 for killing an enemy, 0.1 for attacking an empty grid, and 0.1 for being attacked or killed. Since the proposed F2A2-DDPG/TD3/SAC/COMA are similar to MADDPG, they can only be run on a small scale due to global critic. Hence F2A2 is extended to a large-scale off-policy fully-decentralized actor-critic method, i.e. F2A2-ISAC, which combines F2A2-SAC with IQL (Tan, 1993). It differs from F2A2-SAC only in the input accepted by the critic. The former requires observations and (estimated) policies of all agents, while F2A2-ISAC only requires individual local observation and individual policy. Therefore, the MOA module is no longer required in F2A2-ISAC. Although the MAgent environment scale is enormous, each agent’s task is

etc	MAAC	F2A2-SAC					
	MAAC	F2A2-SAC-0	F2A2-SAC-20	F2A2-SAC-40	F2A2-SAC-60	F2A2-SAC-80	F2A2-SAC-100
	148(±8)	23(±4)	37(±4)	45(±5)	112(±4)	142(±8)	147(±9)
rt	MAAC	F2A2-SAC					
	MAAC	F2A2-SAC-0	F2A2-SAC-20	F2A2-SAC-40	F2A2-SAC-60	F2A2-SAC-80	F2A2-SAC-100
	121(±5)	52(±7)	52(±7)	78(±5)	126(±6)	135(±5)	139(±6)
3m	MAAC	F2A2-SAC					
	MAAC	F2A2-SAC-0	F2A2-SAC-20	F2A2-SAC-40	F2A2-SAC-60	F2A2-SAC-80	F2A2-SAC-100
	0.88(±0.11)	0.65(±0.08)	0.68(±0.07)	0.70(±0.08)	0.76(±0.09)	0.91(±0.10)	0.92(±0.09)
8m	COMA	F2A2-COMA					
	COMA	F2A2-COMA-0	F2A2-COMA-20	F2A2-COMA-40	F2A2-COMA-60	F2A2-COMA-80	F2A2-COMA-100
	0.97(±0.01)	0.83(±0.01)	0.82(±0.01)	0.85(±0.02)	0.84(±0.01)	0.94(±0.01)	0.96(±0.02)
2s3z	MAAC	F2A2-SAC					
	MAAC	F2A2-SAC-0	F2A2-SAC-20	F2A2-SAC-40	F2A2-SAC-60	F2A2-SAC-80	F2A2-SAC-100
	0.48(±0.02)	0.35(±0.02)	0.35(±0.03)	0.41(±0.02)	0.43(±0.02)	0.55(±0.01)	0.57(±0.02)

Table 4: The performance comparison of each algorithm under different communication topologies. The numbers following the F2A2 algorithm represent the proportions of edges added in the ring topology, and 0 means that no edges exist in the communication topology. “etc” and “rt” represent the *Cooperative Treasure Collection* task and the *Rover Tower* task in the *Cooperative MPE* environment, respectively. The data in the *Cooperative MPE* environment represents the mean episode reward of all agents in the last $1k$ timesteps; And “3m”, “8m”, “2s3z” represent different maps in *SC II* environment. The data in the *SC II* environment represents the mean test win rate of all agents in the last $10k$ timesteps. Numbers in parentheses indicate 95% confidence intervals, obtained under 5 random seeds and blacked numbers indicate best results.

relatively simple so that F2A2-ISAC can achieve good results. The parameters in actor and critic for all agents are shared, similar to (Gupta et al., 2017).

In Figure 7(a), as the number of agents increases, the cost curves have a similar trend, which indicates that the decentralized learning ability is not affected by the large-scale setting. In Figure 7(b), the convergence time does not increase significantly as the number of agents grows. The curve shows that the proposed algorithms have good scalability. Algorithms also learn some interesting patterns. In figure 7(c), in the early stage of confrontation, the rear agents explore the environment because they cannot directly engage the enemy; in 7(d), when a small maniple of enemies escape, agents are splitted to chase; in 7(e), the large forces use the means of encirclement to conquest.

6.4 More Analysis

F2A2 reduce the information transmission during the decentralized algorithm learning. We only consider the information transmission between the agents of the decentralized algorithm here. For the centralized algorithm, all the information transmission is carried out on one physical machine, so it is meaningless to count the amount of information transmission. We count the appropriate parameter amount and information transmission amount of the decentralized algorithm involved in the experiment in the two simulation environments. After statistics, it can be seen that F2A2 dramatically reduces the transfer

	MAAC	F2A2-SAC					
ctc	MAAC	F2A2-SAC- $\mathcal{N}(\infty, 0)$	F2A2-SAC- $\mathcal{N}(2000, 1000)$	F2A2-SAC- $\mathcal{N}(200, 100)$	F2A2-SAC- $\mathcal{N}(20, 10)$	F2A2-SAC- $\mathcal{N}(8, 4)$	F2A2-SAC- $\mathcal{N}(1, 0)$
	148(±8)	23(±4)	24(±4)	33(±7)	139(±7)	146(±9)	147(±9)
rt	MAAC	F2A2-SAC					
	MAAC	F2A2-SAC- $\mathcal{N}(\infty, 0)$	F2A2-SAC- $\mathcal{N}(2000, 1000)$	F2A2-SAC- $\mathcal{N}(200, 100)$	F2A2-SAC- $\mathcal{N}(20, 10)$	F2A2-SAC- $\mathcal{N}(8, 4)$	F2A2-SAC- $\mathcal{N}(1, 0)$
	121(±5)	52(±7)	52(±8)	56(±8)	136(±6)	139(±5)	139(±6)
3m	MAAC	F2A2-SAC					
	MAAC	F2A2-SAC- $\mathcal{N}(\infty, 0)$	F2A2-SAC- $\mathcal{N}(2000, 1000)$	F2A2-SAC- $\mathcal{N}(200, 100)$	F2A2-SAC- $\mathcal{N}(20, 10)$	F2A2-SAC- $\mathcal{N}(8, 4)$	F2A2-SAC- $\mathcal{N}(1, 0)$
	0.88(±0.11)	0.65(±0.08)	0.65(±0.07)	0.74(±0.06)	0.91(±0.09)	0.91(±0.10)	0.92(±0.09)
8m	COMA	F2A2-COMA					
	COMA	F2A2-COMA- $\mathcal{N}(\infty, 0)$	F2A2-COMA- $\mathcal{N}(2000, 1000)$	F2A2-COMA- $\mathcal{N}(200, 100)$	F2A2-COMA- $\mathcal{N}(20, 10)$	F2A2-COMA- $\mathcal{N}(8, 4)$	F2A2-COMA- $\mathcal{N}(1, 0)$
	0.97(±0.01)	0.83(±0.01)	0.83(±0.02)	0.83(±0.02)	0.93(±0.02)	0.95(±0.01)	0.96(±0.02)
2s3z	MAAC	F2A2-SAC					
	MAAC	F2A2-SAC- $\mathcal{N}(\infty, 0)$	F2A2-SAC- $\mathcal{N}(2000, 1000)$	F2A2-SAC- $\mathcal{N}(200, 100)$	F2A2-SAC- $\mathcal{N}(20, 10)$	F2A2-SAC- $\mathcal{N}(8, 4)$	F2A2-SAC- $\mathcal{N}(1, 0)$
	0.48(±0.02)	0.35(±0.02)	0.37(±0.01)	0.42(±0.01)	0.53(±0.01)	0.57(±0.02)	0.57(±0.02)

Table 5: The performance comparison of each algorithm under different parameter exchange frequencies. The normal distribution followed by the F2A2 algorithm indicates that each agent samples the time interval for executing the next *consensus step* from the normal distribution after each execution of the current *consensus step*, where ∞ indicates the predefined maximum executable update step of the algorithm. “ctc” and “rt” represent the *Cooperative Treasure Collection* task and the *Rover Tower* task in the *Cooperative MPE* environment, respectively. The data in the *Cooperative MPE* environment represents the mean episode reward of all agents in the last $1k$ timesteps; And “3m”, “8m”, “2s3z” represent different maps in *SC II* environment. The data in the *SC II* environment represents the mean test win rate of all agents in the last $10k$ timesteps. Numbers in parentheses indicate 95% confidence intervals, obtained under 5 random seeds and blacked numbers indicate best results.

of parameters in addition to shared parameters caused by policy sharing and only requires one-tenth of the transmissions.

The PDHG-type methods have advantages over the BCGD-type methods. To verify the superiority of the PDHG-type method, we select the corresponding optimal algorithms in two multi-agent environments and conduct comparative experiments. It can be seen from Figure 8(a) and Figure 8(b) that the PDHG-type jointly optimization methods (*-with-JOINT) have a significant performance improvement compared to the BCGD-type separately optimization methods (*-w/o-JOINT).

The MOA module further improves the performance of the decentralized MARL algorithm. In the previous section, we argue that the MOA module can prevent the agent from overfitting other agents’ policies by introducing noise, and increase exploration, thereby ultimately improving algorithm performance. Moreover, in some scenarios, it can exceed the centralized algorithms. To quantify the effect of the MOA module, the comparative experiments are also conducted. The experimental results show in Figure 8(b) and Figure 8(a) are clearly support the above conclusions.

Finally, to delve into why the F2A2 algorithm can outperform the centralized algorithm in some cases, we conduct a fine-grained comparison between centralized and decentralized methods for various levels of communication budgets. Specifically, the communication budget of the algorithm is constrained in the following two ways:

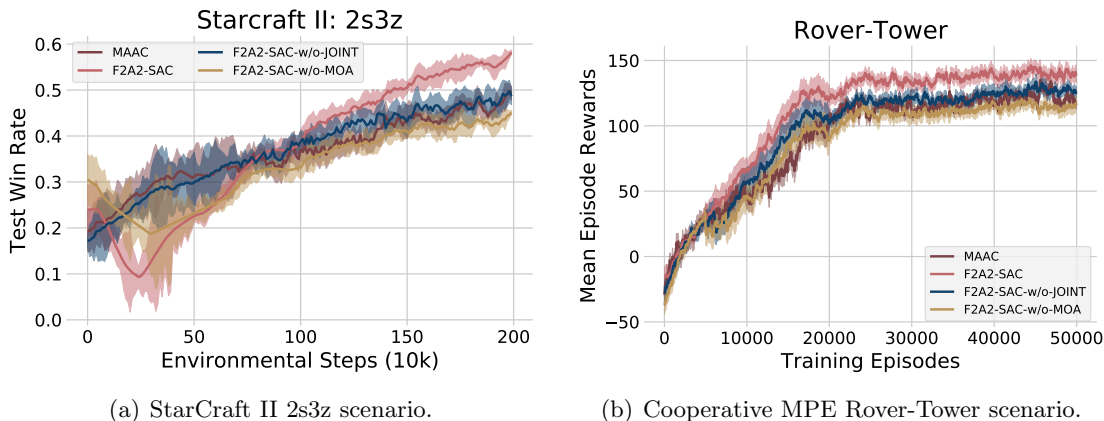


Figure 8: The ablation study of F2A2-SAC algorithm in two multi-agent environments.

- Communication topology.** The above experiments assume that the communication topology is a fully connected graph. In order to add the communication budget, we construct the communication topology as follows. For experiments under each random seed, we first sort all agents in random order to construct a ring topology (guarantee that there is exactly one connected component in the whole graph). Then randomly add edges between agents in different predefined proportions at each update step, and the upper limit of the number of edges is the total number of edges in the fully connected graph. Furthermore, we also consider an extreme case where there are no edges in the graph.
- Parameter exchange frequency.** At each *consensus step*, each agent samples the time interval to execute the next *consensus step* from a normal distribution. The mean of this normal distribution is set to different values from small to large to represent the change in communication budget from large to small. The variance of the normal distribution is fixed, and for each communication budget, different agents sample from the same normal distribution. We also consider an extreme case where the mean of the normal distribution is greater than the total number of update steps in the algorithm, and the variance is 0.

For each task, we select the best performing centralized algorithm and its corresponding decentralized version for comparison. The experimental results are shown in Table 4 and Table 5. It can be seen from the table that with the increase of the communication budget, the performance of the F2A2 algorithms present a step-wise growth rather than a smooth and gradual improvement. Moreover, it can be seen from the last few columns of the table that the F2A2 algorithm can still maintain good performance under the weak communication budget constraint.

7. Conclusion and Limitations

A flexible fully decentralized approximate actor-critic MARL framework is devised to achieve applicability and scalability in interactive multi-agent environments. A primal-dual optimization and joint actor-critic learning are carefully designated to guarantee full decen-

tralization and scalability, with agents modeling to increase the robustness. The proposed approach can even exceed state-of-the-art centralized algorithms in various categories and various scales simulated cooperative environments. In the future, we plan to introduce communication in the training process to promote more efficient cooperation to be more adaptable to complex scenarios.

Below we briefly analyze the limitations of F2A2. We ensured the Markov property of the policy through the enrichment of the observation space, avoiding the use of non-Markovian and history-dependent policy classes. Although this approach facilitates theoretical analysis and engineering implementation, the resulting *vast belief space* and *infinite hierarchy of belief* render the Markov policy class intractable in more complex problems, further constraining its generality. We shed light on the limitations that affect this policy class concerning crucial factors such as the error of the approximated information state and the infinite hierarchy belief representation capabilities in Appendix I. Careful consideration of these factors is essential to enhance the tractability and generality of the Markovian policy within the enriched observation space.

Acknowledgments

This work was supported in part by the National Key Research and Development Program of China (No. 2020AAA0107400), STCSM (22QB1402100), NSFC (No. 12071145), Shenzhen Science and Technology Program (JCYJ20210324120011032), Postdoctoral Science Foundation of China (2022M723039), a grant from China Academy of LVT, and a grant from Shenzhen Institute of Artificial Intelligence and Robotics for Society.

Appendix A. Environments

Cooperative Treasure Collection. The cooperative environment in Figure 9(a) involves 8 total agents, 6 of which are "treasure hunters" and 2 of which are "treasure banks", which each correspond to a different color of treasure. The role of the hunters is to collect the treasure of any color, which re-spawn randomly upon being collected (with a total of 6), and then "deposit" the treasure into the correctly colored "bank". The role of each bank is to gather as much treasure as possible from the hunters simply. All agents can see each others' positions concerning their own. Hunters receive a global cost for the successful collection of treasure, and all agents receive a global cost for the depositing of treasure. Hunters are additionally penalized for colliding with each other. As such, the task contains a mixture of shared and individual costs.

Rover Tower. The environment in Figure 9(b) involves 8 total agents, 4 of which are "rovers" and another 4 which are "towers". In each episode, rovers and towers are randomly paired. The pair is punished by the distance of the rover to its goal. The task can be thought of as a navigation task on an alien planet with limited infrastructure and low visibility. The rovers are unable to see in their surroundings. They must rely on communication from the towers, which can locate the rovers and their destinations and send one of five discrete communication messages to their paired rover. Note that communication is highly restricted and different from centralized policy approaches (Jiang and Lu, 2018), which allow for the free transfer of continuous information among policies. In our setup, the communication is integrated into the environment (in the tower's action space and the rover's observation space), rather than being explicitly part of the model. It is limited to a few discrete signals.

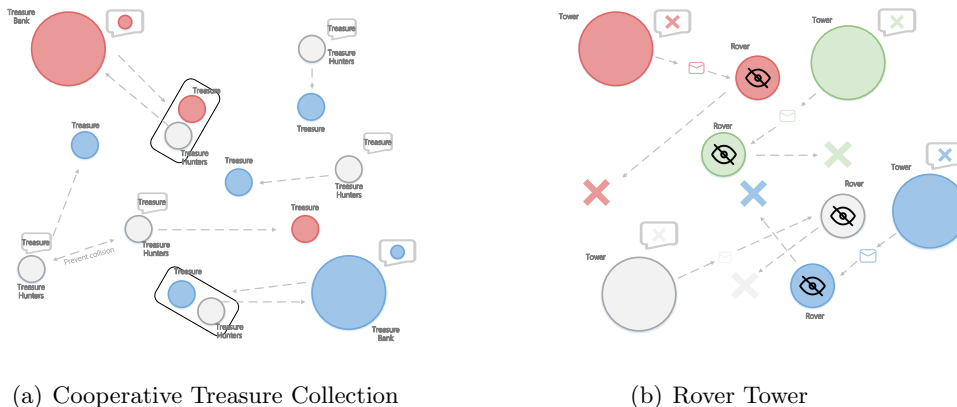


Figure 9: The Cooperative Multi-agent Particle Environments.

StarCraft II Map 2s3z. This map contains mixed unit types, where both the learnable agent and the built-in AI each control two Stalkers and three Zealots. Stalkers are ranged-attack units that take heavy damage from melee-type Zealots. Consequently, a winning strategy needs to dynamically coordinate between letting one's Zealots attack enemy Stalkers and backtrack to defend one's Stalkers against enemy Zealots.

StarCraft II Map 3m and 8m. The first, map 3m, presents both sides with three Marines, which are medium-ranged infantry units. The coordination challenge on this map is to reduce enemy firepower as quickly as possible by focusing unit fire to defeat each enemy unit in turn. Secondly, map 8m scales this task up to eight Marines on both sides. The relatively large number of agents involved poses additional scalability challenges.

Appendix B. Relationship with Trust-region Methods

The proposed F2A2 algorithm contains two learning processes: the learning of multi-agent policies and the learning of MOA modules. Moreover, the performance of the latter will affect the former. Once the MOA modules are too inaccurate, the entire F2A2 algorithm training process will fall into a vicious circle. In order to make the training of the entire system more robust, the intuitive idea is to make the multi-agent policies update more conservative, which can effectively improve the accuracy of MOA modules. Trust-region methods, such as Proximal Policy Optimization (PPO) (Schulman et al., 2017) have been demonstrated as efficient and robust RL algorithms via maximally searching the new policy in a trust region. Coincidentally, in comparing the PPO with the F2A2, we found that the proposed F2A2 framework is closely related to trust-region methods. This means that F2A2 naturally has the conservativeness of policy updating by jointly optimizing actors and critics. More specifically, the simplified objective of PPO is as follows,

$$L(s, a, \theta_k, \theta) = \min \left(\frac{\pi_\theta(a|s)}{\pi_{\theta_k}(a|s)} A^{\pi_{\theta_k}(s, a)}, g(\epsilon, A^{\pi_{\theta_k}(s, a)}) \right),$$

where

$$g(\epsilon, A) = \begin{cases} (1 + \epsilon)A, & A \geq 0; \\ (1 - \epsilon)A, & A < 0. \end{cases}$$

PPO deals with the advantage function $A^\pi(s, a)$, while our F2A2 deals with the gradient of value $Q^\pi(o, a)$. Note that PPO and F2A2 have a clipping regularizer to prevent the new policy beyond the trust region. The hyperparameter ϵ corresponds to the distance that the new policy can go away and still profits the objective (Schulman et al., 2017). Besides, the coefficient $\alpha_1 - 2\alpha_2\delta$ in the gradient of F2A2 makes the actor update more conservative when the current critic is not accurate enough (and the accuracy of the current critic is greatly affected by the accuracy of the MOA modules). To sum up, our F2A2 framework has a similar protective trust-region mechanism, which enhances the robustness of our algorithm.

Appendix C. Policy Gradient Theorem in POMDPs

After enriching the observation space, we can naturally extend the policy gradient theorem under MDP to POMDP based on Mao et al. (2020) as the basis for the subsequent theoretical derivation of this paper. In the previous version of manuscript, we omitted this derivation process and make an addition here. Notice that here o refer to the enriched observation unless specified. Let $p(\tau; \theta)$ denote the probability distribution of trajectories τ under policy π_θ , i.e.,

$$f(\tau; \theta) d\tau := \mathcal{P}(s_0) \prod_{t=0}^{\infty} \mathcal{E}(o_t | s_t) \pi_\theta(a_t | o_t) \mathcal{C}(c_t | o_t, a_t) \mathcal{P}(s_{t+1} | s_t, a_t) d\tau,$$

For a trajectory τ , a random variable $R(\tau)$ represent the cumulative γ -discounted costs in τ , i.e., $R(\tau) = \sum_{t=0}^{\infty} \gamma^t c_t$. Therefore, solving a POMDP is to find a policy π_θ that minimizes the following optimization objectives

$$\eta(\theta) := \mathbb{E}_{\tau \sim f}[R(\tau)] = \int_{\tau} f(\tau; \theta) R(\tau) d\tau \longrightarrow \theta^* \in \arg \min_{\theta \in \Theta} \eta(\theta).$$

And let $V_\theta(o_t) := \mathbb{E}_{\tau \sim f} \left[\sum_{t'=t}^{\infty} \gamma^{t'-t} c_{t'} \mid o_t \right]$, $Q_\theta(o_t, a_t) := \mathbb{E}_{\tau \sim f} \left[\sum_{t'=t}^{\infty} \gamma^{t'-t} c_{t'} \mid o_t, a_t \right]$. Now we extend the main Policy Gradient Theorem (Sutton et al., 2000) to episodic POMDPs, derive by simple modifications in their predecessors.

Theorem C.1 (Policy Gradient in POMDPs) For a given policy π_θ on a POMDP, $\nabla_\theta \eta(\theta) = \mathbb{E}_{\tau \sim f} \left[\sum_{t=0}^{\infty} \nabla_\theta \log \pi_\theta(a_t \mid o_t) Q_\theta(o_t, a_t) \right]$.

Proof For a given policy π_θ , lets restate the value function at a given enriched observation o_0 : $V_\theta(o_0) = \int_{a_0} \pi_\theta(a_0 \mid o_0) Q_\theta(o_0, a_0) da_0$. For this value function, we compute the gradient with respect to parameters θ , i.e.,

$$\nabla_\theta V_\theta(o_0) = \int_{a_0} \nabla_\theta \pi_\theta(a_0 \mid o_0) Q_\theta(o_0, a_0) da_0 + \int_{a_0} \pi_\theta(a_0 \mid o_0) \nabla_\theta Q_\theta(o_0, a_0) da_0.$$

To further expand this gradient, consider the definition of $Q_\theta(o_0, a_0)$ using Bellman equation (the enriched observation space satisfies Markov property (Mao et al., 2020)),

$$\begin{aligned} Q_\theta(o_0, a_0) &= \int_{s_1} \mathcal{P}(s_1 \mid s_0, a_0) \mathcal{C}(c_0 \mid s_0, a_0, s_1) ds_1 \\ &\quad + \gamma \int_{s_1, o_1} \mathcal{P}(s_1 \mid s_0, a_0) \mathcal{E}(o_1 \mid s_1) V_\theta(o_1) ds_1 do_1, \end{aligned}$$

resulting in, $\nabla_\theta Q_\theta(o_0, a_0) = \gamma \int_{s_1, o_1} \mathcal{P}(s_1 \mid s_0, a_0) \mathcal{E}(o_1 \mid s_1) \nabla_\theta V_\theta(o_1) ds_1 do_1$, since the first term in the definition of $Q_\theta(o_0, a_0)$ does not depend on the parameters θ . Following these steps, we derive $\nabla_\theta V_\theta(o_1)$,

$$\nabla_\theta V_\theta(o_1) = \int_{a_1} \nabla_\theta \pi_\theta(a_1 \mid o_1) Q_\theta(o_1, a_1) da_1 + \int_{a_1} \pi_\theta(a_1 \mid o_1) \nabla_\theta Q_\theta(o_1, a_1) da_1.$$

We recursively compute the gradient of the value functions for later time steps and conclude that,

$$\begin{aligned} \nabla_\theta V_\theta(o_t) &= \int_{a_t} \nabla_\theta \pi_\theta(a_t \mid o_t) Q_\theta(o_t, a_t) da_t \\ &\quad + \gamma \int_{a_t} \pi_\theta(a_t \mid o_t) \left(\int_{s_{t+1}, o_{t+1}} \mathcal{P}(s_{t+1} \mid s_t, a_t) \mathcal{E}(o_{t+1} \mid s_{t+1}) \nabla_\theta V_\theta(o_{t+1}) ds_{t+1} do_{t+1}, \right) da_t. \end{aligned}$$

and repeating this decomposition results in

$$\begin{aligned}
 \nabla_{\theta}\eta(\theta) &= \int_{s_0, o_0} \mathcal{P}(s_0)\mathcal{E}(o_0 | s_0)\nabla_{\theta}V_{\theta}(o_0) ds_0do_0 \\
 &= \int_{s_0, o_0} \mathcal{P}(s_0)\mathcal{E}(o_0 | s_0)\nabla_{\theta} \int_{a_0} \pi_{\theta}(a_0 | o_0) Q_{\theta}(o_0, a_0) da_0 ds_0do_0 \\
 &= \int_{s_0, o_0, a_0} \mathcal{P}(s_0)\mathcal{E}(o_0 | s_0)\nabla_{\theta}\pi_{\theta}(a_0 | o_0) Q_{\theta}(o_0, a_0) da_0 ds_0do_0 \\
 &\quad + \int_{s_0, o_0, a_0} \mathcal{P}(s_0)\mathcal{E}(o_0 | s_0)\pi_{\theta}(a_0 | o_0) \nabla_{\theta}Q_{\theta}(o_0, a_0) da_0 ds_0do_0 \\
 &= \int_{\tau} \sum_{t=0}^{\infty} \gamma^t f(\tau_{0..t-1}, s_t, o_t; \theta) \nabla_{\theta}\pi_{\theta}(a_t | o_t) Q_{\theta}(o_t, a_t) d\tau \\
 &= \int_{\tau} \sum_{t=0}^{\infty} \gamma^t f(\tau_{0..t-1}, s_t, o_t; \theta) \pi_{\theta}(a_t | o_t) \nabla_{\theta} \log \pi_{\theta}(a_t | o_t) Q_{\theta}(o_t, a_t) d\tau \\
 &= \int_{\tau} \sum_{t=0}^{\infty} \gamma^t f(\tau_{0..t}; \theta) \nabla_{\theta} \log \pi_{\theta}(a_t | o_t) Q_{\theta}(o_t, a_t) d\tau \\
 &= \mathbb{E}_{\tau \sim f} \left[\sum_{t=0}^{\infty} \gamma^t \nabla_{\theta} \log \pi_{\theta}(a_t | o_t) Q_{\theta}(o_t, a_t) \right].
 \end{aligned}$$

■

It is worth noting that in most open-source implementations (Dhariwal et al., 2017; Raffin et al., 2021) of reinforcement learning algorithms based on policy gradient theorem, the term γ^t in the policy gradient is ignored, and a biased policy gradient estimator $\mathbb{E}_{\tau \sim f} [\sum_{t=0}^{\infty} \nabla_{\theta} \log \pi_{\theta}(a_t | o_t) Q_{\theta}(o_t, a_t)]$ is obtained. However, it has remained the most popular estimator of the policy gradient due to its effectiveness when applied to practical problems. The precise reason for this effectiveness, especially in the episodic setting, remains an open question (Nota and Thomas, 2020). An in-depth exploration of this issue is beyond the scope of this paper, and we still use a biased policy gradient estimator in all algorithm implementations.

It is worth noting that the on-policy policy gradient have two different but equivalent forms: *trajectory-oriented* and *state-oriented*. The detailed derivation process of the former in the POMDP has been developed above. We now derive the second form in the POMDP, which is also the standard form of Sutton’s policy gradient theorem. Since this form is more compact, this paper mainly derives the relevant policy gradient based on it. Specifically, we can write the optimization objective of reinforcement learning in another equivalent form

$$J(\theta) = \int_{s, o} d_{\pi}(s)\mathcal{E}(o | s)V_{\theta}(o)dsdo = \int_{s, o, a} d_{\pi}(s)\mathcal{E}(o | s)\pi_{\theta}(a | o)Q_{\theta}(o, a)dsdoda,$$

where d_{π} represents the distribution of the state-occupancy measure of policy π . Then we have following theorem:

Theorem C.2 (Policy Gradient in POMDPs) For a given policy π_θ on a POMDP, $\nabla_\theta J(\theta) \propto \mathbb{E}_\pi [\nabla_\theta \log \pi_\theta(a | o) Q_\theta(o, a)]$, where \mathbb{E}_π refers to $\mathbb{E}_{s \sim d_\pi, o \sim \mathcal{E}, a \sim \pi_\theta}$ when both state, enriched observation and action distributions follow the policy π_θ (on policy).

Proof We first start with the derivative of the state value function:

$$\begin{aligned}
 & \nabla_\theta V_\theta(o) \\
 = & \nabla_\theta \left(\int_a \pi_\theta(a | o) Q_\theta(o, a) da \right) \\
 = & \int_a (\nabla_\theta \pi_\theta(a | o) Q_\theta(o, a) + \pi_\theta(a | o) \nabla_\theta Q_\theta(o, a)) da \\
 = & \int_a (\nabla_\theta \pi_\theta(a | o) Q_\theta(o, a) + \pi_\theta(a | o) \\
 & \quad \nabla_\theta \left(\int_{s', o', c} \mathcal{P}(s' | s, a) \mathcal{E}(o' | s') (\mathcal{C}(c | s, a, s') + V_\theta(o')) \right) ds' do' dc) da \\
 = & \int_a \left(\nabla_\theta \pi_\theta(a | o) Q_\theta(o, a) + \pi_\theta(a | o) \left(\int_{s', o'} \mathcal{P}(s' | s, a) \mathcal{E}(o' | s') \nabla_\theta V_\theta(o') \right) ds' do' \right) da.
 \end{aligned}$$

This equation has a nice recursive form and the future value function $V_\theta(o')$ can be repeated unrolled by following the same equation.

Let's consider the following visitation sequence and label the probability of transitioning from observation o to observation x with policy π_θ after k step as $\rho^\pi(o \rightarrow x, k)$.

$$o \xrightarrow{o \sim \pi_\theta(\cdot | o)} o' \xrightarrow{a \sim \pi_\theta(\cdot | o')} o'' \xrightarrow{a \sim \pi_\theta(\cdot | o'')} \dots$$

- When $k = 0$: $\rho^\pi(o \rightarrow o, k = 0) = 1$.
- When $k = 1$, we scan through all possible actions and sum up the transition probabilities to the target observation: $\rho^\pi(o \rightarrow o', k = 1) = \int_{a, s'} \pi_\theta(a | o) \mathcal{P}(s' | s, a) \mathcal{E}(o' | s') da ds'$.
- Imagine that the goal is to go from observation o to x after $k+1$ steps while following policy π_θ . We can first travel from o to a middle point o' (any observation can be a middle point, $o' \in \mathcal{O}$) after k steps and then go to the final observation x during the last step. In this way, we are able to update the visitation probability recursively: $\rho^\pi(o \rightarrow x, k+1) = \int_{o'} \rho^\pi(o \rightarrow o', k) \rho^\pi(o' \rightarrow x, 1) do'$.

Then we go back to unroll the recursive representation of $\nabla_\theta V_\theta(o)$. Let $\phi(o) = \int_a \nabla_\theta \pi_\theta(a | o) Q_\theta(o, a) da$ to simplify the maths. If we keep on extending $\nabla_\theta V_\theta(\cdot)$ infinitely, it is easy to find out that we can transition from the starting observation o to any observation after any number of steps in this unrolling process and by summing up all the visitation probabilities, we get $\nabla_\theta V_\theta(o)$.

$$\begin{aligned}
 & \nabla_\theta V_\theta(o) \\
 = & \phi(o) + \int_a \left(\pi_\theta(a | o) \left(\int_{s', o'} \mathcal{P}(s' | s, a) \mathcal{E}(o' | s') \nabla_\theta V_\theta(o') \right) ds' do' \right) da \\
 = & \phi(o) + \int_{s', o', a} \pi_\theta(a | o) \mathcal{P}(s' | s, a) \mathcal{E}(o' | s') \nabla_\theta V_\theta(o') ds' do' da
 \end{aligned}$$

$$\begin{aligned}
 &= \phi(o) + \int_{o'} \rho^\pi(o \rightarrow o', 1) \nabla_\theta V_\theta(o') do' \\
 &= \phi(o) + \int_{o'} \rho^\pi(o \rightarrow o', 1) \left[\phi(o') + \int_{o''} \rho^\pi(o' \rightarrow o'', 1) \nabla_\theta V_\theta(o'') do'' \right] do' \\
 &= \phi(o) + \int_{o'} \rho^\pi(o \rightarrow o', 1) \phi(o') do' + \int_{o''} \rho^\pi(o \rightarrow o'', 2) \nabla_\theta V_\theta(o'') do'' \\
 &= \phi(o) + \int_{o'} \rho^\pi(o \rightarrow o', 1) \phi(o') do' + \int_{o''} \rho^\pi(o \rightarrow o'', 2) \phi(o'') do'' + \\
 &\quad \int_{o'''} \rho^\pi(o \rightarrow o''', 3) \nabla_\theta V_\theta(o''') do''' \\
 &= \int_x \sum_{k=0}^{\infty} \rho^\pi(o \rightarrow x, k) \phi(x) dx
 \end{aligned}$$

The nice rewriting above allows us to exclude the derivative of Q -value function, $\nabla_\theta Q_\theta(o, a)$. By plugging it into the objective function $J(\theta)$, we are getting the following:

$$\begin{aligned}
 \nabla_\theta J(\theta) &= \int_{s_0, o_0} \mathcal{P}(s_0) \mathcal{E}(o_0 | s_0) \nabla_\theta V_\theta(o_0) ds_0 do_0 \\
 &= \int_{s_0, o_0} \mathcal{P}(s_0) \mathcal{E}(o_0 | s_0) \int_{o'} \sum_{k=0}^{\infty} \rho^\pi(o_0 \rightarrow o', k) \phi(o') do' ds_0 do_0 \\
 &= \int_{s_0, o_0} \mathcal{P}(s_0) \mathcal{E}(o_0 | s_0) \int_{o'} \eta(o') \phi(o') do' ds_0 do_0 \\
 &= \int_{s_0, o_0} \mathcal{P}(s_0) \mathcal{E}(o_0 | s_0) \left(\int_{o'} \eta(o') do' \right) \int_{o'} \frac{\eta(o')}{\int_{o'} \eta(o') do'} \phi(o') do' ds_0 do_0 \\
 &\propto \int_{s_0, o_0} \mathcal{P}(s_0) \mathcal{E}(o_0 | s_0) \int_{o'} \frac{\eta(o')}{\int_{o'} \eta(o') do'} \phi(o') do' ds_0 do_0 \\
 &= \int_{s_0, o_0} \mathcal{P}(s_0) \mathcal{E}(o_0 | s_0) \int_{o'} d_\pi(o') \phi(o') do' ds_0 do_0 \\
 &= \int_{o'} d_\pi(o') \phi(o') do' = \int_{s, o, a} d_\pi(s) \mathcal{E}(o | s) \nabla_\theta \pi_\theta(a | o) Q_\theta(o, a) ds do da \\
 &= \int_{s, o, a} d_\pi(s) \mathcal{E}(o | s) \pi_\theta(a | o) \frac{\nabla_\theta \pi_\theta(a | o)}{\pi_\theta(a | o)} Q_\theta(o, a) ds do da \\
 &= \mathbb{E}_\pi [\nabla_\theta \log \pi_\theta(a | o) Q_\theta(o, a)].
 \end{aligned}$$

■

It can be seen that the terms within the two different policy gradient expectations are the same (when γ^t is ignored), but the distributions on which the expectations are based

are different. The former is based on trajectory distribution and the latter is based on stationary state distribution.

When there are partial observations, policy gradient theorem can be a bit subtle and the derivations in this work might require more careful scrutiny. Fortunately, after proper enrichment of the observation space of the original POMDP, we can naturally extend the policy gradient theory under MDP to POMDP with minor modification, and use it as the basis for the derivation of the policy gradient of the proposed MARL algorithm under POSG in this paper.

Appendix D. Missing Proofs

D.1 Proof of the Proposition 3

Proof We proof the first equation first. Extend the COMA algorithm, we have (for convenience here we suppose state space, observation space and action space are discrete)

$$\begin{aligned}
 J_{\text{actor}}^i(\mathbf{w}_{in}, \tilde{\mathbf{w}}_{sh}^i) &= \mathbf{E}_{s \sim d_{\pi}, \mathbf{o} \sim \mathcal{E}, \mathbf{a} \sim \pi} \left[Q_{\tilde{\phi}_{sh}^i}^{\pi, i}(\mathbf{o}, \mathbf{a}) - \mathcal{B}(\mathbf{o}, \mathbf{a}^i) \right] \\
 &= \sum_s d_{\pi}(s) \sum_{\mathbf{o}} \mathcal{E}(\mathbf{o}|s) \sum_{\mathbf{a}} \pi_{\Psi_{in}}(\mathbf{a}|\mathbf{o}) \left[Q_{\tilde{\phi}_{sh}^i}^{\pi, i}(\mathbf{o}, \mathbf{a}) - \mathcal{B}(\mathbf{o}, \mathbf{a}^i) \right] \\
 J_{\text{critic}}^i(\mathbf{w}_{in}, \tilde{\mathbf{w}}_{sh}^i) &= \mathbf{E}_{s \sim d_{\pi}, \mathbf{o} \sim \mathcal{E}, \mathbf{a} \sim \pi} \left[\left(Q_{\tilde{\phi}_{sh}^i}^{\pi, i}(\mathbf{o}, \mathbf{a}) - Q_{tg}^{\pi} \right)^2 \right] \\
 &= \sum_s d_{\pi}(s) \sum_{\mathbf{o}} \mathcal{E}_s^{\mathbf{o}} \sum_{\mathbf{a}} \pi_{\Psi_{in}}(\mathbf{a}|\mathbf{o}) \left[\left(Q_{\tilde{\phi}_{sh}^i}^{\pi, i}(\mathbf{o}, \mathbf{a}) - Q_{tg}^{\pi} \right)^2 \right],
 \end{aligned}$$

where the specific form of Q_{tg}^{π} is described in Section 4.2. We hypothesis joint policy $\pi_{\Psi_{in}}$ is the product of local policy functions $\prod_{i=1}^n \pi_{\psi_{in}^i}^i$. Hence **the actor part** gradient w.r.t. each parameter ψ_{in}^i becomes:

$$\begin{aligned}
 \nabla_{\psi_{in}^i} J_{\text{actor}}^i(\mathbf{w}_{in}, \tilde{\mathbf{w}}_{sh}^i) &= \nabla_{\psi_{in}^i} \sum_s d_{\pi}(s) \sum_{\mathbf{o}} \mathcal{E}(\mathbf{o}|s) \sum_{\mathbf{a}} \pi_{\Psi_{in}}(\mathbf{a}|\mathbf{o}) \left[Q_{\tilde{\phi}_{sh}^i}^{\pi, i}(\mathbf{o}, \mathbf{a}) - \mathcal{B}(\mathbf{o}, \mathbf{a}^i) \right] \\
 &= \sum_s d_{\pi}(s) \sum_{\mathbf{o}} \mathcal{E}(\mathbf{o}|s) \sum_{\mathbf{a}} \nabla_{\psi_{in}^i} \pi_{\Psi_{in}}(\mathbf{a}|\mathbf{o}) \left[Q_{\tilde{\phi}_{sh}^i}^{\pi, i}(\mathbf{o}, \mathbf{a}) - \mathcal{B}(\mathbf{o}, \mathbf{a}^i) \right] \\
 &= \sum_s d_{\pi}(s) \sum_{\mathbf{o}} \mathcal{E}(\mathbf{o}|s) \sum_{\mathbf{a}} \pi_{\Psi_{in}}(\mathbf{a}|\mathbf{o}) \nabla_{\psi_{in}^i} \log \pi_{\psi_{in}^i}^i(a^i|o^i) \left[Q_{\tilde{\phi}_{1,sh}^i}^{\pi, i}(\mathbf{o}, \mathbf{a}) - \mathcal{B}(\mathbf{o}, \mathbf{a}^i) \right] \\
 &= \mathbf{E}_{s \sim d_{\pi}, \mathbf{o} \sim \mathcal{E}, \mathbf{a} \sim \pi} \left[\nabla_{\psi_{in}^i} \log \pi_{\psi_{in}^i}^i(a^i|o^i) \left(Q_{\tilde{\phi}_{1,sh}^i}^{\pi, i}(\mathbf{o}, \mathbf{a}) - \mathcal{B}(\mathbf{o}, \mathbf{a}^i) \right) \right].
 \end{aligned}$$

For the critic part, the gradient sampled by current policy is calculated by:

$$\nabla_{\psi_{in}^i} J_{\text{critic}}^i(\mathbf{w}_{in}, \tilde{\mathbf{w}}_{sh}^i) = \mathbf{E}_{s \sim d_{\pi}, \mathbf{o} \sim \mathcal{E}, \mathbf{a} \sim \pi} \left[\nabla_{\psi_{in}^i} \log \pi_{\psi_{in}^i}^i(a^i|o^i) \delta^2 \right].$$

where $\delta = Q_{\tilde{\phi}_{sh}^i}^{\pi, i}(\mathbf{o}, \mathbf{a}) - Q_{tg}^{\pi}$.

Finally, we get the following on-policy joint gradient w.r.t. actor parameters:

$$\begin{aligned} \nabla_{\psi_{in}^i} J_{ac}^i(\mathbf{w}_{in}, \tilde{\mathbf{w}}_{sh}^i) &= \alpha_1 \nabla_{\psi_{in}^i} J_{actor}^i(\mathbf{w}_{in}, \tilde{\mathbf{w}}_{sh}^i) + \alpha_2 \nabla_{\psi_{in}^i} J_{critic}^i(\mathbf{w}_{in}, \tilde{\mathbf{w}}_{sh}^i) \\ &= \mathbf{E}_{s \sim d_\pi, \mathbf{o} \sim \mathcal{E}, \mathbf{a} \sim \pi} \left[\alpha_1 \nabla_{\psi_{in}^i} \log \pi_{\psi_{in}^i}^i(a^i | \mathbf{o}^i) \left(Q_{\tilde{\phi}_{sh}^i}^{\pi, i}(\mathbf{o}, \mathbf{a}) - \mathcal{B}(\mathbf{o}, \mathbf{a}^i) + \frac{\alpha_2}{\alpha_1} \delta^2 \right) \right]. \end{aligned}$$

The above is the proof of the first equation, **below we prove the second. For actor part**, we have

$$\begin{aligned} \nabla_{\tilde{\phi}_{sh}^i} J_{actor}^i(\mathbf{w}_{in}, \tilde{\mathbf{w}}_{sh}^i) &= \nabla_{\tilde{\phi}_{sh}^i} \sum_s d_\pi(s) \sum_{\mathbf{o}} \mathcal{E}(\mathbf{o} | s) \sum_{\mathbf{a}} \pi_{\Psi_{in}}(\mathbf{a} | \mathbf{o}) \left[Q_{\tilde{\phi}_{sh}^i}^{\pi, i}(\mathbf{o}, \mathbf{a}) - \mathcal{B}(\mathbf{o}, \mathbf{a}^i) \right] \\ &= \sum_s d_\pi(s) \sum_{\mathbf{o}} \mathcal{E}(\mathbf{o} | s) \sum_{\mathbf{a}} \pi_{\Psi_{in}}(\mathbf{a} | \mathbf{o}) \nabla_{\tilde{\phi}_{sh}^i} Q_{\tilde{\phi}_{sh}^i}^{\pi, i}(\mathbf{o}, \mathbf{a}) \\ &= \mathbf{E}_{s \sim d_\pi, \mathbf{o} \sim \mathcal{E}, \mathbf{a} \sim \pi} \left[\nabla_{\tilde{\phi}_{sh}^i} Q_{\tilde{\phi}_{sh}^i}^{\pi, i}(\mathbf{o}, \mathbf{a}) \right]. \end{aligned}$$

For critic part, we have

$$\begin{aligned} \nabla_{\tilde{\phi}_{sh}^i} J_{critic}^i(\mathbf{w}_{in}, \tilde{\mathbf{w}}_{sh}^i) &= \nabla_{\tilde{\phi}_{1, sh}^i} \sum_s d_\pi(s) \sum_{\mathbf{o}} \mathcal{E}_s^{\mathbf{o}} \sum_{\mathbf{a}} \pi_{\Psi_{in}}(\mathbf{a} | \mathbf{o}) \left(Q_{\tilde{\phi}_{sh}^i}^{\pi, i}(\mathbf{o}, \mathbf{a}) - Q_{tg}^{\pi} \right)^2 \\ &= \sum_s d_\pi(s) \sum_{\mathbf{o}} \mathcal{E}_s^{\mathbf{o}} \sum_{\mathbf{a}} \pi_{\Psi_{in}}(\mathbf{a} | \mathbf{o}) 2\delta \nabla_{\tilde{\phi}_{sh}^i} Q_{\tilde{\phi}_{sh}^i}^{\pi, i}(\mathbf{o}, \mathbf{a}) \\ &= \mathbf{E}_{s \sim d_\pi, \mathbf{o} \sim \mathcal{E}, \mathbf{a} \sim \pi} \left[2\delta \nabla_{\tilde{\phi}_{sh}^i} Q_{\tilde{\phi}_{sh}^i}^{\pi, i}(\mathbf{o}, \mathbf{a}) \right]. \end{aligned}$$

Finally, we get following on-policy joint gradient w.r.t. the first critic parameters

$$\begin{aligned} \nabla_{\tilde{\phi}_{sh}^i} J_{ac}^i(\mathbf{w}_{in}, \tilde{\mathbf{w}}_{sh}^i) &= \alpha_1 \nabla_{\tilde{\phi}_{sh}^i} J_{actor}^i(\mathbf{w}_{in}, \tilde{\mathbf{w}}_{sh}^i) + \alpha_2 \nabla_{\tilde{\phi}_{sh}^i} J_{critic}^i(\mathbf{w}_{in}, \tilde{\mathbf{w}}_{sh}^i) \\ &= \mathbf{E}_{s \sim d_\pi, \mathbf{o} \sim \mathcal{E}, \mathbf{a} \sim \pi} \left[(\alpha_1 + 2\alpha_2\delta) \nabla_{\tilde{\phi}_{sh}^i} Q_{\tilde{\phi}_{sh}^i}^{\pi, i}(\mathbf{o}, \mathbf{a}) \right]. \end{aligned}$$

■

D.2 Proof of Proposition 4

Proof We proof the first equation first. Extend the DDPG algorithm, we have⁴

$$\begin{aligned} J_{actor}^i(\mathbf{w}_{in}, \tilde{\mathbf{w}}_{sh}^i) &= \mathbf{E}_{s \sim d_0, \mathbf{o} \sim \mathcal{E}, \mathbf{a} \sim \pi} \left[Q_{\tilde{\phi}_{sh}^i}^{\pi, i}(\mathbf{o}, \mathbf{a}) \right] \\ &= \sum_s d_0(s) \sum_{\mathbf{o}} \mathcal{E}(\mathbf{o} | s) \sum_{\mathbf{a}} \pi_{\Psi_{in}}(\mathbf{a} | \mathbf{o}) \left[Q_{\tilde{\phi}_{sh}^i}^{\pi, i}(\mathbf{o}, \mathbf{a}) \right] \\ J_{critic}^i(\mathbf{w}_{in}, \tilde{\mathbf{w}}_{sh}^i) &= \mathbf{E}_{s \sim d_0, \mathbf{o} \sim \mathcal{E}, \mathbf{a} \sim \pi_0} \left[\left(Q_{\tilde{\phi}_{sh}^i}^{\pi, i}(\mathbf{o}, \mathbf{a}) - Q_{tg}^{\pi} \right)^2 \right] \end{aligned}$$

4. Actually, there is no expectation of action because the policy is deterministic, i.e., $\mathbf{a} = \pi_{\Psi_{in}}(\mathbf{o})$. However, in order to keep the form of the gradient consistent with other algorithms, we still retain the expectation of the action here, which does not affect the calculation of the gradient.

$$= \sum_s d_0(s) \sum_{\mathbf{o}} \mathcal{E}_{0,s}^{\mathbf{o}} \sum_{\mathbf{a}} \pi_{\mathbf{0}}(\mathbf{a}|\mathbf{o}) \left[\left(Q_{\tilde{\phi}_{sh}^i}^{\pi,i}(\mathbf{o}, \mathbf{a}) - Q_{tg}^{\pi} \right)^2 \right],$$

where $Q_{tg}^{\pi} = \sum_{s'} \mathcal{P}_{s,\mathbf{a}}^{s'} \left(\mathcal{C}_{s,\mathbf{a}}^{i,s'} + \gamma \sum_{\mathbf{o}'} \mathcal{E}_{0,s'}^{\mathbf{o}'} \sum_{\mathbf{a}'} \pi_{\Psi_{in}}(\mathbf{a}'|\mathbf{o}') Q_{\tilde{\phi}_{sh}^i}^{\pi,i}(\mathbf{o}', \mathbf{a}') \right)$. Note that the joint policy $\pi_{\Psi_{in}}$ is a deterministic policy.

We hypothesis joint policy $\pi_{\Psi_{in}}$ is the product of local policy functions $\prod_{i=1}^n \pi_{\psi_{in}^i}$. Hence the gradient concerning each parameter ψ_{in}^i becomes

$$\begin{aligned} \nabla_{\psi_{in}^i} J_{\text{actor}}^i(\mathbf{w}_{in}, \tilde{\mathbf{w}}_{sh}^i) &= \nabla_{\psi_{in}^i} \sum_s d_0(s) \sum_{\mathbf{o}} \mathcal{E}(\mathbf{o}|s) \sum_{\mathbf{a}} \pi_{\Psi_{in}}(\mathbf{a}|\mathbf{o}) \left[Q_{\tilde{\phi}_{sh}^i}^{\pi,i}(\mathbf{o}, \mathbf{a}) \right] \\ &= \sum_s d_0(s) \sum_{\mathbf{o}} \mathcal{E}(\mathbf{o}|s) \sum_{\mathbf{a}} \pi_{\Psi_{in}}(\mathbf{a}|\mathbf{o}) \left[\nabla_{\psi_{in}^i} Q_{\tilde{\phi}_{sh}^i}^{\pi,i}(\mathbf{o}, \mathbf{a}) \right] \\ &= \sum_s d_0(s) \sum_{\mathbf{o}} \mathcal{E}(\mathbf{o}|s) \sum_{\mathbf{a}} \pi_{\Psi_{in}}(\mathbf{a}|\mathbf{o}) \left[\nabla_{\psi_{in}^i} \pi_{\psi_{in}^i}^i(a^i|o^i) \nabla_{a^i} Q_{\tilde{\phi}_{sh}^i}^{\pi,i}(\mathbf{o}, \mathbf{a}) \right] \\ &= \mathbf{E}_{s \sim d_0, \mathbf{o} \sim \mathcal{E}, \mathbf{a} \sim \pi} \left[\nabla_{\psi_{in}^i} \pi_{\psi_{in}^i}^i(a^i|o^i) \nabla_{a^i} Q_{\tilde{\phi}_{sh}^i}^{\pi,i}(\mathbf{o}, \mathbf{a}) \right]. \end{aligned}$$

Note that here we need to resample to calculate the unbiased policy gradient so that we can't use the off-policy data directly. We solve the above problem by importance sampling. **For off-policy data saved in experience replay buffer** we have

$$\nabla_{\psi_{in}^i} J_{\text{actor}}^i(\mathbf{w}_{in}, \tilde{\mathbf{w}}_{sh}^i) = \mathbf{E}_{s \sim d_0, \mathbf{o} \sim \mathcal{E}, \mathbf{a} \sim \pi_{\mathbf{0}}} \left[\left(\frac{\pi_{\psi_{in}^i}^i(a^i|o^i)}{\pi_{\mathbf{0}}^i(a^i|o^i)} \right) \nabla_{\psi_{in}^i} \pi_{\psi_{in}^i}^i(a^i|o^i) \nabla_{a^i} Q_{\tilde{\phi}_{sh}^i}^{\pi,i}(\mathbf{o}, \mathbf{a}) \right] = 0.$$

combined the gradient calculated by the data which is resampled by current policy, we get the joint off-policy policy gradient associated with the actor:

$$\nabla_{\psi_{in}^i} J_{\text{actor}}^i(\mathbf{w}_{in}, \tilde{\mathbf{w}}_{sh}^i) = \mathbf{E}_{s \sim d_0, \mathbf{o} \sim \mathcal{E}, \mathbf{a} \sim \pi} \left[\nabla_{\psi_{in}^i} \pi_{\psi_{in}^i}^i(a^i|o^i) \nabla_{a^i} Q_{\tilde{\phi}_{sh}^i}^{\pi,i}(\mathbf{o}, \mathbf{a}) \right].$$

On the contrary, we directly use the off-policy data when calculate the value function gradient, so that we can't use the above resampled data. **For the critic part**, we use above trick in reverse. **For off-policy data saved in experience replay buffer** we have

$$\nabla_{\psi_{in}^i} J_{\text{critic}}^i(\mathbf{w}_{in}, \tilde{\mathbf{w}}_{sh}^i) = \nabla_{\psi_{in}^i} \sum_s d_0(s) \sum_{\mathbf{o}} \mathcal{E}_{0,s}^{\mathbf{o}} \sum_{\mathbf{a}} \pi_{\mathbf{0}}(\mathbf{a}|\mathbf{o}) \left[\left(Q_{\tilde{\phi}_{sh}^i}^{\pi,i}(\mathbf{o}, \mathbf{a}) - Q_{tg}^{\pi} \right)^2 \right] = 0.$$

Combining the gradient calculated by data which is resampled by current policy, we get the joint off-policy policy gradient associates with the critic part⁵:

$$\nabla_{\psi_{in}^i} J_{\text{critic}}^i(\mathbf{w}_{in}, \tilde{\mathbf{w}}_{sh}^i) = \mathbf{E}_{s \sim d_0, \mathbf{o} \sim \mathcal{E}, \mathbf{a} \sim \pi} \left[2\delta \left(\frac{\pi_{\mathbf{0}}^i(a^i|o^i)}{\pi_{\psi_{in}^i}^i(a^i|o^i)} \right) \nabla_{\psi_{in}^i} \pi_{\psi_{in}^i}^i(a^i|o^i) \left(\nabla_{a^i} Q_{\tilde{\phi}_{sh}^i}^{\pi,i}(\mathbf{o}, \mathbf{a}) \right) \right].$$

5. Since the policy is deterministic, we cannot directly get the specific values of the denominator in the importance ratio during implementation. Fortunately, we can draw on a probabilistic reinforcement learning framework (Levine, 2018) to estimate the denominator.

where $\delta = Q_{\tilde{\phi}_{sh}^i}^{\pi, i}(\mathbf{o}, \mathbf{a}) - Q_{tg}^{\pi}$. We denote the clipped importance sampling term $\min\left(\epsilon, \frac{\pi_{\psi_{in}^i}^i(a^i|o^i)}{\pi_0^i(a^i|o^i)}\right)$ as $CIM_{\epsilon}(\pi_{\psi_{in}^i}^i; \pi_0^i)$ and $\min\left(\epsilon, \frac{\pi_0^i(a^i|o^i)}{\pi_{\psi_{in}^i}^i(a^i|o^i)}\right)$ as $CIM_{\epsilon}(\pi_0^i; \pi_{\psi_{in}^i}^i)$. Finally, we get following off-policy joint gradient w.r.t. actor parameters

$$\begin{aligned} \nabla_{\psi_{in}^i} J_{ac}^i(\mathbf{w}_{in}, \tilde{\mathbf{w}}_{sh}^i) &= \alpha_1 \nabla_{\psi_{in}^i} J_{actor}^i(\mathbf{w}_{in}, \tilde{\mathbf{w}}_{sh}^i) + \alpha_2 \nabla_{\psi_{in}^i} J_{critic}^i(\mathbf{w}_{in}, \tilde{\mathbf{w}}_{sh}^i) \\ &= \mathbf{E}_{s \sim d_0, \mathbf{o} \sim \mathcal{E}, \mathbf{a} \sim \pi} \left[(\alpha_1 + 2\alpha_2 \delta \left(\frac{\pi_0^i(a^i|o^i)}{\pi_{\psi_{in}^i}^i(a^i|o^i)} \right)) \nabla_{\psi_{in}^i} \pi_{\psi_{in}^i}^i(a^i|o^i) \nabla_{\mathbf{a}^i} Q_{\tilde{\phi}_{sh}^i}^{\pi, i}(\mathbf{o}, \mathbf{a}) \right] \\ &\simeq \mathbf{E}_{s \sim d_0, \mathbf{o} \sim \mathcal{E}, \mathbf{a} \sim \pi} \left[(\alpha_1 + 2\alpha_2 \delta CIM_{\epsilon}(\pi_0^i; \pi_{\psi_{in}^i}^i)) \nabla_{\psi_{in}^i} \pi_{\psi_{in}^i}^i(a^i|o^i) \nabla_{\mathbf{a}^i} Q_{\tilde{\phi}_{sh}^i}^{\pi, i}(\mathbf{o}, \mathbf{a}) \right]. \end{aligned}$$

The above is the proof of the first equation, **below we prove the second. For actor part, we first calculate the gradient use the data resampled by the current policy**

$$\begin{aligned} \nabla_{\tilde{\phi}_{sh}^i} J_{actor}^i(\mathbf{w}_{in}, \tilde{\mathbf{w}}_{sh}^i) &= \nabla_{\tilde{\phi}_{sh}^i} \sum_s d_0(s) \sum_{\mathbf{o}} \mathcal{E}(\mathbf{o}|s) \sum_{\mathbf{a}} \pi_{\Psi_{in}}(\mathbf{a}|\mathbf{o}) \left[Q_{\tilde{\phi}_{sh}^i}^{\pi, i}(\mathbf{o}, \mathbf{a}) \right] \\ &= \sum_s d_0(s) \sum_{\mathbf{o}} \mathcal{E}(\mathbf{o}|s) \sum_{\mathbf{a}} \pi_{\Psi_{in}}(\mathbf{a}|\mathbf{o}) \nabla_{\tilde{\phi}_{sh}^i} Q_{\tilde{\phi}_{sh}^i}^{\pi, i}(\mathbf{o}, \mathbf{a}) \\ &= \mathbf{E}_{s \sim d_0, \mathbf{o} \sim \mathcal{E}, \mathbf{a} \sim \pi} \left[\nabla_{\tilde{\phi}_{sh}^i} Q_{\tilde{\phi}_{sh}^i}^{\pi, i}(\mathbf{o}, \mathbf{a}) \right]. \end{aligned}$$

For calculate the gradient use off-policy data we also introduce importance sampling, then we have

$$\nabla_{\tilde{\phi}_{sh}^i} J_{actor}^i(\mathbf{w}_{in}, \tilde{\mathbf{w}}_{sh}^i) = \mathbf{E}_{s \sim d_0, \mathbf{o} \sim \mathcal{E}, \mathbf{a} \sim \pi_0} \left[\left(\frac{\pi_{\psi_{in}^i}^i(a^i|o^i)}{\pi_0^i(a^i|o^i)} \right) \nabla_{\tilde{\phi}_{sh}^i} Q_{\tilde{\phi}_{sh}^i}^{\pi, i}(\mathbf{o}, \mathbf{a}) \right].$$

We then combine the above two part gradient:

$$\begin{aligned} \nabla_{\tilde{\phi}_{sh}^i} J_{actor}^i(\mathbf{w}_{in}, \tilde{\mathbf{w}}_{sh}^i) &= \mathbf{E}_{s \sim d_0, \mathbf{o} \sim \mathcal{E}, \mathbf{a} \sim \pi} \left[\nabla_{\tilde{\phi}_{sh}^i} Q_{\tilde{\phi}_{sh}^i}^{\pi, i}(\mathbf{o}, \mathbf{a}) \right] + \\ &\quad \mathbf{E}_{s \sim d_0, \mathbf{o} \sim \mathcal{E}, \mathbf{a} \sim \pi_0} \left[\left(\frac{\pi_{\psi_{in}^i}^i(a^i|o^i)}{\pi_0^i(a^i|o^i)} \right) \nabla_{\tilde{\phi}_{sh}^i} Q_{\tilde{\phi}_{sh}^i}^{\pi, i}(\mathbf{o}, \mathbf{a}) \right]. \end{aligned}$$

For critic part, we first calculate the gradient use the off-policy data

$$\begin{aligned} \nabla_{\tilde{\phi}_{sh}^i} J_{critic}^i(\mathbf{w}_{in}, \tilde{\mathbf{w}}_{sh}^i) &= \nabla_{\tilde{\phi}_{sh}^i} \sum_s d_0(s) \sum_{\mathbf{o}} \mathcal{E}_{0,s}^{\mathbf{o}} \sum_{\mathbf{a}} \pi_0(\mathbf{a}|\mathbf{o}) \left[\left(Q_{\tilde{\phi}_{sh}^i}^{\pi, i}(\mathbf{o}, \mathbf{a}) - Q_{tg}^{\pi} \right)^2 \right] \\ &= \sum_s d_0(s) \sum_{\mathbf{o}} \mathcal{E}_{0,s}^{\mathbf{o}} \sum_{\mathbf{a}} \pi_0(\mathbf{a}|\mathbf{o}) 2\delta \nabla_{\tilde{\phi}_{sh}^i} Q_{\tilde{\phi}_{sh}^i}^{\pi, i}(\mathbf{o}, \mathbf{a}) \\ &= \mathbf{E}_{s \sim d_0, \mathbf{o} \sim \mathcal{E}, \mathbf{a} \sim \pi_0} \left[2\delta \nabla_{\tilde{\phi}_{sh}^i} Q_{\tilde{\phi}_{sh}^i}^{\pi, i}(\mathbf{o}, \mathbf{a}) \right]. \end{aligned}$$

Then, for resampled data, we have:

$$\nabla_{\tilde{\phi}_{sh}^i} J_{\text{critic}}^i(\mathbf{w}_{in}, \tilde{\mathbf{w}}_{sh}^i) \simeq \mathbf{E}_{s \sim d_0, \mathbf{o} \sim \mathcal{E}, \mathbf{a} \sim \pi} \left[2\delta \left(\frac{\pi_0^i(a^i | o^i)}{\pi_{\psi_{in}^i}^i(a^i | o^i)} \right) \nabla_{\tilde{\phi}_{sh}^i} Q_{\tilde{\phi}_{sh}^i}^{\pi, i}(\mathbf{o}, \mathbf{a}) \right].$$

Combining the above two parts' gradient, we have:

$$\begin{aligned} \nabla_{\tilde{\phi}_{sh}^i} J_{\text{critic}}^i(\mathbf{w}_{in}, \tilde{\mathbf{w}}_{sh}^i) &= \mathbf{E}_{s \sim d_0, \mathbf{o} \sim \mathcal{E}, \mathbf{a} \sim \pi_0} \left[2\delta \nabla_{\tilde{\phi}_{sh}^i} Q_{\tilde{\phi}_{sh}^i}^{\pi, i}(\mathbf{o}, \mathbf{a}) \right] + \\ &\quad \mathbf{E}_{s \sim d_0, \mathbf{o} \sim \mathcal{E}, \mathbf{a} \sim \pi} \left[2\delta \left(\frac{\pi_0^i(a^i | o^i)}{\pi_{\psi_{in}^i}^i(a^i | o^i)} \right) \nabla_{\tilde{\phi}_{sh}^i} Q_{\tilde{\phi}_{sh}^i}^{\pi, i}(\mathbf{o}, \mathbf{a}) \right]. \end{aligned}$$

Finally, we get following off-policy joint gradient w.r.t. the critic parameters

$$\begin{aligned} \nabla_{\tilde{\phi}_{sh}^i} J_{ac}^i(\mathbf{w}_{in}, \tilde{\mathbf{w}}_{sh}^i) &= \alpha_1 \nabla_{\tilde{\phi}_{sh}^i} J_{\text{actor}}^i(\mathbf{w}_{in}, \tilde{\mathbf{w}}_{sh}^i) + \alpha_2 \nabla_{\tilde{\phi}_{sh}^i} J_{\text{critic}}^i(\mathbf{w}_{in}, \tilde{\mathbf{w}}_{sh}^i) \\ &= \mathbf{E}_{s \sim d_0, \mathbf{o} \sim \mathcal{E}, \mathbf{a} \sim \pi_0} \left[\left(\alpha_1 \left(\frac{\pi_{\psi_{in}^i}^i(a^i | o^i)}{\pi_0^i(a^i | o^i)} \right) + 2\alpha_2 \delta \right) \nabla_{\tilde{\phi}_{sh}^i} Q_{\tilde{\phi}_{sh}^i}^{\pi, i}(\mathbf{o}, \mathbf{a}) \right] + \\ &\quad \mathbf{E}_{s \sim d_0, \mathbf{o} \sim \mathcal{E}, \mathbf{a} \sim \pi} \left[\left(\alpha_1 + 2\alpha_2 \delta \left(\frac{\pi_0^i(a^i | o^i)}{\pi_{\psi_{in}^i}^i(a^i | o^i)} \right) \right) \nabla_{\tilde{\phi}_{sh}^i} Q_{\tilde{\phi}_{sh}^i}^{\pi, i}(\mathbf{o}, \mathbf{a}) \right] \\ &\simeq \mathbf{E}_{s \sim d_0, \mathbf{o} \sim \mathcal{E}, \mathbf{a} \sim \pi_0} \left[\left(\alpha_1 \text{CIM}_\epsilon(\pi_{\psi_{in}^i}^i; \pi_0^i) + 2\alpha_2 \delta \right) \nabla_{\tilde{\phi}_{sh}^i} Q_{\tilde{\phi}_{sh}^i}^{\pi, i}(\mathbf{o}, \mathbf{a}) \right] + \\ &\quad \mathbf{E}_{s \sim d_0, \mathbf{o} \sim \mathcal{E}, \mathbf{a} \sim \pi} \left[\left(\alpha_1 + 2\alpha_2 \delta \text{CIM}_\epsilon(\pi_0^i; \pi_{\psi_{in}^i}^i) \right) \nabla_{\tilde{\phi}_{sh}^i} Q_{\tilde{\phi}_{sh}^i}^{\pi, i}(\mathbf{o}, \mathbf{a}) \right]. \end{aligned}$$

■

Appendix E. Extend Proposition 4 to Other Off-policy Algorithm

E.1 F2A2-TD3: Extend Proposition 4 to TD3

The Twin-Delayed Deep Deterministic Policy Gradient algorithm is similar to the DDPG algorithm, just adding a twin Q -value function to a stable training process except for some tricks for implementation. Formally, we can extend it to the fully decentralized multi-agent scenario use the variant of Proposition 4.

Proposition 5 (Off-Policy TD3-Based Joint Gradient) *We set π_0 the data collection policy sampled from experience replay buffer, d_0 represents the distribution of the state-occupancy measure of policy π_0 and δ the TD(0)-error. So the gradient*

of $J_{ac}^i(\mathbf{w}_{in}, \tilde{\mathbf{w}}_{sh}^i)$ is:

$$\begin{aligned}
 & \nabla_{\psi_{in}^i} J_{ac}^i(\mathbf{w}_{in}, \tilde{\mathbf{w}}_{sh}^i) \\
 &= \mathbf{E}_{s \sim d_0, \mathbf{o} \sim \mathcal{E}, \mathbf{a} \sim \pi} \left[\nabla_{\psi_{in}^i} \pi_{\psi_{in}^i}^i(a^i | o^i) \left((\alpha_1 + 2\alpha_2 \delta_1 \left(\frac{\pi_0^i(a^i | o^i)}{\pi_{\psi_{in}^i}^i(a^i | o^i)} \right)) \nabla_{a^i} Q_{\tilde{\phi}_{1,sh}^i}^{\pi,i}(\mathbf{o}, \mathbf{a}) + \right. \right. \\
 & \quad \left. \left. 2\alpha_2 \delta_2 \left(\frac{\pi_0^i(a^i | o^i)}{\pi_{\psi_{in}^i}^i(a^i | o^i)} \right) \nabla_{a^i} Q_{\tilde{\phi}_{2,sh}^i}^{\pi,i}(\mathbf{o}, \mathbf{a}) \right) \right]. \\
 & \nabla_{\tilde{\phi}_{1,sh}^i} J_{ac}^i(\mathbf{w}_{in}, \tilde{\mathbf{w}}_{sh}^i) \\
 &= \mathbf{E}_{s \sim d_0, \mathbf{o} \sim \mathcal{E}, \mathbf{a} \sim \pi_0} \left[\left(\alpha_1 \left(\frac{\pi_{\psi_{in}^i}^i(a^i | o^i)}{\pi_0^i(a^i | o^i)} \right) + 2\alpha_2 \delta_1 \right) \nabla_{\tilde{\phi}_{1,sh}^i} Q_{\tilde{\phi}_{1,sh}^i}^{\pi,i}(\mathbf{o}, \mathbf{a}) \right] + \\
 & \quad \mathbf{E}_{s \sim d_0, \mathbf{o} \sim \mathcal{E}, \mathbf{a} \sim \pi} \left[\left(\alpha_1 + 2\alpha_2 \delta_1 \left(\frac{\pi_0^i(a^i | o^i)}{\pi_{\psi_{in}^i}^i(a^i | o^i)} \right) \right) \nabla_{\tilde{\phi}_{1,sh}^i} Q_{\tilde{\phi}_{1,sh}^i}^{\pi,i}(\mathbf{o}, \mathbf{a}) \right] \\
 & \nabla_{\tilde{\phi}_{2,sh}^i} J_{ac}^i(\mathbf{w}_{in}, \tilde{\mathbf{w}}_{sh}^i) \\
 &= \mathbf{E}_{s \sim d_0, \mathbf{o} \sim \mathcal{E}, \mathbf{a} \sim \pi_0} \left[2\alpha_2 \delta_2 \nabla_{\tilde{\phi}_{2,sh}^i} Q_{\tilde{\phi}_{2,sh}^i}^{\pi,i}(\mathbf{o}, \mathbf{a}) \right] + \\
 & \quad \mathbf{E}_{s \sim d_0, \mathbf{o} \sim \mathcal{E}, \mathbf{a} \sim \pi} \left[2\alpha_2 \delta_2 \left(\frac{\pi_0^i(a^i | o^i)}{\pi_{\psi_{in}^i}^i(a^i | o^i)} \right) \nabla_{\tilde{\phi}_{2,sh}^i} Q_{\tilde{\phi}_{2,sh}^i}^{\pi,i}(\mathbf{o}, \mathbf{a}) \right].
 \end{aligned}$$

Proof We proof the first equation first. Extend the TD3 algorithm, we have

$$\begin{aligned}
 J_{\text{actor}}^i(\mathbf{w}_{in}, \tilde{\mathbf{w}}_{sh}^i) &= \mathbf{E}_{s \sim d_0, \mathbf{o} \sim \mathcal{E}, \mathbf{a} \sim \pi} \left[Q_{\tilde{\phi}_{1,sh}^i}^{\pi,i}(\mathbf{o}, \mathbf{a}) \right] \\
 &= \sum_s d_0(s) \sum_{\mathbf{o}} \mathcal{E}(\mathbf{o} | s) \sum_{\mathbf{a}} \pi_{\Psi_{in}}(\mathbf{a} | \mathbf{o}) \left[Q_{\tilde{\phi}_{1,sh}^i}^{\pi,i}(\mathbf{o}, \mathbf{a}) \right] \\
 J_{\text{critic}}^i(\mathbf{w}_{in}, \tilde{\mathbf{w}}_{sh}^i) &= \mathbf{E}_{s \sim d_0, \mathbf{o} \sim \mathcal{E}, \mathbf{a} \sim \pi_0} \left[\left(Q_{\tilde{\phi}_{1,sh}^i}^{\pi,i}(\mathbf{o}, \mathbf{a}) - Q_{tg}^{\pi} \right)^2 + \left(Q_{\tilde{\phi}_{2,sh}^i}^{\pi,i}(\mathbf{o}, \mathbf{a}) - Q_{tg}^{\pi} \right)^2 \right] \\
 &= \sum_s d_0(s) \sum_{\mathbf{o}} \mathcal{E}_{0,s}^{\mathbf{o}} \sum_{\mathbf{a}} \pi_0(\mathbf{a} | \mathbf{o}) \left[\left(Q_{\tilde{\phi}_{1,sh}^i}^{\pi,i}(\mathbf{o}, \mathbf{a}) - Q_{tg}^{\pi} \right)^2 + \left(Q_{\tilde{\phi}_{2,sh}^i}^{\pi,i}(\mathbf{o}, \mathbf{a}) - Q_{tg}^{\pi} \right)^2 \right],
 \end{aligned}$$

where $Q_{tg}^{\pi} = \sum_{s'} \mathcal{P}_{s,a}^{s'} \left(\mathcal{C}_{s,a}^{i,s'} + \gamma \sum_{\mathbf{o}'} \mathcal{E}_{0,s'}^{\mathbf{o}'} \sum_{\mathbf{a}'} \pi_{\Psi_{in}}(\mathbf{a}' | \mathbf{o}') \left(\min_{j=1,2} Q_{\tilde{\phi}_{j,sh}^i}^{\pi,i}(\mathbf{o}', \mathbf{a}') \right) \right)$. Note that the joint policy $\pi_{\Psi_{in}}$ is a deterministic policy.

We hypothesis joint policy $\pi_{\Psi_{in}}$ is the product of local policy functions $\prod_{i=1}^n \pi_{\psi_{in}^i}$. Hence the gradient concerning each parameter ψ_{in}^i becomes

$$\begin{aligned}
 \nabla_{\psi_{in}^i} J_{\text{actor}}^i(\mathbf{w}_{in}, \tilde{\mathbf{w}}_{sh}^i) &= \nabla_{\psi_{in}^i} \sum_s d_0(s) \sum_{\mathbf{o}} \mathcal{E}(\mathbf{o} | s) \sum_{\mathbf{a}} \pi_{\Psi_{in}}(\mathbf{a} | \mathbf{o}) \left[Q_{\tilde{\phi}_{1,sh}^i}^{\pi,i}(\mathbf{o}, \mathbf{a}) \right] \\
 &= \sum_s d_0(s) \sum_{\mathbf{o}} \mathcal{E}(\mathbf{o} | s) \sum_{\mathbf{a}} \pi_{\Psi_{in}}(\mathbf{a} | \mathbf{o}) \left[\nabla_{\psi_{in}^i} Q_{\tilde{\phi}_{1,sh}^i}^{\pi,i}(\mathbf{o}, \mathbf{a}) \right]
 \end{aligned}$$

$$\begin{aligned}
 &= \sum_s d_0(s) \sum_{\mathbf{o}} \mathcal{E}(\mathbf{o}|s) \sum_{\mathbf{a}} \pi_{\psi_{in}}(\mathbf{a}|\mathbf{o}) \left[\nabla_{\psi_{in}^i} \pi_{\psi_{in}^i}^i(a^i|o^i) \nabla_{\mathbf{a}^i} Q_{\tilde{\phi}_{1,sh}^i}^{\pi,i}(\mathbf{o}, \mathbf{a}) \right] \\
 &= \mathbf{E}_{s \sim d_0, \mathbf{o} \sim \mathcal{E}, \mathbf{a} \sim \pi} \left[\nabla_{\psi_{in}^i} \pi_{\psi_{in}^i}^i(a^i|o^i) \nabla_{\mathbf{a}^i} Q_{\tilde{\phi}_{1,sh}^i}^{\pi,i}(\mathbf{o}, \mathbf{a}) \right].
 \end{aligned}$$

Note that we need to resample to calculate the unbiased policy gradient so that we cannot use the off-policy data directly. We solve the above problem by importance sampling. **For off-policy data saved in experience replay buffer** we have:

$$\nabla_{\psi_{in}^i} J_{\text{actor}}^i(\mathbf{w}_{in}, \tilde{\mathbf{w}}_{sh}^i) = \mathbf{E}_{s \sim d_0, \mathbf{o} \sim \mathcal{E}, \mathbf{a} \sim \pi_0} \left[\left(\frac{\pi_{\psi_{in}^i}^i(a^i|o^i)}{\pi_0^i(a^i|o^i)} \right) \nabla_{\psi_{in}^i} \pi_{\psi_{in}^i}^i(a^i|o^i) \nabla_{\mathbf{a}^i} Q_{\tilde{\phi}_{1,sh}^i}^{\pi,i}(\mathbf{o}, \mathbf{a}) \right] = 0.$$

Combining **the gradient calculated by data which is resampled by current policy**, we get the joint off-policy policy gradient associates with the actor

$$\nabla_{\psi_{in}^i} J_{\text{actor}}^i(\mathbf{w}_{in}, \tilde{\mathbf{w}}_{sh}^i) = \mathbf{E}_{s \sim d_0, \mathbf{o} \sim \mathcal{E}, \mathbf{a} \sim \pi} \left[\nabla_{\psi_{in}^i} \pi_{\psi_{in}^i}^i(a^i|o^i) \nabla_{\mathbf{a}^i} Q_{\tilde{\phi}_{1,sh}^i}^{\pi,i}(\mathbf{o}, \mathbf{a}) \right].$$

On the contrary, we directly use the off-policy data when calculating the value function gradient so that we cannot use the above-resampled data. **For the critic part**, we use the above trick in reverse. **For off-policy data saved in experience replay buffer** we have

$$\begin{aligned}
 &\nabla_{\psi_{in}^i} J_{\text{critic}}^i(\mathbf{w}_{in}, \tilde{\mathbf{w}}_{sh}^i) \\
 &= \nabla_{\psi_{in}^i} \sum_s d_0(s) \sum_{\mathbf{o}} \mathcal{E}_{0,s}^{\mathbf{o}} \sum_{\mathbf{a}} \pi_0(\mathbf{a}|\mathbf{o}) \left[\left(Q_{\tilde{\phi}_{1,sh}^i}^{\pi,i}(\mathbf{o}, \mathbf{a}) - Q_{tg}^{\pi} \right)^2 + \left(Q_{\tilde{\phi}_{2,sh}^i}^{\pi,i}(\mathbf{o}, \mathbf{a}) - Q_{tg}^{\pi} \right)^2 \right] \\
 &= 0,
 \end{aligned}$$

combined **the gradient calculated by the data which is resampled by current policy**, we get the joint off-policy policy gradient associates with the critic part

$$\begin{aligned}
 &\nabla_{\psi_{in}^i} J_{\text{critic}}^i(\mathbf{w}_{in}, \tilde{\mathbf{w}}_{sh}^i) \\
 &= \mathbf{E}_{s \sim d_0, \mathbf{o} \sim \mathcal{E}, \mathbf{a} \sim \pi} \left[2 \left(\frac{\pi_0^i(a^i|o^i)}{\pi_{\psi_{in}^i}^i(a^i|o^i)} \right) \nabla_{\psi_{in}^i} \pi_{\psi_{in}^i}^i(a^i|o^i) \left(\delta_1 \nabla_{\mathbf{a}^i} Q_{\tilde{\phi}_{1,sh}^i}^{\pi,i}(\mathbf{o}, \mathbf{a}) + \delta_2 \nabla_{\mathbf{a}^i} Q_{\tilde{\phi}_{2,sh}^i}^{\pi,i}(\mathbf{o}, \mathbf{a}) \right) \right].
 \end{aligned}$$

where $\delta_j = Q_{\tilde{\phi}_{j,sh}^i}^{\pi,i}(\mathbf{o}, \mathbf{a}) - Q_{tg}^{\pi}$. We denote the clipped importance sampling term $\min \left(\epsilon, \frac{\pi_{\psi_{in}^i}^i(a^i|o^i)}{\pi_0^i(a^i|o^i)} \right)$

as $CIM_{\epsilon}(\pi_{\psi_{in}^i}^i; \pi_0^i)$ and $\min \left(\epsilon, \frac{\pi_0^i(a^i|o^i)}{\pi_{\psi_{in}^i}^i(a^i|o^i)} \right)$ as $CIM_{\epsilon}(\pi_0^i; \pi_{\psi_{in}^i}^i)$. Finally, we get following off-

policy joint gradient w.r.t. actor parameters:

$$\begin{aligned}
 \nabla_{\psi_{in}^i} J_{ac}^i(\mathbf{w}_{in}, \tilde{\mathbf{w}}_{sh}^i) &= \alpha_1 \nabla_{\psi_{in}^i} J_{actor}^i(\mathbf{w}_{in}, \tilde{\mathbf{w}}_{sh}^i) + \alpha_2 \nabla_{\psi_{in}^i} J_{critic}^i(\mathbf{w}_{in}, \tilde{\mathbf{w}}_{sh}^i) \\
 &= \mathbf{E}_{s \sim d_0, \mathbf{o} \sim \mathcal{E}, \mathbf{a} \sim \pi} \left[\nabla_{\psi_{in}^i} \pi_{\psi_{in}^i}^i(a^i | o^i) \left((\alpha_1 + 2\alpha_2 \delta_1 \left(\frac{\pi_0^i(a^i | o^i)}{\pi_{\psi_{in}^i}^i(a^i | o^i)} \right)) \nabla_{a^i} Q_{\tilde{\phi}_{1,sh}^i}^{\pi, i}(\mathbf{o}, \mathbf{a}) + \right. \right. \\
 &\quad \left. \left. 2\alpha_2 \delta_2 \left(\frac{\pi_0^i(a^i | o^i)}{\pi_{\psi_{in}^i}^i(a^i | o^i)} \right) \nabla_{a^i} Q_{\tilde{\phi}_{2,sh}^i}^{\pi, i}(\mathbf{o}, \mathbf{a}) \right) \right]. \\
 &\simeq \mathbf{E}_{s \sim d_0, \mathbf{o} \sim \mathcal{E}, \mathbf{a} \sim \pi} \left[\nabla_{\psi_{in}^i} \pi_{\psi_{in}^i}^i(a^i | o^i) \left((\alpha_1 + 2\alpha_2 \delta_1 CIM_\epsilon(\pi_0^i; \pi_{\psi_{in}^i}^i)) \nabla_{a^i} Q_{\tilde{\phi}_{1,sh}^i}^{\pi, i}(\mathbf{o}, \mathbf{a}) + \right. \right. \\
 &\quad \left. \left. 2\alpha_2 \delta_2 CIM_\epsilon(\pi_0^i; \pi_{\psi_{in}^i}^i) \nabla_{a^i} Q_{\tilde{\phi}_{2,sh}^i}^{\pi, i}(\mathbf{o}, \mathbf{a}) \right) \right].
 \end{aligned}$$

The above is the proof of the first equation, below we prove the second. For actor part, we first calculate the gradient use the data resampled by the current policy

$$\begin{aligned}
 \nabla_{\tilde{\phi}_{1,sh}^i} J_{actor}^i(\mathbf{w}_{in}, \tilde{\mathbf{w}}_{sh}^i) &= \nabla_{\tilde{\phi}_{1,sh}^i} \sum_s d_0(s) \sum_o \mathcal{E}(\mathbf{o} | s) \sum_a \pi_{\Psi_{in}}(\mathbf{a} | \mathbf{o}) \left[Q_{\tilde{\phi}_{1,sh}^i}^{\pi, i}(\mathbf{o}, \mathbf{a}) \right] \\
 &= \sum_s d_0(s) \sum_o \mathcal{E}(\mathbf{o} | s) \sum_a \pi_{\Psi_{in}}(\mathbf{a} | \mathbf{o}) \nabla_{\tilde{\phi}_{1,sh}^i} Q_{\tilde{\phi}_{1,sh}^i}^{\pi, i}(\mathbf{o}, \mathbf{a}) \\
 &= \mathbf{E}_{s \sim d_0, \mathbf{o} \sim \mathcal{E}, \mathbf{a} \sim \pi} \left[\nabla_{\tilde{\phi}_{1,sh}^i} Q_{\tilde{\phi}_{1,sh}^i}^{\pi, i}(\mathbf{o}, \mathbf{a}) \right].
 \end{aligned}$$

For calculating the gradient use off-policy data we also introduce importance sampling, then we have

$$\nabla_{\tilde{\phi}_{1,sh}^i} J_{actor}^i(\mathbf{w}_{in}, \tilde{\mathbf{w}}_{sh}^i) = \mathbf{E}_{s \sim d_0, \mathbf{o} \sim \mathcal{E}, \mathbf{a} \sim \pi_0} \left[\left(\frac{\pi_{\psi_{in}^i}^i(a^i | o^i)}{\pi_0^i(a^i | o^i)} \right) \nabla_{\tilde{\phi}_{1,sh}^i} Q_{\tilde{\phi}_{1,sh}^i}^{\pi, i}(\mathbf{o}, \mathbf{a}) \right].$$

We then combine the above two part gradient:

$$\begin{aligned}
 \nabla_{\tilde{\phi}_{1,sh}^i} J_{actor}^i(\mathbf{w}_{in}, \tilde{\mathbf{w}}_{sh}^i) &= \mathbf{E}_{s \sim d_0, \mathbf{o} \sim \mathcal{E}, \mathbf{a} \sim \pi} \left[\nabla_{\tilde{\phi}_{1,sh}^i} Q_{\tilde{\phi}_{1,sh}^i}^{\pi, i}(\mathbf{o}, \mathbf{a}) \right] + \\
 &\quad \mathbf{E}_{s \sim d_0, \mathbf{o} \sim \mathcal{E}, \mathbf{a} \sim \pi_0} \left[\left(\frac{\pi_{\psi_{in}^i}^i(a^i | o^i)}{\pi_0^i(a^i | o^i)} \right) \nabla_{\tilde{\phi}_{1,sh}^i} Q_{\tilde{\phi}_{1,sh}^i}^{\pi, i}(\mathbf{o}, \mathbf{a}) \right].
 \end{aligned}$$

For critic part, we first calculate the gradient use the off-policy data

$$\begin{aligned}
 &\nabla_{\tilde{\phi}_{1,sh}^i} J_{critic}^i(\mathbf{w}_{in}, \tilde{\mathbf{w}}_{sh}^i) \\
 &= \nabla_{\tilde{\phi}_{1,sh}^i} \sum_s d_0(s) \sum_o \mathcal{E}_{0,s}^o \sum_a \pi_0(\mathbf{a} | \mathbf{o}) \left(\left(Q_{\tilde{\phi}_{1,sh}^i}^{\pi, i}(\mathbf{o}, \mathbf{a}) - Q_{tg} \right)^2 + \left(Q_{\tilde{\phi}_{2,sh}^i}^{\pi, i}(\mathbf{o}, \mathbf{a}) - Q_{tg} \right)^2 \right) \\
 &= \sum_s d_0(s) \sum_o \mathcal{E}_{0,s}^o \sum_a \pi_0(\mathbf{a} | \mathbf{o}) 2\delta_1 \nabla_{\tilde{\phi}_{1,sh}^i} Q_{\tilde{\phi}_{1,sh}^i}^{\pi, i}(\mathbf{o}, \mathbf{a}) \\
 &= \mathbf{E}_{s \sim d_0, \mathbf{o} \sim \mathcal{E}, \mathbf{a} \sim \pi_0} \left[2\delta_1 \nabla_{\tilde{\phi}_{1,sh}^i} Q_{\tilde{\phi}_{1,sh}^i}^{\pi, i}(\mathbf{o}, \mathbf{a}) \right].
 \end{aligned}$$

Then, for resampled data, we have

$$\nabla_{\tilde{\phi}_{1,sh}^i} J_{\text{critic}}^i(\mathbf{w}_{in}, \tilde{\mathbf{w}}_{sh}^i) = \mathbf{E}_{s \sim d_0, \mathbf{o} \sim \mathcal{E}, \mathbf{a} \sim \pi} \left[2\delta_1 \left(\frac{\pi_0^i(a^i|o^i)}{\pi_{\psi_{in}^i}^i(a^i|o^i)} \right) \nabla_{\tilde{\phi}_{1,sh}^i} Q_{\tilde{\phi}_{1,sh}^i}^{\pi,i}(\mathbf{o}, \mathbf{a}) \right].$$

Combining above two parts' gradient, we have:

$$\begin{aligned} \nabla_{\tilde{\phi}_{1,sh}^i} J_{\text{critic}}^i(\mathbf{w}_{in}, \tilde{\mathbf{w}}_{sh}^i) &= \mathbf{E}_{s \sim d_0, \mathbf{o} \sim \mathcal{E}, \mathbf{a} \sim \pi_0} \left[2\delta_1 \nabla_{\tilde{\phi}_{1,sh}^i} Q_{\tilde{\phi}_{1,sh}^i}^{\pi,i}(\mathbf{o}, \mathbf{a}) \right] + \\ &\quad \mathbf{E}_{s \sim d_0, \mathbf{o} \sim \mathcal{E}, \mathbf{a} \sim \pi} \left[2\delta_1 \left(\frac{\pi_0^i(a^i|o^i)}{\pi_{\psi_{in}^i}^i(a^i|o^i)} \right) \nabla_{\tilde{\phi}_{1,sh}^i} Q_{\tilde{\phi}_{1,sh}^i}^{\pi,i}(\mathbf{o}, \mathbf{a}) \right]. \end{aligned}$$

Finally, we get the following off-policy joint gradient w.r.t. the first critic parameters:

$$\begin{aligned} \nabla_{\tilde{\phi}_{1,sh}^i} J_{ac}^i(\mathbf{w}_{in}, \tilde{\mathbf{w}}_{sh}^i) &= \alpha_1 \nabla_{\tilde{\phi}_{1,sh}^i} J_{actor}^i(\mathbf{w}_{in}, \tilde{\mathbf{w}}_{sh}^i) + \alpha_2 \nabla_{\tilde{\phi}_{1,sh}^i} J_{critic}^i(\mathbf{w}_{in}, \tilde{\mathbf{w}}_{sh}^i) \\ &= \mathbf{E}_{s \sim d_0, \mathbf{o} \sim \mathcal{E}, \mathbf{a} \sim \pi_0} \left[\left(\alpha_1 \left(\frac{\pi_{\psi_{in}^i}^i(a^i|o^i)}{\pi_0^i(a^i|o^i)} \right) + 2\alpha_2 \delta_1 \right) \nabla_{\tilde{\phi}_{1,sh}^i} Q_{\tilde{\phi}_{1,sh}^i}^{\pi,i}(\mathbf{o}, \mathbf{a}) \right] + \\ &\quad \mathbf{E}_{s \sim d_0, \mathbf{o} \sim \mathcal{E}, \mathbf{a} \sim \pi} \left[\left(\alpha_1 + 2\alpha_2 \delta_1 \left(\frac{\pi_0^i(a^i|o^i)}{\pi_{\psi_{in}^i}^i(a^i|o^i)} \right) \right) \nabla_{\tilde{\phi}_{1,sh}^i} Q_{\tilde{\phi}_{1,sh}^i}^{\pi,i}(\mathbf{o}, \mathbf{a}) \right] \\ &\simeq \mathbf{E}_{s \sim d_0, \mathbf{o} \sim \mathcal{E}, \mathbf{a} \sim \pi_0} \left[\left(\alpha_1 \text{CIM}_\epsilon(\pi_{\psi_{in}^i}^i; \pi_0^i) + 2\alpha_2 \delta_1 \right) \nabla_{\tilde{\phi}_{1,sh}^i} Q_{\tilde{\phi}_{1,sh}^i}^{\pi,i}(\mathbf{o}, \mathbf{a}) \right] + \\ &\quad \mathbf{E}_{s \sim d_0, \mathbf{o} \sim \mathcal{E}, \mathbf{a} \sim \pi} \left[\left(\alpha_1 + 2\alpha_2 \delta_1 \text{CIM}_\epsilon(\pi_0^i; \pi_{\psi_{in}^i}^i) \right) \nabla_{\tilde{\phi}_{1,sh}^i} Q_{\tilde{\phi}_{1,sh}^i}^{\pi,i}(\mathbf{o}, \mathbf{a}) \right]. \end{aligned}$$

For the third equation and use resampled data to calculate the actor part gradient, we have

$$\nabla_{\tilde{\phi}_{2,sh}^i} J_{actor}^i(\mathbf{w}_{in}, \tilde{\mathbf{w}}_{sh}^i) = \nabla_{\tilde{\phi}_{2,sh}^i} \sum_s d_0(s) \sum_{\mathbf{o}} \mathcal{E}(\mathbf{o}|s) \sum_{\mathbf{a}} \pi_{\Psi_{in}}(\mathbf{a}|\mathbf{o}) \left[Q_{\tilde{\phi}_{1,sh}^i}^{\pi,i}(\mathbf{o}, \mathbf{a}) \right] = 0.$$

Similarly, for off-policy data the gradient also is 0. For critic part the gradient is similar as the second equation

$$\begin{aligned} \nabla_{\tilde{\phi}_{2,sh}^i} J_{critic}^i(\mathbf{w}_{in}, \tilde{\mathbf{w}}_{sh}^i) &= \mathbf{E}_{s \sim d_0, \mathbf{o} \sim \mathcal{E}, \mathbf{a} \sim \pi_0} \left[2\delta_2 \nabla_{\tilde{\phi}_{2,sh}^i} Q_{\tilde{\phi}_{2,sh}^i}^{\pi,i}(\mathbf{o}, \mathbf{a}) \right] + \\ &\quad \mathbf{E}_{s \sim d_0, \mathbf{o} \sim \mathcal{E}, \mathbf{a} \sim \pi} \left[2\delta_2 \left(\frac{\pi_0^i(a^i|o^i)}{\pi_{\psi_{in}^i}^i(a^i|o^i)} \right) \nabla_{\tilde{\phi}_{2,sh}^i} Q_{\tilde{\phi}_{2,sh}^i}^{\pi,i}(\mathbf{o}, \mathbf{a}) \right]. \end{aligned}$$

Finally, we get following off-policy joint gradient w.r.t. the second critic parameters:

$$\begin{aligned} \nabla_{\tilde{\phi}_{2,sh}^i} J_{ac}^i(\mathbf{w}_{in}, \tilde{\mathbf{w}}_{sh}^i) &= \alpha_1 \nabla_{\tilde{\phi}_{2,sh}^i} J_{actor}^i(\mathbf{w}_{in}, \tilde{\mathbf{w}}_{sh}^i) + \alpha_2 \nabla_{\tilde{\phi}_{2,sh}^i} J_{critic}^i(\mathbf{w}_{in}, \tilde{\mathbf{w}}_{sh}^i) \\ &= \mathbf{E}_{s \sim d_0, \mathbf{o} \sim \mathcal{E}, \mathbf{a} \sim \pi_0} \left[2\alpha_2 \delta_2 \nabla_{\tilde{\phi}_{2,sh}^i} Q_{\tilde{\phi}_{2,sh}^i}^{\pi,i}(\mathbf{o}, \mathbf{a}) \right] + \\ &\quad \mathbf{E}_{s \sim d_0, \mathbf{o} \sim \mathcal{E}, \mathbf{a} \sim \pi} \left[2\alpha_2 \delta_2 \left(\frac{\pi_0^i(a^i|o^i)}{\pi_{\psi_{in}^i}^i(a^i|o^i)} \right) \nabla_{\tilde{\phi}_{2,sh}^i} Q_{\tilde{\phi}_{2,sh}^i}^{\pi,i}(\mathbf{o}, \mathbf{a}) \right]. \end{aligned}$$

■

E.2 F2A2-SAC: Extend Proposition 4 to SAC

Soft Actor Critic (SAC) is an algorithm that optimizes a stochastic policy in an off-policy way. A central feature of SAC is entropy regularization. The policy is trained to maximize a trade-off between expected return and entropy, a measure of randomness in the policy. We only used one Q -value function and omitted the estimation of the state-value(V) function. Formally, we can extend it to the fully decentralized multi-agent scenario use the variant of Proposition 4:

Proposition 6 (Off-Policy SAC-Based Joint Gradient) *We set π_0 the data collection policy sampled from experience replay buffer, d_0 represents the distribution of the state-occupancy measure of policy π_0 and δ the TD(0)-error. So the gradient of $J_{ac}^i(\mathbf{w}_{in}, \tilde{\mathbf{w}}_{sh}^i)$ is:*

$$\begin{aligned}
 & \nabla_{\psi_{in}^i} J_{ac}^i(\mathbf{w}_{in}, \tilde{\mathbf{w}}_{sh}^i) \\
 = & \mathbf{E}_{s \sim d_0, \mathbf{o} \sim \mathcal{E}, \mathbf{a} \sim \pi_0} \left[\alpha_1 \left(\frac{\pi_{\psi_{in}^i}^i(a^i | \mathbf{o}^i)}{\pi_0^i(a^i | \mathbf{o}^i)} \right) \nabla_{\psi_{in}^i} \log \pi_{\psi_{in}^i}^i(a^i | \mathbf{o}^i) \right. \\
 & \quad \left. \left(Q_{\tilde{\phi}_{1,sh}^i}^{\pi, i}(\mathbf{o}, \mathbf{a}) - \alpha \log \pi_{\psi_{in}^i}^i(a^i | \mathbf{o}^i) - \mathcal{B}(\mathbf{o}, \mathbf{a} \setminus^i) \right) \right] + \\
 & \mathbf{E}_{s \sim d_0, \mathbf{o} \sim \mathcal{E}, \mathbf{a} \sim \pi} \left[\alpha_1 \nabla_{\psi_{in}^i} \log \pi_{\psi_{in}^i}^i(a^i | \mathbf{o}^i) \right. \\
 & \quad \left. \left(Q_{\tilde{\phi}_{1,sh}^i}^{\pi, i}(\mathbf{o}, \mathbf{a}) - \alpha \log \pi_{\psi_{in}^i}^i(a^i | \mathbf{o}^i) - \mathcal{B}(\mathbf{o}, \mathbf{a} \setminus^i) \right) \right] + \\
 & \mathbf{E}_{s \sim d_0, \mathbf{o} \sim \mathcal{E}, \mathbf{a} \sim \pi} \left[\alpha_2 \left(\frac{\pi_0^i(a^i | \mathbf{o}^i)}{\pi_{\psi_{in}^i}^i(a^i | \mathbf{o}^i)} \right) \nabla_{\psi_{in}^i} \log \pi_{\psi_{in}^i}^i(a^i | \mathbf{o}^i) (\delta_1^2 + \delta_2^2) \right]. \\
 & \nabla_{\tilde{\phi}_{1,sh}^i} J_{ac}^i(\mathbf{w}_{in}, \tilde{\mathbf{w}}_{sh}^i) \\
 = & \mathbf{E}_{s \sim d_0, \mathbf{o} \sim \mathcal{E}, \mathbf{a} \sim \pi_0} \left[\left(\alpha_1 \left(\frac{\pi_{\psi_{in}^i}^i(a^i | \mathbf{o}^i)}{\pi_0^i(a^i | \mathbf{o}^i)} \right) + 2\alpha_2\delta_1 \right) \nabla_{\tilde{\phi}_{1,sh}^i} Q_{\tilde{\phi}_{1,sh}^i}^{\pi, i}(\mathbf{o}, \mathbf{a}) \right] + \\
 & \mathbf{E}_{s \sim d_0, \mathbf{o} \sim \mathcal{E}, \mathbf{a} \sim \pi} \left[\left(\alpha_1 + 2\alpha_2\delta_1 \left(\frac{\pi_0^i(a^i | \mathbf{o}^i)}{\pi_{\psi_{in}^i}^i(a^i | \mathbf{o}^i)} \right) \right) \nabla_{\tilde{\phi}_{1,sh}^i} Q_{\tilde{\phi}_{1,sh}^i}^{\pi, i}(\mathbf{o}, \mathbf{a}) \right]. \\
 & \nabla_{\tilde{\phi}_{2,sh}^i} J_{ac}^i(\mathbf{w}_{in}, \tilde{\mathbf{w}}_{sh}^i) = \mathbf{E}_{s \sim d_0, \mathbf{o} \sim \mathcal{E}, \mathbf{a} \sim \pi_0} \left[2\alpha_2\delta_2 \nabla_{\tilde{\phi}_{2,sh}^i} Q_{\tilde{\phi}_{2,sh}^i}^{\pi, i}(\mathbf{o}, \mathbf{a}) \right] + \\
 & \mathbf{E}_{s \sim d_0, \mathbf{o} \sim \mathcal{E}, \mathbf{a} \sim \pi} \left[2\alpha_2\delta_2 \left(\frac{\pi_0^i(a^i | \mathbf{o}^i)}{\pi_{\psi_{in}^i}^i(a^i | \mathbf{o}^i)} \right) \nabla_{\tilde{\phi}_{2,sh}^i} Q_{\tilde{\phi}_{2,sh}^i}^{\pi, i}(\mathbf{o}, \mathbf{a}) \right].
 \end{aligned}$$

Proof We proof the first equation first. Extend the SAC algorithm, we have (for convenience here we suppose state space, observation space and action space are discrete)

$$\begin{aligned}
 J_{\text{actor}}^i(\mathbf{w}_{in}, \tilde{\mathbf{w}}_{sh}^i) &= \mathbf{E}_{s \sim d_0, \mathbf{o} \sim \mathcal{E}, \mathbf{a} \sim \pi} \left[Q_{\tilde{\phi}_{1,sh}^i}^{\pi, i}(\mathbf{o}, \mathbf{a}) - \alpha \log \pi_{\psi_{in}^i}^i(a^i | o^i) - \mathcal{B}(\mathbf{o}, \mathbf{a}^i) \right] \\
 &= \sum_s d_0(s) \sum_{\mathbf{o}} \mathcal{E}(\mathbf{o} | s) \sum_{\mathbf{a}} \pi_{\Psi_{in}}(\mathbf{a} | \mathbf{o}) \left[Q_{\tilde{\phi}_{1,sh}^i}^{\pi, i}(\mathbf{o}, \mathbf{a}) - \alpha \log \pi_{\psi_{in}^i}^i(a^i | o^i) - \mathcal{B}(\mathbf{o}, \mathbf{a}^i) \right] \\
 J_{\text{critic}}^i(\mathbf{w}_{in}, \tilde{\mathbf{w}}_{sh}^i) &= \mathbf{E}_{s \sim d_0, \mathbf{o} \sim \mathcal{E}, \mathbf{a} \sim \pi_0} \left[\left(Q_{\tilde{\phi}_{1,sh}^i}^{\pi, i}(\mathbf{o}, \mathbf{a}) - Q_{tg}^{\pi} \right)^2 + \left(Q_{\tilde{\phi}_{2,sh}^i}^{\pi, i}(\mathbf{o}, \mathbf{a}) - Q_{tg}^{\pi} \right)^2 \right] \\
 &= \sum_s d_0(s) \sum_{\mathbf{o}} \mathcal{E}_{0,s}^{\mathbf{o}} \sum_{\mathbf{a}} \pi_0(\mathbf{a} | \mathbf{o}) \left[\left(Q_{\tilde{\phi}_{1,sh}^i}^{\pi, i}(\mathbf{o}, \mathbf{a}) - Q_{tg}^{\pi} \right)^2 + \left(Q_{\tilde{\phi}_{2,sh}^i}^{\pi, i}(\mathbf{o}, \mathbf{a}) - Q_{tg}^{\pi} \right)^2 \right],
 \end{aligned}$$

where

$$Q_{tg}^{\pi} = \sum_{s'} \mathcal{P}_{s,a}^{s'} \left(\mathcal{C}_{s,a}^{i,s'} + \gamma \sum_{\mathbf{o}'} \mathcal{E}_{0,s'}^{\mathbf{o}'} \sum_{\mathbf{a}'} \pi_{\Psi_{in}}(\mathbf{a}' | \mathbf{o}') \left(\min_{j=1,2} Q_{\tilde{\phi}_{j,sh}^i}^{\pi, i}(\mathbf{o}', \mathbf{a}') - \alpha \log \pi_{\psi_{in}^i}^i(a'^i | o'^i) \right) \right).$$

We assume joint policy $\pi_{\Psi_{in}}$ is the product of local policy functions $\prod_{i=1}^n \pi_{\psi_{in}^i}^i$. Hence **the actor part** gradient concerning each parameter ψ_{in}^i becomes:

$$\begin{aligned}
 &\nabla_{\psi_{in}^i} J_{\text{actor}}^i(\mathbf{w}_{in}, \tilde{\mathbf{w}}_{sh}^i) \\
 &= \nabla_{\psi_{in}^i} \sum_s d_0(s) \sum_{\mathbf{o}} \mathcal{E}(\mathbf{o} | s) \sum_{\mathbf{a}} \pi_{\Psi_{in}}(\mathbf{a} | \mathbf{o}) \left[Q_{\tilde{\phi}_{1,sh}^i}^{\pi, i}(\mathbf{o}, \mathbf{a}) - \alpha \log \pi_{\psi_{in}^i}^i(a^i | o^i) - \mathcal{B}(\mathbf{o}, \mathbf{a}^i) \right] \\
 &= \sum_s d_0(s) \sum_{\mathbf{o}} \mathcal{E}(\mathbf{o} | s) \sum_{\mathbf{a}} \nabla_{\psi_{in}^i} \pi_{\Psi_{in}}(\mathbf{a} | \mathbf{o}) \left[Q_{\tilde{\phi}_{1,sh}^i}^{\pi, i}(\mathbf{o}, \mathbf{a}) - \alpha \log \pi_{\psi_{in}^i}^i(a^i | o^i) - \mathcal{B}(\mathbf{o}, \mathbf{a}^i) \right] \\
 &= \mathbf{E}_{s \sim d_0, \mathbf{o} \sim \mathcal{E}, \mathbf{a} \sim \pi} \left[\nabla_{\psi_{in}^i} \log \pi_{\psi_{in}^i}^i(a^i | o^i) \left(Q_{\tilde{\phi}_{1,sh}^i}^{\pi, i}(\mathbf{o}, \mathbf{a}) - \alpha \log \pi_{\psi_{in}^i}^i(a^i | o^i) - \mathcal{B}(\mathbf{o}, \mathbf{a}^i) \right) \right].
 \end{aligned}$$

Note that we need to resample to calculate the unbiased policy gradient so that we cannot use the off-policy data directly. We solve the above problem by importance sampling. **For off-policy data saved in experience replay buffer** we have

$$\begin{aligned}
 \nabla_{\psi_{in}^i} J_{\text{actor}}^i(\mathbf{w}_{in}, \tilde{\mathbf{w}}_{sh}^i) &= \mathbf{E}_{s \sim d_0, \mathbf{o} \sim \mathcal{E}, \mathbf{a} \sim \pi_0} \left[\left(\frac{\pi_{\psi_{in}^i}^i(a^i | o^i)}{\pi_0^i(a^i | o^i)} \right) \nabla_{\psi_{in}^i} \log \pi_{\psi_{in}^i}^i(a^i | o^i) \right. \\
 &\quad \left. \left(Q_{\tilde{\phi}_{1,sh}^i}^{\pi, i}(\mathbf{o}, \mathbf{a}) - \alpha \log \pi_{\psi_{in}^i}^i(a^i | o^i) - \mathcal{B}(\mathbf{o}, \mathbf{a}^i) \right) \right],
 \end{aligned}$$

combined **the gradient calculated by the data which is resampled by current policy**, we get the joint off-policy policy gradient associates with the actor part:

$$\begin{aligned} \nabla_{\psi_{in}^i} J_{actor}^i(\mathbf{w}_{in}, \tilde{\mathbf{w}}_{sh}^i) &= \mathbf{E}_{s \sim d_0, \mathcal{O} \sim \mathcal{E}, \mathbf{a} \sim \pi_0} \left[\left(\frac{\pi_{\psi_{in}^i}^i(a^i | o^i)}{\pi_0^i(a^i | o^i)} \right) \nabla_{\psi_{in}^i} \log \pi_{\psi_{in}^i}^i(a^i | o^i) \right. \\ &\quad \left. \left(Q_{\tilde{\phi}_{1,sh}^i}^{\pi, i}(\mathbf{o}, \mathbf{a}) - \alpha \log \pi_{\psi_{in}^i}^i(a^i | o^i) - \mathcal{B}(\mathbf{o}, \mathbf{a}^i) \right) \right] + \\ &\quad \mathbf{E}_{s \sim d_0, \mathcal{O} \sim \mathcal{E}, \mathbf{a} \sim \pi} \left[\nabla_{\psi_{in}^i} \log \pi_{\psi_{in}^i}^i(a^i | o^i) \left(Q_{\tilde{\phi}_{1,sh}^i}^{\pi, i}(\mathbf{o}, \mathbf{a}) - \alpha \log \pi_{\psi_{in}^i}^i(a^i | o^i) - \mathcal{B}(\mathbf{o}, \mathbf{a}^i) \right) \right]. \end{aligned}$$

On the contrary, we directly use the off-policy data when calculating the value function gradient so that we cannot use the above-resampled data. **For the critic part**, we use the above trick in reverse. **For off-policy data saved in experience replay buffer** we have

$$\begin{aligned} &\nabla_{\psi_{in}^i} J_{critic}^i(\mathbf{w}_{in}, \tilde{\mathbf{w}}_{sh}^i) \\ &= \nabla_{\psi_{in}^i} \sum_s d_0(s) \sum_{\mathcal{O}} \mathcal{E}_{0,s}^{\mathcal{O}} \sum_{\mathbf{a}} \pi_0(\mathbf{a} | \mathcal{O}) \left(\left(Q_{\tilde{\phi}_{1,sh}^i}^{\pi, i}(\mathbf{o}, \mathbf{a}) - Q_{tg}^{\pi} \right)^2 + \left(Q_{\tilde{\phi}_{2,sh}^i}^{\pi, i}(\mathbf{o}, \mathbf{a}) - Q_{tg}^{\pi} \right)^2 \right) \\ &= 0, \end{aligned}$$

combined **the gradient calculated by the data which is resampled by current policy**, we get the joint off-policy policy gradient associates with the critic part

$$\nabla_{\psi_{in}^i} J_{critic}^i(\mathbf{w}_{in}, \tilde{\mathbf{w}}_{sh}^i) = \mathbf{E}_{s \sim d_0, \mathcal{O} \sim \mathcal{E}, \mathbf{a} \sim \pi} \left[\left(\frac{\pi_0^i(a^i | o^i)}{\pi_{\psi_{in}^i}^i(a^i | o^i)} \right) \nabla_{\psi_{in}^i} \log \pi_{\psi_{in}^i}^i(a^i | o^i) (\delta_1^2 + \delta_2^2) \right].$$

where $\delta_j = Q_{\tilde{\phi}_{j,sh}^i}^{\pi, i}(\mathbf{o}, \mathbf{a}) - Q_{tg}^{\pi}$. We denote the clipped importance sampling term $\min \left(\epsilon, \frac{\pi_{\psi_{in}^i}^i(a^i | o^i)}{\pi_0^i(a^i | o^i)} \right)$

as $CIM_{\epsilon}(\pi_{\psi_{in}^i}^i; \pi_0^i)$ and $\min \left(\epsilon, \frac{\pi_0^i(a^i | o^i)}{\pi_{\psi_{in}^i}^i(a^i | o^i)} \right)$ as $CIM_{\epsilon}(\pi_0^i; \pi_{\psi_{in}^i}^i)$. Finally, we get the following off-policy joint gradient w.r.t. actor parameters:

$$\begin{aligned} &\nabla_{\psi_{in}^i} J_{ac}^i(\mathbf{w}_{in}, \tilde{\mathbf{w}}_{sh}^i) \\ &= \alpha_1 \nabla_{\psi_{in}^i} J_{actor}^i(\mathbf{w}_{in}, \tilde{\mathbf{w}}_{sh}^i) + \alpha_2 \nabla_{\psi_{in}^i} J_{critic}^i(\mathbf{w}_{in}, \tilde{\mathbf{w}}_{sh}^i) \\ &= \mathbf{E}_{s \sim d_0, \mathcal{O} \sim \mathcal{E}, \mathbf{a} \sim \pi_0} \left[\alpha_1 \left(\frac{\pi_{\psi_{in}^i}^i(a^i | o^i)}{\pi_0^i(a^i | o^i)} \right) \nabla_{\psi_{in}^i} \log \pi_{\psi_{in}^i}^i(a^i | o^i) \right. \\ &\quad \left. \left(Q_{\tilde{\phi}_{1,sh}^i}^{\pi, i}(\mathbf{o}, \mathbf{a}) - \alpha \log \pi_{\psi_{in}^i}^i(a^i | o^i) - \mathcal{B}(\mathbf{o}, \mathbf{a}^i) \right) \right] + \\ &\quad \mathbf{E}_{s \sim d_0, \mathcal{O} \sim \mathcal{E}, \mathbf{a} \sim \pi} \left[\alpha_1 \nabla_{\psi_{in}^i} \log \pi_{\psi_{in}^i}^i(a^i | o^i) \left(Q_{\tilde{\phi}_{1,sh}^i}^{\pi, i}(\mathbf{o}, \mathbf{a}) - \alpha \log \pi_{\psi_{in}^i}^i(a^i | o^i) - \mathcal{B}(\mathbf{o}, \mathbf{a}^i) \right) \right] + \end{aligned}$$

$$\begin{aligned}
 & \mathbf{E}_{s \sim d_0, \mathbf{o} \sim \mathcal{E}, \mathbf{a} \sim \pi} \left[\alpha_2 \left(\frac{\pi_0^i(a^i | o^i)}{\pi_{\psi_{in}^i}^i(a^i | o^i)} \right) \nabla_{\psi_{in}^i} \log \pi_{\psi_{in}^i}^i(a^i | o^i) (\delta_1^2 + \delta_2^2) \right] \\
 \simeq & \mathbf{E}_{s \sim d_0, \mathbf{o} \sim \mathcal{E}, \mathbf{a} \sim \pi_0} \left[\alpha_1 \text{CIM}_\epsilon(\pi_{\psi_{in}^i}^i; \pi_0^i) \nabla_{\psi_{in}^i} \log \pi_{\psi_{in}^i}^i(a^i | o^i) \right. \\
 & \left. \left(Q_{\tilde{\phi}_{1,sh}^i}^{\pi, i}(\mathbf{o}, \mathbf{a}) - \alpha \log \pi_{\psi_{in}^i}^i(a^i | o^i) - \mathcal{B}(\mathbf{o}, \mathbf{a}^i) \right) \right] + \\
 & \mathbf{E}_{s \sim d_0, \mathbf{o} \sim \mathcal{E}, \mathbf{a} \sim \pi} \left[\alpha_1 \nabla_{\psi_{in}^i} \log \pi_{\psi_{in}^i}^i(a^i | o^i) \left(Q_{\tilde{\phi}_{1,sh}^i}^{\pi, i}(\mathbf{o}, \mathbf{a}) - \alpha \log \pi_{\psi_{in}^i}^i(a^i | o^i) - \mathcal{B}(\mathbf{o}, \mathbf{a}^i) \right) \right] + \\
 & \mathbf{E}_{s \sim d_0, \mathbf{o} \sim \mathcal{E}, \mathbf{a} \sim \pi} \left[\alpha_2 \text{CIM}_\epsilon(\pi_0^i; \pi_{\psi_{in}^i}^i) \nabla_{\psi_{in}^i} \log \pi_{\psi_{in}^i}^i(a^i | o^i) (\delta_1^2 + \delta_2^2) \right].
 \end{aligned}$$

The above is the proof of the first equation, **below we prove the second. For actor part, we first calculate the gradient use the data resampled by the current policy**

$$\begin{aligned}
 \nabla_{\tilde{\phi}_{1,sh}^i} J_{\text{actor}}^i(\mathbf{w}_{in}, \tilde{\mathbf{w}}_{sh}^i) &= \nabla_{\tilde{\phi}_{1,sh}^i} \sum_s d_0(s) \sum_{\mathbf{o}} \mathcal{E}(\mathbf{o} | s) \sum_{\mathbf{a}} \pi_{\Psi_{in}}(\mathbf{a} | \mathbf{o}) \\
 & \quad \left[Q_{\tilde{\phi}_{1,sh}^i}^{\pi, i}(\mathbf{o}, \mathbf{a}) - \alpha \log \pi_{\psi_{in}^i}^i(a^i | o^i) - \mathcal{B}(\mathbf{o}, \mathbf{a}^i) \right] \\
 &= \sum_s d_0(s) \sum_{\mathbf{o}} \mathcal{E}(\mathbf{o} | s) \sum_{\mathbf{a}} \pi_{\Psi_{in}}(\mathbf{a} | \mathbf{o}) \nabla_{\tilde{\phi}_{1,sh}^i} Q_{\tilde{\phi}_{1,sh}^i}^{\pi, i}(\mathbf{o}, \mathbf{a}) \\
 &= \mathbf{E}_{s \sim d_0, \mathbf{o} \sim \mathcal{E}, \mathbf{a} \sim \pi} \left[\nabla_{\tilde{\phi}_{1,sh}^i} Q_{\tilde{\phi}_{1,sh}^i}^{\pi, i}(\mathbf{o}, \mathbf{a}) \right].
 \end{aligned}$$

For calculate the gradient use off-policy data we also introduce importance sampling, then we have

$$\nabla_{\tilde{\phi}_{1,sh}^i} J_{\text{actor}}^i(\mathbf{w}_{in}, \tilde{\mathbf{w}}_{sh}^i) = \mathbf{E}_{s \sim d_0, \mathbf{o} \sim \mathcal{E}, \mathbf{a} \sim \pi_0} \left[\left(\frac{\pi_{\psi_{in}^i}^i(a^i | o^i)}{\pi_0^i(a^i | o^i)} \right) \nabla_{\tilde{\phi}_{1,sh}^i} Q_{\tilde{\phi}_{1,sh}^i}^{\pi, i}(\mathbf{o}, \mathbf{a}) \right].$$

We then combine the above two part gradient:

$$\begin{aligned}
 \nabla_{\tilde{\phi}_{1,sh}^i} J_{\text{actor}}^i(\mathbf{w}_{in}, \tilde{\mathbf{w}}_{sh}^i) &= \mathbf{E}_{s \sim d_0, \mathbf{o} \sim \mathcal{E}, \mathbf{a} \sim \pi} \left[\nabla_{\tilde{\phi}_{1,sh}^i} Q_{\tilde{\phi}_{1,sh}^i}^{\pi, i}(\mathbf{o}, \mathbf{a}) \right] + \\
 & \quad \mathbf{E}_{s \sim d_0, \mathbf{o} \sim \mathcal{E}, \mathbf{a} \sim \pi_0} \left[\left(\frac{\pi_{\psi_{in}^i}^i(a^i | o^i)}{\pi_0^i(a^i | o^i)} \right) \nabla_{\tilde{\phi}_{1,sh}^i} Q_{\tilde{\phi}_{1,sh}^i}^{\pi, i}(\mathbf{o}, \mathbf{a}) \right].
 \end{aligned}$$

For critic part, we first calculate the gradient use the off-policy data

$$\begin{aligned}
 \nabla_{\tilde{\phi}_{1,sh}^i} J_{\text{critic}}^i(\mathbf{w}_{in}, \tilde{\mathbf{w}}_{sh}^i) &= \nabla_{\tilde{\phi}_{1,sh}^i} \sum_s d_0(s) \sum_{\mathbf{o}} \mathcal{E}_{0,s}^o \sum_{\mathbf{a}} \pi_0(\mathbf{a} | \mathbf{o}) \\
 & \quad \left(\left(Q_{\tilde{\phi}_{1,sh}^i}^{\pi, i}(\mathbf{o}, \mathbf{a}) - Q_{tg}^\pi \right)^2 + \left(Q_{\tilde{\phi}_{2,sh}^i}^{\pi, i}(\mathbf{o}, \mathbf{a}) - Q_{tg}^\pi \right)^2 \right) \\
 &= \sum_s d_0(s) \sum_{\mathbf{o}} \mathcal{E}_{0,s}^o \sum_{\mathbf{a}} \pi_0(\mathbf{a} | \mathbf{o}) 2\delta_1 \nabla_{\tilde{\phi}_{1,sh}^i} Q_{\tilde{\phi}_{1,sh}^i}^{\pi, i}(\mathbf{o}, \mathbf{a}) \\
 &= \mathbf{E}_{s \sim d_0, \mathbf{o} \sim \mathcal{E}, \mathbf{a} \sim \pi_0} \left[2\delta_1 \nabla_{\tilde{\phi}_{1,sh}^i} Q_{\tilde{\phi}_{1,sh}^i}^{\pi, i}(\mathbf{o}, \mathbf{a}) \right].
 \end{aligned}$$

Then, for resampled data, we have

$$\nabla_{\tilde{\phi}_{1,sh}^i} J_{\text{critic}}^i(\mathbf{w}_{in}, \tilde{\mathbf{w}}_{sh}^i) = \mathbf{E}_{s \sim d_0, \mathbf{o} \sim \mathcal{E}, \mathbf{a} \sim \pi} \left[2\delta_1 \left(\frac{\pi_0^i(a^i|o^i)}{\pi_{\psi_{in}^i}^i(a^i|o^i)} \right) \nabla_{\tilde{\phi}_{1,sh}^i} Q_{\tilde{\phi}_{1,sh}^i}^{\pi,i}(\mathbf{o}, \mathbf{a}) \right].$$

Combining above two part gradient, we obtain:

$$\begin{aligned} \nabla_{\tilde{\phi}_{1,sh}^i} J_{\text{critic}}^i(\mathbf{w}_{in}, \tilde{\mathbf{w}}_{sh}^i) &= \mathbf{E}_{s \sim d_0, \mathbf{o} \sim \mathcal{E}, \mathbf{a} \sim \pi_0} \left[2\delta_1 \nabla_{\tilde{\phi}_{1,sh}^i} Q_{\tilde{\phi}_{1,sh}^i}^{\pi,i}(\mathbf{o}, \mathbf{a}) \right] + \\ &\quad \mathbf{E}_{s \sim d_0, \mathbf{o} \sim \mathcal{E}, \mathbf{a} \sim \pi} \left[2\delta_1 \left(\frac{\pi_0^i(a^i|o^i)}{\pi_{\psi_{in}^i}^i(a^i|o^i)} \right) \nabla_{\tilde{\phi}_{1,sh}^i} Q_{\tilde{\phi}_{1,sh}^i}^{\pi,i}(\mathbf{o}, \mathbf{a}) \right]. \end{aligned}$$

Finally, we get following off-policy joint gradient w.r.t. the first critic parameters

$$\begin{aligned} \nabla_{\tilde{\phi}_{1,sh}^i} J_{ac}^i(\mathbf{w}_{in}, \tilde{\mathbf{w}}_{sh}^i) &= \alpha_1 \nabla_{\tilde{\phi}_{1,sh}^i} J_{actor}^i(\mathbf{w}_{in}, \tilde{\mathbf{w}}_{sh}^i) + \alpha_2 \nabla_{\tilde{\phi}_{1,sh}^i} J_{critic}^i(\mathbf{w}_{in}, \tilde{\mathbf{w}}_{sh}^i) \\ &= \mathbf{E}_{s \sim d_0, \mathbf{o} \sim \mathcal{E}, \mathbf{a} \sim \pi_0} \left[\left(\alpha_1 \left(\frac{\pi_{\psi_{in}^i}^i(a^i|o^i)}{\pi_0^i(a^i|o^i)} \right) + 2\alpha_2 \delta_1 \right) \nabla_{\tilde{\phi}_{1,sh}^i} Q_{\tilde{\phi}_{1,sh}^i}^{\pi,i}(\mathbf{o}, \mathbf{a}) \right] + \\ &\quad \mathbf{E}_{s \sim d_0, \mathbf{o} \sim \mathcal{E}, \mathbf{a} \sim \pi} \left[\left(\alpha_1 + 2\alpha_2 \delta_1 \left(\frac{\pi_0^i(a^i|o^i)}{\pi_{\psi_{in}^i}^i(a^i|o^i)} \right) \right) \nabla_{\tilde{\phi}_{1,sh}^i} Q_{\tilde{\phi}_{1,sh}^i}^{\pi,i}(\mathbf{o}, \mathbf{a}) \right] \\ &\simeq \mathbf{E}_{s \sim d_0, \mathbf{o} \sim \mathcal{E}, \mathbf{a} \sim \pi_0} \left[\left(\alpha_1 \text{CIM}_\epsilon(\pi_{\psi_{in}^i}^i; \pi_0^i) + 2\alpha_2 \delta_1 \right) \nabla_{\tilde{\phi}_{1,sh}^i} Q_{\tilde{\phi}_{1,sh}^i}^{\pi,i}(\mathbf{o}, \mathbf{a}) \right] + \\ &\quad \mathbf{E}_{s \sim d_0, \mathbf{o} \sim \mathcal{E}, \mathbf{a} \sim \pi} \left[\left(\alpha_1 + 2\alpha_2 \delta_1 \text{CIM}_\epsilon(\pi_0^i; \pi_{\psi_{in}^i}^i) \right) \nabla_{\tilde{\phi}_{1,sh}^i} Q_{\tilde{\phi}_{1,sh}^i}^{\pi,i}(\mathbf{o}, \mathbf{a}) \right]. \end{aligned}$$

For the third equation and use resampled data to calculate the actor part gradient, we have

$$\begin{aligned} \nabla_{\tilde{\phi}_{2,sh}^i} J_{actor}^i(\mathbf{w}_{in}, \tilde{\mathbf{w}}_{sh}^i) &= \nabla_{\tilde{\phi}_{2,sh}^i} \sum_s d_0(s) \sum_{\mathbf{o}} \mathcal{E}(\mathbf{o}|s) \sum_{\mathbf{a}} \pi_{\Psi_{in}}(\mathbf{a}|\mathbf{o}) \left[Q_{\tilde{\phi}_{1,sh}^i}^{\pi,i}(\mathbf{o}, \mathbf{a}) - \alpha \log \pi_{\psi_{in}^i}^i(a^i|o^i) - \mathcal{B}(\mathbf{o}, \mathbf{a}) \right] \\ &= 0 \end{aligned}$$

Similarly, for off-policy data the gradient also is 0. For critic part the gradient is similar as the second equation

$$\begin{aligned} \nabla_{\tilde{\phi}_{2,sh}^i} J_{critic}^i(\mathbf{w}_{in}, \tilde{\mathbf{w}}_{sh}^i) &= \mathbf{E}_{s \sim d_0, \mathbf{o} \sim \mathcal{E}, \mathbf{a} \sim \pi_0} \left[2\delta_2 \nabla_{\tilde{\phi}_{2,sh}^i} Q_{\tilde{\phi}_{2,sh}^i}^{\pi,i}(\mathbf{o}, \mathbf{a}) \right] + \\ &\quad \mathbf{E}_{s \sim d_0, \mathbf{o} \sim \mathcal{E}, \mathbf{a} \sim \pi} \left[2\delta_2 \left(\frac{\pi_0^i(a^i|o^i)}{\pi_{\psi_{in}^i}^i(a^i|o^i)} \right) \nabla_{\tilde{\phi}_{2,sh}^i} Q_{\tilde{\phi}_{2,sh}^i}^{\pi,i}(\mathbf{o}, \mathbf{a}) \right]. \end{aligned}$$

Finally, we get following off-policy joint gradient w.r.t. the second critic parameters

$$\nabla_{\tilde{\phi}_{2,sh}^i} J_{ac}^i(\mathbf{w}_{in}, \tilde{\mathbf{w}}_{sh}^i) = \alpha_1 \nabla_{\tilde{\phi}_{2,sh}^i} J_{actor}^i(\mathbf{w}_{in}, \tilde{\mathbf{w}}_{sh}^i) + \alpha_2 \nabla_{\tilde{\phi}_{2,sh}^i} J_{critic}^i(\mathbf{w}_{in}, \tilde{\mathbf{w}}_{sh}^i)$$

$$\begin{aligned}
 &= \mathbf{E}_{s \sim d_0, o \sim \mathcal{E}, a \sim \pi_0} \left[2\alpha_2 \delta_2 \nabla_{\tilde{\phi}_{2,sh}^i} Q_{\tilde{\phi}_{2,sh}^i}^{\pi,i}(\mathbf{o}, \mathbf{a}) \right] + \\
 &\quad \mathbf{E}_{s \sim d_0, o \sim \mathcal{E}, a \sim \pi} \left[2\alpha_2 \delta_2 \left(\frac{\pi_0^i(a^i | o^i)}{\pi_{\psi_{in}^i}^i(a^i | o^i)} \right) \nabla_{\tilde{\phi}_{2,sh}^i} Q_{\tilde{\phi}_{2,sh}^i}^{\pi,i}(\mathbf{o}, \mathbf{a}) \right] \\
 &\simeq \mathbf{E}_{s \sim d_0, o \sim \mathcal{E}, a \sim \pi_0} \left[2\alpha_2 \delta_2 \nabla_{\tilde{\phi}_{2,sh}^i} Q_{\tilde{\phi}_{2,sh}^i}^{\pi,i}(\mathbf{o}, \mathbf{a}) \right] + \\
 &\quad \mathbf{E}_{s \sim d_0, o \sim \mathcal{E}, a \sim \pi} \left[2\alpha_2 \delta_2 \text{CIM}_\epsilon(\pi_0^i; \pi_{\psi_{in}^i}^i) \nabla_{\tilde{\phi}_{2,sh}^i} Q_{\tilde{\phi}_{2,sh}^i}^{\pi,i}(\mathbf{o}, \mathbf{a}) \right].
 \end{aligned}$$

■

Appendix F. The Specific Forms of Instantiation Algorithms

The specific forms of F2A2-COMA, F2A2-DDPG, F2A2-TD3 and F2A2-SAC are shown in Figure 10.

Appendix G. Training procedure

Cooperative MPE. For all algorithms except for COMA and F2A2-COMA, our training procedure consists of performing 12 parallel rollouts, and adding a tuple of $(o_t, a_t, c_t, o_{t+1})_{1 \dots N}$ to a replay buffer (with maximum length 1e6) for each timestep. We reset each environment after every 100 steps for *Cooperative Treasure Collection* and 25 steps for *Rover Tower* (across all rollouts), we perform 4 updates for the all actors and critics. For each update, we sample mini-batches of 1024 timesteps from the replay buffer (for COMA and F2A2-COMA, we used the most recent 1024 timesteps) and then perform gradient descent on the corresponding loss objective, using Adam (Kingma and Ba, 2015) as the optimizer with a learning rate of 0.001. After the updates are complete, we update the target network parameters (if there are) to move toward our learned parameters, as in (Haarnoja et al., 2018): $\bar{\Theta} = (1 - \tau)\bar{\Theta} + \tau\Theta$, where τ is the soft update rate (set to 0.005). We use a discount factor, γ , of 0.99. All networks use a hidden dimension of 128 and Leaky Rectified Linear Units as the nonlinearity. We use 0.01 as our temperature setting for MAAC and F2A2-SAC. We use 4 attention heads in MAAC and F2A2-SAC.

StarCraft II. All policies are implemented as two-layer recurrent neural networks (GRUs) with 64 hidden units, while the critic is the same as the settings of *Cooperative Multi-agent Particle Environments*. For exploration, we use a bounded softmax distribution in which the agent samples from a softmax over the policy logits with probability $(1 - \epsilon)$ and samples randomly with probability ϵ . We anneal ϵ from 0.5 to 0.01 across the first 50k environment steps. Episodes are collected using eight parallel SCII environments. Optimization is carried out on a single GPU with Adam and a learning rate of 0.0005 for both the agents and the critics. The policies are fully unrolled and updated in a large mini-batch of TB entries, where $T = 60$ and $B = 8$. By contrast, the critic is optimized in small mini-batches of size 8, one for each time-step, looping backward in time. We found that this stabilized and

accelerated training compared to full batch updates for the critic. The target network for the critic is updated after every 200 critic update.

MAgent Environment. For the MAgent environment, we fix the model of the enemy agents to IQL after a fixed number (2000) of self-play training rounds. The learning rate for actors and critics is 0.0001. The discount factor is set to 0.95, and the mini-batch size is 128. The size of the replay buffer is 500000. And the temperature for F2A2-ISAC is 0.08.

While for the algorithm to model other agents, we first collect 1000 trajectories randomly to construct the past trajectories dataset and fixed it. We keep the most recent 1200 state-action pairs of all agents; thus, the mini-batches size of the impromptu prediction net is also 1200. The char-net and mental-net are 2-layer GRU networks with 64 hidden dimensions, and two prediction net are MLP networks with 128 hidden dimensions. The length of the former is fixed to the length of the episode, 25, and the latter is fixed at 5. The update frequency of the whole network is the same as the reinforcement learning part, and the mini-batches size of the natural prediction net is 1600. The optimizer and learning rate are also set the same as those in the reinforcement learning part.

Except that the F2A2-ISAC algorithm uses the Tensorflow (Abadi et al., 2015) framework, the other algorithms we proposed use the PyTorch (Paszke et al., 2017) framework. We run all the experiments on a machine with 44 CPU cores, 128G RAM, and 4 Nvidia 1080Ti GPUs. We use the original papers ’ open-source code for all of the simulation environments involved in the experiments.

Appendix H. More Case Studies

In the following, we will discuss more state-of-the-art MARL algorithms to motivate the above optimization formulation (4.5) and the following proposed algorithm framework. There contain to the communication learning algorithm (Jiang and Lu, 2018) and common knowledge based algorithm (Schroeder de Witt et al., 2019).

H.1 Attentional Communication (ATOC)

The communication learning algorithms advocate learning collaborative policies through communication. However, for early communication learning algorithms(Foerster et al., 2016; Peng et al., 2017; Sukhbaatar et al., 2016), information sharing among all agents or in predefined communication architectures, these methods adopt can be problematic. When there is a large number of agents, agents cannot differentiate valuable information that helps cooperative decision making from globally shared information, and hence communication barely helps and could even jeopardize the learning of cooperation. Moreover, in real-world applications, it is costly that all agents communicate with each other since receiving a large amount of information requires high bandwidth and incurs long delay and high computational complexity(Jiang and Lu, 2018).

ATOC(Jiang and Lu, 2018) propose an attentional communication model to solve the above problems. ATOC introduces a shared attention unit so that the communication architecture between agents can dynamically change according to needs; at the same time, in a self-organized communication group, the message is generated through a shared Bi-LSTM

\mathbf{w}_{in}		\mathbf{w}_{sh}		Functions: J_{actor}^i and J_{critic}^i	
ψ_{in}^i	ϕ_{in}^i	ψ_{sh}^i	ϕ_{sh}^i	$J_{\text{actor}}^i(\psi_{sh}, \phi_{sh})$	$J_{\text{critic}}^i(\psi_{in}, \phi_{sh})$
\emptyset	\emptyset	$\{\theta_{enc}, \theta_{ram}, \theta_{lstm}, \theta_{actor}\}$	$\pi^i(o^i)$ (MLP)	$\mathbb{E} \left[Q_{\phi_{sh}}^{\pi, i}(o^i, \pi_{\psi_{sh}}^i(o^i)) \right]$	$\mathbb{E} \left[\left(Q_{\phi_{sh}}^{\pi, i}(o^i, a^i) - Q_{tg}^{\pi, i} \right)^2 \right]$

Table 6: Algorithm elements about ATOC.

network, the message generation module. In a word, ATOC sharing all actor parameters and all critic parameters for all agents, which means $\phi_{sh} = \emptyset$ and $\psi_{sh} = \emptyset$. The ATOC algorithm uses the DDPG(Lillicrap et al., 2016) as the backbone, and the objective function for each agent i in actor and critic phases can be reformulated as

$$\begin{aligned}
 J_{\text{actor}}^i(\psi_{sh}, \phi_{sh}) &= \mathbb{E} \left[Q_{\phi_{sh}}^{\pi, i}(o^i, \pi_{\psi_{sh}}^i(o^i)) \right], \\
 J_{\text{critic}}^i(\psi_{sh}, \phi_{sh}) &= \mathbb{E} \left[\left(Q_{\phi_{sh}}^{\pi, i}(o^i, a^i) - Q_{tg}^{\pi, i} \right)^2 \right],
 \end{aligned}$$

where the expectations are taken on $o^i, a^i, c^i, o'^i \sim \mathcal{D}$ because the DDPG algorithm are off-policy algorithm, and \mathcal{D} is the shared experience replay buffer; $\pi := \left\{ \pi_{\psi_{sh}}^i \right\}$ represents the joint policy of all agents; ψ_{sh} and ϕ_{sh} represent the sharing actor parameter and critic parameter respectively; $Q_{tg}^{\pi, i} := c^i + \gamma Q_{\phi_{sh}}^{\pi, i}(o', \pi_{\psi_{sh}}^i(o'))$;

The detailed calculation process of $\pi_{\psi_{sh}}^i(o^i)$ is as follows: Firstly, the observation of agent i is encoded by an observation encoder which is parameterized by θ_{enc} . Then the observation embedding is fed into a recurrent attention model, which is parameterized by θ_{ram} , and output a two-valued variable, indicating whether or not agent i becomes the founder of the communication group. If the output value is 1, then agent i becomes the founder of a new communication group; otherwise, agent i chooses an existing communication group to join according to a pre-defined strategy. After that, all messages in the communication group are fed into an LSTM network which is parameterized by θ_{lstm} , and the output corresponding to the agent i is used as a fusion representation of the message sent by the remaining agents in the communication group. Finally, this fusion representation is fed into the actor-network, which is parameterized by θ_{actor} , to obtain the final policy. It can be seen that $\psi_{sh} := \{\theta_{enc}, \theta_{ram}, \theta_{lstm}, \theta_{actor}\}$ and $\psi_{in}^i := \emptyset$, ϕ_{sh} is the critic sharing between all agents and $\phi_{in}^i := \emptyset$.

The overall optimization problem specified from (4.5) can be formulated as

$$\min_{\psi_{sh}, \phi_{sh}} \alpha_1 \sum_{i=1}^n \mathbb{E} \left[Q_{\phi_{sh}}^{\pi, i}(o^i, \pi_{\psi_{sh}}^i(o^i)) \right] + \alpha_2 \sum_{i=1}^n \mathbb{E} \left[\left(Q_{\phi_{sh}}^{\pi, i}(o^i, a^i) - Q_{tg}^{\pi, i} \right)^2 \right]. \quad (\text{H.1})$$

Note that here we omit the parameters regularization term $\mathcal{R}(\mathbf{w}_{in}, \mathbf{w}_{sh})$ in (4.5). In the practical implementation of the ATOC algorithm, the regularization of the parameters is generally implemented by L_2 regularization or gradient norm clipping. All detailed elements about ATOC are summarized in Table 6, and the algorithm framework of ATOC for

formulation (4.5) is show as the following form

$$\left\{ \begin{array}{l} \phi_{sh}^{k+1} = \phi_{sh}^k - \alpha_1 \frac{1}{M} \sum_{m=1}^M \sum_{i=1}^n \nabla_{\phi_{sh}} \left(Q_{\phi_{sh}^k}^{\pi^{k,i}}(o^{i,m}, a^{i,m}) - Q_{tg}^{\pi^{k,i}} \right)^2 \\ \approx \arg \min_{\phi_{sh}} \alpha_1 \sum_{i=1}^n \mathbb{E} \left[\left(Q_{\phi_{sh}}^{\pi,i}(o^i, a^i) - Q_{tg}^{\pi,i} \right)^2 \right], \\ \{\theta_{enc}, \theta_{actor}\}^{k+1} = \{\theta_{enc}, \theta_{actor}\}^k - \alpha_2 \frac{1}{M} \sum_{m=1}^M \sum_{i=1}^n \\ \nabla_{\{\theta_{enc}, \theta_{actor}\}} \left(Q_{\phi_{sh}^{k+1}}^{\pi,i}(o^{i,m}, \pi_{\psi_{sh}^k}^i(o^{i,m})) \right) \\ \approx \arg \min_{\{\theta_{enc}, \theta_{actor}\}} \alpha_2 \sum_{i=1}^n \mathbb{E} \left[Q_{\phi_{sh}}^{\pi,i}(o^i, \pi_{\psi_{sh}}^i(o^i)) \right], \\ \theta_{lstm}^{k+1} = \theta_{lstm}^k - \alpha_2 \frac{1}{M} \sum_{m=1}^M \sum_{i=1}^n \nabla_{\theta_{lstm}} \left(Q_{\phi_{sh}^{k+1}}^{\pi,i}(o^{i,m}, \pi_{\psi_{sh}^k}^i(o^{i,m})) \right) \\ \approx \arg \min_{\theta_{lstm}} \alpha_2 \sum_{i=1}^n \mathbb{E} \left[Q_{\phi_{sh}}^{\pi,i}(o^i, \pi_{\psi_{sh}}^i(o^i)) \right], \end{array} \right.$$

which can be considered to employ block coordinate gradient descent on formulation (H.1).

H.2 Multi-Agent Common Knowledge (MACKRL)

For the cooperative POSG, in the absence of common knowledge, complex decentralized coordination has to rely on implicit communication, i.e., observing each other’s actions or their effects (Heider and Simmel, 1944; Rasouli et al., 2017). However, implicit communication protocols for complex coordination problems are challenging to learn and, as they typically require multiple timesteps to execute, can limit the agility of control during execution (Tian et al., 2018). By contrast, coordination based on common knowledge is simultaneous and does not require learning communication protocols (Halpern and Moses, 1990).

MACKRL (Schroeder de Witt et al., 2019) is a novel stochastic policy actor-critic algorithm that can learn complex coordination policies end-to-end by exploiting common knowledge between groups of agents at the appropriate level. MACKRL uses a hierarchical policy tree in order to select the right level of coordination dynamically. Specifically, MACKRL is learning a centralized actor and centralized critic. However, due to the use of hierarchical policies, the centralized actor can be calculated efficiently. In other words, just like ATOC, MACKRL sharing all actor parameters and all critic parameters for all agents, which means $\phi_{sh} = \emptyset$ and $\psi_{sh} = \emptyset$.

MACKRL is based on *Central-V* (Foerster et al., 2018a) and approximately solves the Eq. (4.5) by iteratively optimizing the same two subproblems as above algorithms, and the specific form of J_{actor}^i and J_{critic}^i of agent i in MACKRL algorithm are as follows:

$$\begin{aligned} J_{actor}^i(\psi_{sh}, \phi_{sh}) &= \mathbb{E} \left[Q_{\phi_{sh}}^{\pi,i}(\mathbf{o}, \mathbf{a}) \right], \\ J_{critic}^i(\psi_{sh}, \phi_{sh}) &= \mathbb{E} \left[\left(Q_{\phi_{sh}}^{\pi,i}(\mathbf{o}, \mathbf{a}) - Q_{tg}^{\pi,i} \right)^2 \right], \end{aligned}$$

where the expectations are taken on $s \sim d_{\Psi}$, $\mathbf{o} \sim \mathcal{E}$, $\mathbf{a} \sim \boldsymbol{\pi}_{\Psi}$ and $\Psi := \{\psi_{sh}\}$; $\boldsymbol{\pi} := \boldsymbol{\pi}_{\psi_{sh}}$ represents the joint policy of all agents; ψ_{sh} and ϕ_{sh} represent the sharing actor parameter and critic parameter respectively; MACKRL also uses TD(λ) algorithm to learn the shared

\mathbf{w}_{in}		\mathbf{w}_{sh}		Functions: J_{actor}^i and J_{critic}^i	
ψ_{in}^i	ϕ_{in}^i	ψ_{sh}^i	ϕ_{sh}^i	$J_{actor}^i(\psi_{sh}, \phi_{sh})$	$J_{critic}^i(\psi_{sh}, \phi_{sh})$
\emptyset	\emptyset	$\{\psi_{group}, \psi_{agent}\}$ (GRUs)	The parameter of Q and V (MLP)	$\mathbb{E} \left[Q_{\phi_{sh}}^{\pi, i}(\mathbf{o}, \mathbf{a}) \right]$	$\mathbb{E} \left[\left(Q_{\phi_{sh}}^{\pi, i}(\mathbf{o}, \mathbf{a}) - Q_{tg}^{\pi, i} \right)^2 \right]$

Table 7: Algorithm information about MACKRL.

critic and it will not be repeated here because it was mentioned in the previous introduction of the COMA algorithm.

The way that MACKRL effectively calculates joint policy is similar to hierarchical reinforcement learning. It divides joint policy into common-knowledge-based group-level policies and common-knowledge-based agent-level policies, so ψ_{sh} can be decomposed into ψ_{group} and ψ_{agent} . MACKRL allows multi-agent policies to introduce common knowledge while training end-to-end efficiently naturally (see the original paper (Schroeder de Witt et al., 2019) for details).

The overall optimization problem specified from (4.5) can be formulated as

$$\min_{\psi_{sh}, \phi_{sh}} \alpha_1 \sum_{i=1}^n \mathbb{E} \left[Q_{\phi_{sh}}^{\pi, i}(\mathbf{o}, \mathbf{a}) \right] + \alpha_2 \sum_{i=1}^n \mathbb{E} \left[\left(Q_{\phi_{sh}}^{\pi, i}(\mathbf{o}, \mathbf{a}) - Q_{tg}^{\pi, i} \right)^2 \right]. \quad (\text{H.2})$$

Note that here we omit the parameters regularization term $\mathcal{R}(\mathbf{w}_{in}, \mathbf{w}_{sh})$ in (4.5). In the practical implementation of the MACKRL algorithm, the regularization of the parameters is generally implemented by L_2 regularization or gradient norm clipping. All detailed elements about MACKRL are summarized in Table 7, and the algorithm framework of MACKRL for formulation (4.5) is show as the following form

$$\left\{ \begin{array}{l} \phi_{sh}^{k+1} = \phi_{sh}^k - \alpha_1 \frac{1}{M} \sum_{m=1}^M \sum_{i=1}^n \nabla_{\phi_{sh}} \left(Q_{\phi_{sh}^k}^{\pi^k, i}(\mathbf{o}^m, \mathbf{a}^m) - Q_{tg}^{\pi^k, i} \right)^2 \\ \quad \approx \arg \min_{\phi_{sh}} \alpha_1 \sum_{i=1}^n \mathbb{E} \left[\left(Q_{\phi_{sh}}^{\pi, i}(\mathbf{o}, \mathbf{a}) - Q_{tg}^{\pi, i} \right)^2 \right], \\ \psi_{sh}^{k+1} = \psi_{sh}^k - \alpha_2 \frac{1}{M} \sum_{m=1}^M \sum_{i=1}^n \nabla_{\psi_{sh}} \left(Q_{\phi_{sh}^{k+1}}^{\pi^k, i}(\mathbf{o}^m, \mathbf{a}^m) \right) \\ \quad \approx \arg \min_{\psi_{sh}} \alpha_2 \sum_{i=1}^n \mathbb{E} \left[Q_{\phi_{sh}}^{\pi, i}(\mathbf{o}, \mathbf{a}) \right], \end{array} \right.$$

which can be considered to employ block coordinate gradient descent on formulation (H.2).

H.3 Multi-Agent Deep Deterministic Policy Gradient (MADDPG)

MADDPG algorithm adopting the framework of CTDE. Thus, it allows the policies to use extra information to ease training so long as it is not used at test time. However, it is unnatural to do this with Q-learning, as the Q function generally cannot contain different information at training and test time. Thus, MADDPG proposes a simple extension of actor-critic methods based on the DDPG algorithm, where the critic is augmented with extra information about other agents' policies.

Since the MADDPG algorithm is directly extended on the DDPG, the difference from the DDPG algorithm is that additional information of the other agents is introduced when

\mathbf{w}_{in}		\mathbf{w}_{sh}		Functions: J_{actor}^i and J_{critic}^i	
ψ_{in}^i	ϕ_{in}^i	ψ_{sh}^i	ϕ_{sh}^i	$J_{\text{actor}}^i(\{\psi_{in}^i\}, \phi_{sh})$	$J_{\text{critic}}^i(\{\psi_{in}^i\}, \phi_{sh})$
$\pi^i(\mathbf{o}^i)$	\emptyset	\emptyset	$Q(\mathbf{o}, \mathbf{a})$	$\mathbb{E} \left[Q_{\phi_{sh}}^{\pi, i}(\mathbf{o}, \mathbf{a}) \right]$	$\mathbb{E} \left[\left(Q_{\phi_{sh}}^{\pi, i}(\mathbf{o}, \mathbf{a}) - Q_{tg}^{\pi, i} \right)^2 \right]$

Table 8: Algorithm information about MADDPG.

train the centralized critic. For each agent, it is still essentially completing a single-agent task. Therefore, there is no shared parameter between agents, which means $\psi_{sh} = \emptyset$ and $\phi_{sh} = \emptyset$. Because this paper focuses on the collaboration problem, minor modifications to MADDPG are made here. That is, all agents share a centralized Q-value function, which means $\phi_{in}^i = \emptyset$. The objective function for each agent i in actor and critic phases can be reformulated as

$$\begin{aligned} J_{\text{actor}}^i(\{\psi_{in}^i\}, \phi_{sh}) &= \mathbb{E} \left[Q_{\phi_{sh}}^{\pi, i}(\mathbf{o}, \mathbf{a}) \right], \\ J_{\text{critic}}^i(\{\psi_{in}^i\}, \phi_{sh}) &= \mathbb{E} \left[\left(Q_{\phi_{sh}}^{\pi, i}(\mathbf{o}, \mathbf{a}) - Q_{tg}^{\pi, i} \right)^2 \right], \end{aligned}$$

where the expectations are taken on $\mathbf{o}, \mathbf{a}, \mathbf{c}, \mathbf{o}' \sim \mathcal{D}$ because the DDPG algorithm are off-policy algorithm, and \mathcal{D} is the shared experience replay buffer; $\boldsymbol{\pi} := \{\pi_{\psi_{in}^i}^i\}$ represents the policy of each agent; ψ_{in}^i and ϕ_{sh} represent the independent actor parameter and shared critic parameter respectively; $Q_{tg}^{\pi, i} := c^i + \gamma Q_{\phi_{sh}}^{\pi, i}(\mathbf{o}', \boldsymbol{\pi}(\mathbf{o}'))$.

The overall optimization problem specified from (4.5) can be formulated as

$$\min_{\psi_{sh}, \phi_{sh}} \alpha_1 \sum_{i=1}^n \mathbb{E} \left[Q_{\phi_{sh}}^{\pi, i}(\mathbf{o}, \mathbf{a}) \right] + \alpha_2 \sum_{i=1}^n \mathbb{E} \left[\left(Q_{\phi_{sh}}^{\pi, i}(\mathbf{o}, \mathbf{a}) - Q_{tg}^{\pi, i} \right)^2 \right]. \quad (\text{H.3})$$

All detailed information about MADDPG is summarized in Table 8, and the algorithm framework of MADDPG for formulation (4.5) is shown as the following form

$$\left\{ \begin{array}{l} \phi_{sh}^{k+1} = \phi_{sh}^k - \alpha_1 \frac{1}{M} \sum_{m=1}^M \sum_{i=1}^N \nabla_{\phi_{sh}} \left(Q_{\phi_{sh}^k}^{\pi^k, i}(\mathbf{o}^m, \mathbf{a}^m) - Q_{tg}^{\pi^k, i} \right)^2 \\ \quad \approx \arg \min_{\phi_{sh}} \alpha_1 \sum_{i=1}^n \mathbb{E} \left[\left(Q_{\phi_{sh}}^{\pi, i}(\mathbf{o}, \mathbf{a}) - Q_{tg}^{\pi, i} \right)^2 \right], \\ \{\psi_{in}^i\}^{k+1} = \{\psi_{in}^i\}^k - \alpha_2 \frac{1}{M} \sum_{m=1}^M \sum_{n=1}^N \nabla_{\{\psi_{in}^i\}} \left(Q_{\phi_{sh}^{k+1}}^{\pi^k, i}(\mathbf{o}^m, \mathbf{a}^m) - Q_{tg}^{\pi^k, i} \right)^2 \\ \quad \approx \arg \min_{\{\psi_{in}^i\}} \alpha_2 \sum_{i=1}^n \mathbb{E} \left[\left(Q_{\phi_{sh}}^{\pi, i}(\mathbf{o}, \mathbf{a}) - Q_{tg}^{\pi, i} \right)^2 \right], \end{array} \right.$$

which can be considered to employ block coordinate gradient descent for formulation (H.3).

H.4 Counterfactual Multi-Agent Policy Gradients (COMA)

In fully cooperative POSG, joint actions typically generate only the global cost (i.e., the same cost function sharing between agents), making it difficult for each agent to deduce its contribution to the team's success. In some cases, it is possible to design an individual cost function for each agent. However, these costs are not generally available in a collaborative

\mathbf{w}_{in}		\mathbf{w}_{sh}		Functions: J_{actor}^i and J_{critic}^i	
ψ_{in}^i	ϕ_{in}^i	ψ_{sh}	ϕ_{sh}	$J_{\text{actor}}^i(\{\psi_{in}^i\}, \phi_{sh})$	$J_{\text{critic}}^i(\{\psi_{in}^i\}, \phi_{sh})$
$\pi^i(o^i)$ (GRUs ⁶)	\emptyset	\emptyset	$Q(\mathbf{o}, \mathbf{a})$ and $V(\mathbf{o})$ (MLP ⁷)	$\mathbb{E} \left[Q_{\phi_{sh}}^{\pi, i}(\mathbf{o}, \mathbf{a}) - \mathcal{B}(\mathbf{o}, \mathbf{a}^{\setminus i}) \right]$	$\mathbb{E} \left[\left(Q_{\phi_{sh}}^{\pi, i}(\mathbf{o}, \mathbf{a}) - Q_{tg}^{\pi, i} \right)^2 \right]$

Table 9: Algorithm elements about COMA.

setting and usually fail to encourage individual agents to sacrifice to obtain better global performance. This will substantially impede multi-agent learning in challenging tasks, even with a relatively small number of agents (Foerster et al., 2018a). This crucial challenge is called *multi-agent credit assignment* problem (Chang et al., 2004). COMA (Foerster et al., 2018a) solves this problem by learning a centralized critic with a *counterfactual baseline* which is inspired by *difference rewards* (Wolpert and Tumer, 2002).

In the COMA algorithm, for the reason that a centralized critic is employed, so that all the parameters in critic are shared, which means $\phi_{in}^i = \emptyset$ for all agents. However there is no shared parameter in actors, which means $\psi_{sh} = \emptyset$. Further, the objective function for each agent i in actor and critic phases can be reformulated as

$$\begin{aligned} J_{\text{actor}}^i(\{\psi_{in}^i\}, \phi_{sh}) &= \mathbb{E} \left[Q_{\phi_{sh}}^{\pi, i}(\mathbf{o}, \mathbf{a}) - \mathcal{B}(\mathbf{o}, \mathbf{a}^{\setminus i}) \right], \\ J_{\text{critic}}^i(\{\psi_{in}^i\}, \phi_{sh}) &= \mathbb{E} \left[\left(Q_{\phi_{sh}}^{\pi, i}(\mathbf{o}, \mathbf{a}) - Q_{tg}^{\pi, i} \right)^2 \right], \end{aligned}$$

where where the expectations are taken on $s \sim d_{\Psi}$, $\mathbf{o} \sim \mathcal{E}$, $\mathbf{a} \sim \pi_{\Psi}$ and $\Psi := \{\psi_{in}^i\}$; $\pi := \left\{ \pi_{\psi_{in}^i}^i \right\}$ represents the joint policy of all agents; ϕ_{sh} represents the sharing parameter of the centralized critic; ψ_{in}^i represents the independent actor parameter of each agent; COMA introduces TD(λ) (Sutton and Barto, 2018) for critic learning, thus $Q_{tg}^{\pi, i}$ is also denoted as $G_t^{\lambda} := (1 - \lambda) \sum_{n=1}^{\infty} \lambda^{n-1} G_t^{(n)}$; when $n = 1$, we have $G_t^1 := \mathbf{E}_{s' \sim \mathcal{P}, \mathbf{o}' \sim \mathcal{E}} \left(\mathcal{C}_{s, \mathbf{a}}^{s'} + \gamma V_{\phi_{sh}}^{\pi, i}(\mathbf{o}') \right)$, and V represents the approximated *state-value function* (Sutton and Barto, 2018); $\mathcal{B}(\mathbf{o}, \mathbf{a}^{\setminus i})$ denotes the *multi-agent counterfactual baseline* that is used to solve the credit assignment problem, $\mathcal{B}(\mathbf{o}, \mathbf{a}^{\setminus i}) = \mathbb{E}_{a^i \sim \pi^i(o^i)} \left[Q_{\phi_{sh}}^{\pi, i}(\mathbf{o}, (a^i, \mathbf{a}^{\setminus i})) \right]$. The overall optimization problem specified from (4.5) can be formulated as

$$\min_{\{\psi_{in}^i\}, \phi_{sh}} \alpha_1 \sum_{i=1}^n \mathbb{E} \left[Q_{\phi_{sh}}^{\pi, i}(\mathbf{o}, \mathbf{a}) - \mathcal{B}(\mathbf{o}, \mathbf{a}^{\setminus i}) \right] + \alpha_2 \sum_{i=1}^n \mathbb{E} \left[\left(Q_{\phi_{sh}}^{\pi, i}(\mathbf{o}, \mathbf{a}) - Q_{tg}^{\pi, i} \right)^2 \right]. \quad (\text{H.4})$$

Note that here we omit the parameters regularization term $\mathcal{R}(\mathbf{w}_{in}, \mathbf{w}_{sh})$ in (4.5). In the practical implementation of the COMA algorithm, the regularization of the parameters is generally implemented by L_2 regularization or gradient norm clipping. All detailed elements about COMA are summarized in Table 9. Further the algorithm framework of COMA

algorithm is shown follows, i.e.,

$$\left\{ \begin{array}{l} \phi_{sh}^{k+1} = \phi_{sh}^k - \alpha_1 \frac{1}{M} \sum_{m=1}^M \sum_{i=1}^n \nabla_{\phi_{sh}} \left(Q_{\phi_{sh}^k}^{\pi^k, i}(\mathbf{o}^m, \mathbf{a}^m) - Q_{tg}^{\pi^k, i} \right)^2 \\ \approx \arg \min_{\phi_{sh}} \alpha_1 \sum_{i=1}^n \mathbb{E} \left[\left(Q_{\phi_{sh}}^{\pi, i}(\mathbf{o}, \mathbf{a}) - Q_{tg}^{\pi, i} \right)^2 \right], \\ \{\psi_{in}^{i, k+1}\} = \{\psi_{in}^{i, k}\} - \alpha_2 \frac{1}{M} \sum_{m=1}^M \sum_{i=1}^n \nabla_{\{\psi_{in}^i\}} \left(Q_{\phi_{sh}^{k+1}}^{\pi^k, i}(\mathbf{o}^m, \mathbf{a}^m) - \mathcal{B}(\mathbf{o}^m, \mathbf{a}^{\setminus i, m}) \right) \\ \approx \arg \min_{\{\psi_{in}^i\}} \alpha_2 \sum_{i=1}^n \mathbb{E} \left[Q_{\phi_{sh}}^{\pi, i}(\mathbf{o}, \mathbf{a}) - \mathcal{B}(\mathbf{o}, \mathbf{a}^{\setminus i}) \right], \end{array} \right.$$

which can be considered to employ block coordinate gradient descent on formulation (H.4).

H.5 Multi-Actor-Attention-Critic (MAAC)

A large-scale cooperative multi-agent system is complicated and unnecessary for an agent to surveillance all agents' states and behavior. Meanwhile, the decision of each agent may only be affected by these strongly related agents, not all agents. Considering too much other agents' information will make proper signals inevitably submerged in the background noise. Therefore, the multiple attention actor-critic (MAAC) algorithm (Iqbal and Sha, 2019) introduces the effective attention mechanism to avoid the instability problem of estimating other agents' policy in MADDPG (Lowe et al., 2017). MAAC follows the learning procedure of centralized training with decentralized execution. Based on the popular Soft-Actor-Critic (SAC) algorithm (Haarnoja et al., 2018), MAAC considers an additional attention layer to avoid directly using other agents' policies, and the policy is determined by maximizing a trade-off between expected return and the entropy regularization.

The MAAC algorithm introduces a shared attention module in the modeling of each agent's critic, so that ϕ_{in}^i and ϕ_{sh} both *non-empty* parameter sets. Moreover, same as COMA, there is no shared parameter in actors, which means $\psi_{sh} = \emptyset$. The objective function for each agent i in actor and critic phases can be reformulated as

$$\begin{aligned} J_{\text{actor}}^i(\{\psi_{in}^i\}, \phi_{in}^i, \phi_{sh}) &= \mathbb{E} \left[Q_{\phi_{in}^i, \phi_{sh}}^{\pi, i}(\mathbf{o}, \mathbf{a}) + \alpha \mathcal{H}(\cdot | \pi_{\psi_{in}^i}^i(o^i)) - \mathcal{B}(\mathbf{o}, \mathbf{a}^{\setminus i}) \right], \\ J_{\text{critic}}^i(\{\psi_{in}^i\}, \phi_{in}^i, \phi_{sh}) &= \mathbb{E} \left[\left(Q_{\phi_{in}^i, \phi_{sh}}^{\pi, i}(\mathbf{o}, \mathbf{a}) - Q_{tg}^{\pi, i} \right)^2 \right], \end{aligned}$$

where the expectations are take on $\mathbf{o}, \mathbf{a}, \mathbf{c}, \mathbf{o}' \sim \mathcal{D}$ because the SAC algorithm are off-policy algorithm, and \mathcal{D} is the shared experience replay buffer; $\pi := \left\{ \pi_{\psi_{in}^i}^i \right\}$ represents the joint policy of all agents; ϕ_{sh} represents the sharing attention parameter of the centralized critic and ϕ_{in}^i represents the rest independent critic parameter of each agent; ψ_{in}^i represents the independent actor parameter of each agent; $\mathcal{H}(\cdot | \pi_{\psi_{in}^i}^i(o^i))$ denotes the entropy of the policy at state o^i ; $\mathcal{B}(\mathbf{o}, \mathbf{a}^{\setminus i})$ is the multi-agent counterfactual baseline that is same as COMA; $Q_{tg}^{\pi, i}$ is defined as

$$c_i + \gamma \mathbb{E}_{\mathbf{a}' \sim \pi(\mathbf{o}')} \left[Q_{\phi_{in}^i, \phi_{sh}}^{\pi, i}(\mathbf{o}', \mathbf{a}') - \alpha \log \left(\pi_{\psi_{in}^i}^i(a'_i | o'_i) \right) \right].$$

Finally, the detailed form of critic with attention mechanism is denoted as

$$Q_{\phi_{in}^i, \phi_{sh}}^{\pi, i}(\mathbf{o}, \mathbf{a}) = \zeta^i \left(F^i(o^i, a^i), \sum_{h=1}^H \sum_{j \neq i} \alpha_h^j \Upsilon(V_h F^j(o^j, a^j)) \right),$$

\mathbf{w}_{in}		\mathbf{w}_{sh}		Functions: J_{actor}^i and J_{critic}^i	
ψ_{in}^i	ϕ_{in}^i	ψ_{sh}^i	ϕ_{sh}^i	$J_{actor}^i(\{\psi_{in}^i\}, \phi_{in}^i, \phi_{sh}^i)$	$J_{critic}^i(\{\psi_{in}^i\}, \phi_{in}^i, \phi_{sh}^i)$
$\pi^i(o^i)$ (MLP)	ζ^i, F^i	\emptyset	$\{V_h, W_h^{key}, W_h^{que}\}_{h=1}^H$	$\mathbb{E} \left[Q_{\phi_{in}^i, \phi_{sh}^i}^{\pi, i}(\mathbf{o}, \mathbf{a}) + \alpha \mathcal{H}(\cdot \pi_{\psi_{in}^i}^i(o^i)) - \mathcal{B}(\mathbf{o}, \mathbf{a}^i) \right]$	$\mathbb{E} \left[\left(Q_{\phi_{in}^i, \phi_{sh}^i}^{\pi, i}(\mathbf{o}, \mathbf{a}) - Q_{tg}^{\pi, i} \right)^2 \right]$

Table 10: Algorithm elements about MAAC.

where $\phi_{in}^i := \{\zeta^i, F^i\}$ and ζ^i, F^i are two-layer multi-layer perceptron (MLP) and one-layer MLP encoding function respectively. Υ denotes a specific nonlinear activation function, and

$$\phi_{sh} := \left\{ V_h, W_h^{key}, W_h^{que} \right\}_{h=1}^H,$$

where V_h represents attention module parameters and $\alpha_h^{j \neq i} \propto \exp((F^{j \neq i})^T (W_h^{key})^T W_h^{que} F^i)$ represents the attention factor. H denotes the number of attention heads.

The overall optimization problem specified from (4.5) can be formulated as

$$\begin{aligned} \min_{\{\psi_{in}^i\}, \{\phi_{in}^i\}, \phi_{sh}} \alpha_1 \sum_{i=1}^n \mathbb{E} \left[Q_{\phi_{in}^i, \phi_{sh}}^{\pi, i}(\mathbf{o}, \mathbf{a}) + \alpha \mathcal{H}(\cdot | \pi_{\psi_{in}^i}^i(o^i)) - \mathcal{B}(\mathbf{o}, \mathbf{a}^i) \right] + \\ \alpha_2 \sum_{i=1}^n \mathbb{E} \left[\left(Q_{\phi_{in}^i, \phi_{sh}}^{\pi, i}(\mathbf{o}, \mathbf{a}) - Q_{tg}^{\pi, i} \right)^2 \right]. \end{aligned} \quad (\text{H.5})$$

Note that here we omit the parameters regularization term $\mathcal{R}(\mathbf{w}_{in}, \mathbf{w}_{sh})$ in (4.5). In the practical implementation of the MAAC algorithm, the regularization of the parameters is generally implemented by L_2 regularization or gradient norm clipping. All detailed elements about MAAC are summarized in Table 10.

Futher the algorithm framework of MAAC for formulation (H.5) is shown as follows,

$$\left\{ \begin{aligned} \left\{ V_h, W_h^{key}, W_h^{que} \right\}^{k+1} &= \left\{ V_h, W_h^{key}, W_h^{que} \right\}^k - \alpha_1 \frac{1}{M} \sum_{m=1}^M \sum_{i=1}^n \\ &\quad \nabla_{\{V_h, W_h^{key}, W_h^{que}\}} \left(Q_{\phi_{in}^i, \phi_{sh}^i}^{\pi^k, i}(\mathbf{o}^m, \mathbf{a}^m) - Q_{tg}^{\pi^k, i} \right)^2 \\ &\quad \approx \arg \min_{\{V_h, W_h^{key}, W_h^{que}\}} \alpha_1 \sum_{i=1}^n \mathbb{E} \left[\left(Q_{\phi_{in}^i, \phi_{sh}^i}^{\pi, i}(\mathbf{o}, \mathbf{a}) - Q_{tg}^{\pi, i} \right)^2 \right], \\ \left\{ \zeta^i, F^i \right\}^{k+1} &= \left\{ \zeta^i, F^i \right\}^k - \alpha_1 \frac{1}{M} \sum_{m=1}^M \nabla_{\{\zeta^i, F^i\}} \left(Q_{\phi_{in}^i, \phi_{sh}^i}^{\pi^k, i}(\mathbf{o}^m, \mathbf{a}^m) - Q_{tg}^{\pi^k, i} \right)^2 \\ &\quad \approx \arg \min_{\{\zeta^i, F^i\}} \alpha_1 \sum_{i=1}^n \mathbb{E} \left[\left(Q_{\phi_{in}^i, \phi_{sh}^i}^{\pi, i}(\mathbf{o}, \mathbf{a}) - Q_{tg}^{\pi, i} \right)^2 \right], \\ \left\{ \psi_{in}^i, k+1 \right\} &= \left\{ \psi_{in}^i, k \right\} - \alpha_2 \frac{1}{M} \sum_{m=1}^M \sum_{i=1}^n \\ &\quad \nabla_{\{\psi_{in}^i\}} \left(Q_{\phi_{in}^i, \phi_{sh}^i}^{\pi^k, i}(\mathbf{o}^m, \mathbf{a}^m) + \alpha \mathcal{H}(\cdot | \pi_{\psi_{in}^i}^i(o^{i,m})) - \mathcal{B}(\mathbf{o}^m, \mathbf{a}^i, m) \right) \\ &\quad \approx \arg \min_{\{\psi_{in}^i\}} \alpha_2 \sum_{i=1}^n \mathbb{E} \left[Q_{\phi_{in}^i, \phi_{sh}^i}^{\pi, i}(\mathbf{o}, \mathbf{a}) + \alpha \mathcal{H}(\cdot | \pi_{\psi_{in}^i}^i(o^i)) - \mathcal{B}(\mathbf{o}, \mathbf{a}^i) \right], \end{aligned} \right.$$

which can be considered to employ block coordinate gradient descent on formulation (H.5).

Appendix I. Limitations

We ensured the Markov property of the policy through the enrichment of the observation space, avoiding the use of non-Markovian and history-dependent policy classes. Although

this approach facilitates theoretical analysis and engineering implementation, the resulting *vast belief space* and *infinite hierarchy of belief* render the Markov policy class intractable in more complex problems, further constraining its generality. We shed light on the limitations that affect this policy class concerning crucial factors such as the error of the approximated information state and the infinite hierarchy belief representation capabilities. Careful consideration of these factors is essential to enhance the tractability and generality of the Markovian policy within the enriched observation space.

I.1 Error of the Approximated Information State

Lifting POMDPs to an MDP over the belief states would require estimating the belief states based on the whole history. This estimation is prone to the curse of dimensionality due to the ample space of belief states. To analyze the theoretical impact of this curse on algorithm performance, we adopt the perspective of *information state* based on existing work (Subramanian et al., 2022).

It is pertinent to mention that the concepts and theories related to the information state in POSG (Mao et al., 2020) are an extension of those in POMDP (Subramanian et al., 2022). A comprehensive description of these concepts would require a significant amount of space, and it is not the primary focus of this paper. Instead, we will concentrate on a specific agent in a POSG and gradually deduce the limitations of Markovian policies from a single-agent perspective. This approach allows us to omit the discussion of *common* and *private information* and their relevant content (Tavafoghi et al., 2021; Mao et al., 2020) while retaining the generalizability of our analysis. Concretely, we have the following definition (for a single agent in POSG):

Definition 3 (Subramanian et al. 2022) Let $\{Z_t\}_{t=1}^T$ be a pre-specified collection of Banach spaces, $H_t := \{\mathbf{O}_{1:t-1}, \mathbf{A}_{1:t-1}, \}$ be the past (joint) observations and actions. A collection $\{\sigma_t : H_t \rightarrow Z_t\}_{t=1}^T$ of history compression functions is called an information state generator and $\{Z_t\}_{t=1}^T$ is called an information state, if the process $\{Z_t\}_{t=1}^T$, where $Z_t = \sigma_t(H_t)$, satisfies the following properties:

(P1) Sufficient for performance evaluation, i.e., for any timestep t , any realization h_t of H_t and any choice \mathbf{a}_t of \mathbf{A}_t , we have

$$\mathbb{E}[R_t \mid H_t = h_t, \mathbf{A}_t = \mathbf{a}_t] = \mathbb{E}[R_t \mid Z_t = \sigma_t(h_t), \mathbf{A}_t = \mathbf{a}_t].$$

(P2) Sufficient to predict itself, i.e., for any time t , any realization h_t of H_t and any choice \mathbf{a}_t of \mathbf{A}_t , we have that for any Borel subset B of Z_{t+1} ,

$$\mathbb{P}(Z_{t+1} \in B \mid H_t = h_t, \mathbf{A}_t = \mathbf{a}_t) = \mathbb{P}(Z_{t+1} \in B \mid Z_t = \sigma_t(h_t), \mathbf{A}_t = \mathbf{a}_t).$$

It is evident that the history H_t is a trivial information state for any partially observed model Subramanian et al. (2022). In F2A2, we enrich the observation space based on all histories to create an information state space. However, for computational efficiency, F2A2 learns a Markovian policy in the approximated information state space. To accomplish this, we first perform a finite memory compression (FMC) of the entire history by maintaining a fixed window of the local history as an approximate information state. We then use LSTM, a type of recurrent neural network, to encode the truncated/approximated information state.

Nevertheless, truncating and embedding the information state leads to errors. As per theoretical results from Subramanian et al. (2022) and Mao et al. (2020), the distance between the optimal Markovian policy learned in the approximated information state space and the actual optimal Markovian policy is bounded by $(T - t + 1)(\varepsilon + L_V\delta)$, where L_V represents the upper bound of the Lipschitz constant of the policy and value function. The values of (ε, δ) are defined as follows:

Definition 4 Let $\{\hat{Z}_t\}_{t=1}^T$ be a pre-specified collection of Banach spaces, \mathfrak{F} be a probability metrics with a ζ -structure (Zolotarev, 1983), and (ε, δ) be pre-specified positive real numbers. A collection $\{\hat{\sigma}_t : H_t \rightarrow \hat{Z}_t\}_{t=1}^T$ of history compression functions, along with approximate update kernels $\{\hat{P}_t : \hat{Z}_t \times \mathbf{A} \rightarrow \Delta(\hat{Z}_{t+1})\}_{t=1}^T$ and reward approximation functions $\{\hat{r}_t : \hat{Z}_t \times \mathbf{A} \rightarrow \mathbb{R}\}_{t=1}^T$, is called an (ε, δ) approximate information state generator and $\{Z_t\}_{t=1}^T$ is called an approximate information state if the process $\{\hat{Z}_t\}_{t=1}^T$, where $\hat{Z}_t = \hat{\sigma}_t(H_t)$, satisfies the following properties:
(AP1) Sufficient for approximate performance evaluation, i.e., for any time t , any realization h_t of H_t and any choice \mathbf{a}_t of \mathbf{A}_t , we have

$$|\mathbb{E}[R_t \mid H_t = h_t, \mathbf{A}_t = \mathbf{a}_t] - \hat{r}_t(\hat{\sigma}_t(h_t), \mathbf{a}_t)| \leq \varepsilon$$

(AP2) Sufficient to predict itself approximately, i.e., for any time t , any realization h_t of H_t , any choice \mathbf{a}_t of \mathbf{A}_t , and for any Borel subset B of \hat{Z}_{t+1} , define $\mu_t(B) := \mathbb{P}(\hat{Z}_{t+1} \in B \mid H_t = h_t, \mathbf{A}_t = \mathbf{a}_t)$ and $\nu_t(B) := \hat{P}_t(B \mid \hat{\sigma}_t(h_t), \mathbf{a}_t)$; then

$$d_{\mathfrak{F}}(\mu_t, \nu_t) \leq \delta.$$

The theoretical upper bound for the (ε, δ) values remains uncertain due to the need for quantifying the expressiveness of the LSTM. As the state space, action space, and task horizon increase, the error caused by using FMC to estimate the information state will likely amplify, resulting in a more significant gap between the policy converged by F2A2 and the optimal policy. Therefore, our future research will focus on more efficient methods of estimating the information state based on existing work Mao et al. (2020).

Furthermore, the *infinite hierarchy of beliefs* in the POSG provides an alternative perspective that can be used to analyze the upper bound for the (ε, δ) values and the limitations of the Markovian policy, as discussed in the next section.

1.2 Infinite Hierarchy Belief Representation Capabilities

In the preceding section, our analysis begins with the information state. It aims to theoretically examine how the curse of dimensionality arises from the enriched observation space in POSG and its impact on the Markovian policy. Bayesian RL frameworks represent another approach to addressing partially observed problems that are discussed in the literature (Zamir, 2008; Ross et al., 2007; Katt et al., 2018; Foerster et al., 2018b). These frameworks

maintain a posterior distribution over environment models, where at each step, a model is sampled from the posterior, and its corresponding optimal policy is learned and executed.

In practice, it is reasonable for each agent in a POSG to compute a belief state that captures their uncertainty about the environment’s state, following the logic of POMDPs. In this context, such a belief state is called a ”zeroth-order belief,” as discussed in Moreno et al. (2021). When viewed from a single agent’s perspective, assuming that other agents’ policies are fixed, the POSG can be regarded as a POMDP. The presence of other agents transforms these agents into part of the environment, causing the unknown state of the world to encompass not only the environment state but also the knowledge state of other agents. Consequently, agent i must create a belief about the other agent’s beliefs, known as ”first-order” beliefs. This recursive process can continue, allowing agents to form ”second-order” beliefs about other agents’ ”first-order” beliefs, and so on. Ultimately, each agent maintains an ”infinite hierarchy” of beliefs.

It is apparent that F2A2 must possess the ability to express the infinite hierarchy belief if it uses finite memory compression and LSTM models to generate small (ε, δ) values. However, representing an infinite hierarchy of beliefs is theoretically unfeasible. Fortunately, agents do not need to consider an infinite hierarchy of beliefs to exhibit robust decision-making and generalization abilities. This fact is exemplified by human beings, with the Keynes Beauty Contest (Keynes, 1937) being a classic illustration. In a simplified version of this experiment, players are asked to select a number between 0 and 100, with the player whose guess is closest to half of the average declared the winner. According to the theory of Nash equilibrium (NE), selecting 0 is the only rational choice for each player in the Keynes Beauty Contest. This reasoning is as follows: ”If all players make random guesses, then the average of those guesses would be 50 (level-0). Therefore, I should guess at most $1/2 * 50 = 25$ (level- 1). If other players think similarly, I should not guess more than $1/2 * 25 = 13$ (level-2),” and so on. In this manner, the level of beliefs can continue to develop until all players choose 0, which is the only NE in this game.

However, experimental evidence reveals that most human players select numbers between 13 and 25 (Coricelli and Nagel, 2009), contradicting this theoretical result. This discrepancy is because not all human players exhibit perfect rationality; they are bounded by the levels of recursion they prefer to reason with (i.e., bounded rationality). Furthermore, in constrained scenarios, Wen et al. (2021) demonstrated that endowing agents with higher-level reasoning capabilities did not significantly enhance performance.

To this end, we can comprehend the limitations of Markovian policies based on the previous analysis by examining whether finite memory compression (FMC) and LSTM can represent a finite hierarchy of beliefs. In this regard, Moreno et al. (2021) introduces a scalable approach to approximate hierarchical belief structures using recursive deep generative models and leverages these belief models to obtain representations useful for complex tasks. For belief approximation at each level, they employ a Markovian policy based on finite memory compression and two independent RNNs. This implies that multi-layer RNNs may be capable of representing a finite hierarchy of beliefs. However, in F2A2, only a single-layer LSTM is used in the Markovian policy, which likely imposes significant limitations on representing multi-layer beliefs and leads to large (ε, δ) values. Therefore, in future work, we intend to conduct more comprehensive research on designing a more effective network structure to estimate multi-layer beliefs.

References

- Martín Abadi, Ashish Agarwal, Paul Barham, Eugene Brevdo, Zhifeng Chen, Craig Citro, Greg S. Corrado, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Ian Goodfellow, Andrew Harp, Geoffrey Irving, Michael Isard, Yangqing Jia, et al. TensorFlow: Large-scale machine learning on heterogeneous systems, 2015. URL <http://tensorflow.org/>. Software available from tensorflow.org.
- Jeffrey L Adler and Victor J Blue. A cooperative multi-agent transportation management and route guidance system. *Transportation Research Part C: Emerging Technologies*, 10(5-6):433–454, 2002.
- Gürdal Arslan and S. Yüksel. Decentralized q-learning for stochastic teams and games. *IEEE Transactions on Automatic Control*, 62:1545–1558, 2017.
- Kamyar Azizzadenesheli, Alessandro Lazaric, and Animashree Anandkumar. Reinforcement learning of pomdps using spectral methods. In *CLT*, 2016.
- Kamyar Azizzadenesheli, Yisong Yue, and Animashree Anandkumar. Policy gradient in partially observable environments: Approximation and convergence. *arXiv preprint arXiv:1810.07900*, 2018.
- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. Neural machine translation by jointly learning to align and translate. In *ICLR*, 2015.
- Etienne Barnard. Temporal-difference methods and markov models. *IEEE Transactions on Systems, Man, and Cybernetics*, 23(2):357–365, 1993.
- Jonathan Baxter and Peter L Bartlett. Infinite-horizon policy-gradient estimation. *Journal of Artificial Intelligence Research*, 15:319–350, 2001.
- Jonathan Baxter, Peter L Bartlett, et al. Reinforcement learning in pomdp’s via direct gradient ascent. In *ICML*, 2000.
- Stephen Boyd, Neal Parikh, Eric Chu, Borja Peleato, and Jonathan Eckstein. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends® in Machine Learning*, 3(1):1–122, 2011.
- Ronen I Brafman and Moshe Tennenholtz. Learning to coordinate efficiently: A model-based approach. *Journal of Artificial Intelligence Research*, 19:11–23, 2003.
- L. Cassano, S. Alghunaim, and A. Sayed. Team policy learning for multi-agent reinforcement learning. In *ICASSP*, 2019.
- Lucas Cassano, Kun Yuan, and Ali H. Sayed. Multi-agent fully decentralized off-policy learning with linear convergence rates. *arXiv preprint arXiv:1810.07792*, 2018.
- Pablo Samuel Castro, Prakash Panangaden, and Doina Precup. Equivalence relations in fully and partially observable markov decision processes. In *IJCAI*, 2009.

- Antonin Chambolle and Thomas Pock. A first-order primal-dual algorithm for convex problems with applications to imaging. *Journal of Mathematical Imaging and Vision*, 40(1):120–145, 2011.
- Yu-Han Chang, Tracey Ho, and Leslie P Kaelbling. All learning is local: Multi-agent learning in global reward games. In *NeurIPS*, 2004.
- Filippos Christianos, Georgios Papoudakis, A. Rahman, and Stefano V. Albrecht. Scaling multi-agent reinforcement learning with selective parameter sharing. In *ICML*, 2021.
- Tianshu Chu, Sandeep Chinchali, and S. Katti. Multi-agent reinforcement learning for networked system control. In *ICLR*, 2020.
- Kamil Ciosek and Shimon Whiteson. Expected policy gradients. In *AAAI*, 2018.
- Caroline Claus and Craig Boutilier. The dynamics of reinforcement learning in cooperative multiagent systems. In *AAAI*, 1998.
- Giorgio Coricelli and Rosemarie Nagel. Neural correlates of depth of strategic reasoning in medial prefrontal cortex. *Proceedings of the National Academy of Sciences*, 106(23):9163–9168, 2009.
- Bo Dai, Albert Shaw, Lihong Li, Lin Xiao, Niao He, Zhen Liu, Jianshu Chen, and Le Song. SBEED: Convergent reinforcement learning with nonlinear function approximation. In *ICML*, 2018.
- Prafulla Dhariwal, Christopher Hesse, Oleg Klimov, Alex Nichol, Matthias Plappert, Alec Radford, John Schulman, Szymon Sidor, Yuhuai Wu, and Peter Zhokhov. Openai baselines. <https://github.com/openai/baselines>, 2017.
- Thinh Doan, Siva Maguluri, and Justin Romberg. Finite-time analysis of distributed TD(0) with linear function approximation on multi-agent reinforcement learning. In *ICML*, 2019a.
- Thinh T Doan, Siva Theja Maguluri, and Justin Romberg. Convergence rates of distributed TD(0) with linear function approximation for multi-agent reinforcement learning. *arXiv preprint arXiv:1902.07393*, 2019b.
- Jesse Farebrother, Marlos C Machado, and Michael Bowling. Generalization and regularization in DQN. *arXiv preprint arXiv:1810.00123*, 2018.
- Jakob Foerster, Ioannis Alexandros Assael, Nando de Freitas, and Shimon Whiteson. Learning to communicate with deep multi-agent reinforcement learning. In *NeurIPS*, 2016.
- Jakob N Foerster, Gregory Farquhar, Triantafyllos Afouras, Nantas Nardelli, and Shimon Whiteson. Counterfactual multi-agent policy gradients. In *AAAI*, 2018a.
- Jakob N. Foerster, H. Francis Song, Edward Hughes, Neil Burch, Iain Dunning, Shimon Whiteson, Matthew M. Botvinick, and Michael H. Bowling. Bayesian action decoder for deep multi-agent reinforcement learning. In *ICML*, 2018b.

- Scott Fujimoto, Herke Hoof, and David Meger. Addressing function approximation error in actor-critic methods. In *ICML*, 2018.
- Niko A. Grupen, B. Selman, and Daniel D. Lee. Fairness for cooperative multi-agent learning with equivariant policies. *ArXiv*, abs/2106.05727, 2021.
- Owen Guillermo. *Game Theory: Second Edition*. Academic Press, Orlando, Florida, 1982. ISBN 0125311508.
- Jayesh K Gupta, Maxim Egorov, and Mykel Kochenderfer. Cooperative multi-agent control using deep reinforcement learning. In *AAMAS*, 2017.
- Tuomas Haarnoja, Haoran Tang, Pieter Abbeel, and Sergey Levine. Reinforcement learning with deep energy-based policies. In *ICML*, 2017.
- Tuomas Haarnoja, Aurick Zhou, Kristian Hartikainen, George Tucker, Sehoon Ha, Jie Tan, Vikash Kumar, Henry Zhu, Abhishek Gupta, Pieter Abbeel, and Sergey Levine. Soft actor-critic algorithms and applications. *arXiv preprint arXiv:1812.05905*, 2018.
- Joseph Y Halpern and Yoram Moses. Knowledge and common knowledge in a distributed environment. *Journal of the ACM*, 37(3):549–587, 1990.
- Eric A Hansen, Daniel S Bernstein, and Shlomo Zilberstein. Dynamic programming for partially observable stochastic games. In *AAAI*, 2004.
- Matthew Hausknecht and Peter Stone. Deep recurrent Q-learning for partially observable MDPs. In *AAAI Fall Symposium Series*, 2015.
- Fritz Heider and Marianne Simmel. An experimental study of apparent behavior. *The American Journal of Psychology*, 57(2):243–259, 1944.
- Mingyi Hong, Zhi-Quan Luo, and Meisam Razaviyayn. Convergence analysis of alternating direction method of multipliers for a family of nonconvex problems. *SIAM Journal on Optimization*, 26(1):337–364, 2016.
- Mingyi Hong, Hoi-To Wai, Zhaoran Wang, and Zhuoran Yang. A two-timescale framework for bilevel optimization: Complexity analysis and application to actor-critic. *ArXiv*, abs/2007.05170, 2020.
- Jueming Hu, Zhe Xu, Weichang Wang, Guannan Qu, Yutian Pang, and Yongming Liu. Decentralized graph-based multi-agent reinforcement learning using reward machines. *arXiv preprint arXiv:2110.00096*, 2021.
- Shariq Iqbal and Fei Sha. Actor-attention-critic for multi-agent reinforcement learning. In *ICML*, 2019.
- Eric Jang, Shixiang Gu, and Ben Poole. Categorical reparametrization with gumble-softmax. In *ICLR*, 2017.

- Natasha Jaques, Angeliki Lazaridou, Edward Hughes, Çağlar Gülçehre, Pedro A Ortega, DJ Strouse, Joel Z Leibo, and Nando de Freitas. Social influence as intrinsic motivation for multi-agent deep reinforcement learning. In *ICML*, 2019.
- Stanislaw Jastrzebski, Zachary Kenton, Devansh Arpit, Nicolas Ballas, Asja Fischer, Yoshua Bengio, and Amos Storkey. Three factors influencing minima in SGD. *arXiv preprint arXiv:1711.04623*, 2017.
- Jiechuan Jiang and Zongqing Lu. Learning attentional communication for multi-agent cooperation. In *NeurIPS*, 2018.
- Jiechuan Jiang, Chen Dun, and Zongqing Lu. Graph convolutional reinforcement learning for multi-agent cooperation. *arXiv preprint arXiv:1810.09202*, 2018.
- Chi Jin, Sham Kakade, Akshay Krishnamurthy, and Qinghua Liu. Sample-efficient reinforcement learning of undercomplete pomdps. In *NeurIPS*, 2020.
- Spiros Kapetanakis and D. Kudenko. Reinforcement learning of coordination in cooperative multi-agent systems. In *AAAI*, 2002.
- Soumya Kar, José MF Moura, and H Vincent Poor. QD-learning: A collaborative distributed strategy for multi-agent reinforcement learning through consensus+innovations. *IEEE Transactions on Signal Processing*, 61(7):1848–1862, 2013.
- Sammie Katt, Frans A. Oliehoek, and Chris Amato. Bayesian reinforcement learning in factored pomdps. In *AAMAS*, 2018.
- John Maynard Keynes. The general theory of employment. *The quarterly journal of economics*, 51(2):209–223, 1937.
- Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *ICLR*, 2015.
- Vijay R Konda and John N Tsitsiklis. Actor-critic algorithms. In *NeurIPS*, 2000.
- Akshay Krishnamurthy, Alekh Agarwal, and John Langford. Pac reinforcement learning with rich observations. In *NeurIPS*, 2016.
- M. Lauer and Martin A. Riedmiller. An algorithm for distributed reinforcement learning in cooperative multi-agent systems. In *ICML*, 2000.
- Donghwan Lee, Hyungjin Yoon, and Naira Hovakimyan. Primal-dual algorithm for distributed reinforcement learning: Distributed GTD. In *CDC*, 2018.
- Sergey Levine. Reinforcement learning and control as probabilistic inference: Tutorial and review. *arXiv preprint arXiv:1805.00909*, 2018.
- Shihui Li, Yi Wu, Xinyue Cui, Honghua Dong, Fei Fang, and Stuart Russell. Robust multi-agent reinforcement learning via minimax deep deterministic policy gradient. In *AAAI*, 2019.

- Timothy P Lillicrap, Jonathan J Hunt, Alexander Pritzel, Nicolas Heess, Tom Erez, Yuval Tassa, David Silver, and Daan Wierstra. Continuous control with deep reinforcement learning. In *ICLR*, 2016.
- M. Littman. Value-function reinforcement learning in markov games. *Cognitive Systems Research*, 2:55–66, 2001.
- Michael L Littman. Markov games as a framework for multi-agent reinforcement learning. In *ICML*. 1994a.
- Michael L Littman. Memoryless policies: Theoretical limitations and practical results. In *From Animals to Animats 3: Proceedings of the third international conference on simulation of adaptive behavior*, volume 3, page 238, 1994b.
- Qinghua Liu, Alan Chung, Csaba Szepesvári, and Chi Jin. When is partially observable reinforcement learning not scary? In *COLT*, 2022.
- Zhuang Liu, Xuanlin Li, Bingyi Kang, and Trevor Darrell. Regularization matters in policy optimization. *arXiv preprint arXiv:1910.09191*, 2019.
- Ryan Lowe, Yi Wu, Aviv Tamar, Jean Harb, OpenAI Pieter Abbeel, and Igor Mordatch. Multi-agent actor-critic for mixed cooperative-competitive environments. In *NeurIPS*, 2017.
- Songtao Lu, Kaiqing Zhang, Tianyi Chen, Tamer Basar, and Lior Horesh. Decentralized policy gradient descent ascent for safe multi-agent reinforcement learning. In *AAAI*, 2021.
- Sergio Macua, Jianshu Chen, Santiago Zazo, and Ali Sayed. Distributed policy evaluation under multiple behavior strategies. *IEEE Transactions on Automatic Control*, 60(5): 1260–1274, 2015.
- Omid Madani, Steve Hanks, and Anne Condon. On the undecidability of probabilistic planning and infinite-horizon partially observable markov decision problems. In *AAAI*, 1999.
- Weichao Mao, Kaiqing Zhang, Erik Miehling, and Tamer Başar. Information state embedding in partially observable cooperative multi-agent reinforcement learning. In *CDC*, 2020.
- Laëtitia Matignon, G. Laurent, and N. L. Fort-Piat. Hysteretic q-learning :an algorithm for decentralized reinforcement learning in cooperative multi-agent teams. *2007 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 64–69, 2007.
- Laëtitia Matignon, Laurent Jeanpierre, and Abdel-illah Mouaddib. Coordinated multi-robot exploration under communication constraints using decentralized Markov decision processes. In *AAAI*, 2012.
- Guido Montufar, Keyan Ghazi-Zahedi, and Nihat Ay. Geometry and determinism of optimal stationary control in partially observable markov decision processes. *arXiv preprint arXiv:1503.07206*, 2015.

- Pol Moreno, Edward Hughes, Kevin R McKee, Bernardo Avila Pires, and Théophane Weber. Neural recursive belief states in multi-agent reinforcement learning. *arXiv preprint arXiv:2102.02274*, 2021.
- Rémi Munos, Tom Stepleton, Anna Harutyunyan, and Marc Bellemare. Safe and efficient off-policy reinforcement learning. In *NeurIPS*, 2016.
- Arvind Neelakantan, Luke Vilnis, Quoc V Le, Ilya Sutskever, Lukasz Kaiser, Karol Kurach, and James Martens. Adding gradient noise improves learning for very deep networks. *ICLR*, 2017.
- Chris Nota and Philip S Thomas. Is the policy gradient a gradient? In *AAMAS*, 2020.
- Frans A Oliehoek, Matthijs TJ Spaan, and Nikos Vlassis. Optimal and approximate Q-value functions for decentralized POMDPs. *Journal of Artificial Intelligence Research*, 32:289–353, 2008.
- Timothy O’Connor. Emergent Properties. In *The Stanford Encyclopedia of Philosophy*. 2021.
- Liviu Panait, K. Tuyls, and S. Luke. Theoretical advantages of lenient learners: An evolutionary game theoretic perspective. *J. Mach. Learn. Res.*, 9:423–457, 2008.
- Christos H Papadimitriou and John N Tsitsiklis. The complexity of markov decision processes. *Mathematics of operations research*, 12(3):441–450, 1987.
- Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in PyTorch. 2017.
- Peng Peng, Quan Yuan, Ying Wen, Yaodong Yang, Zhenkun Tang, Haitao Long, and Jun Wang. Multiagent bidirectionally-coordinated nets for learning to play Starcraft combat games. *arXiv preprint arXiv:1703.10069*, 2017.
- Martin L Puterman. *Markov decision processes: discrete stochastic dynamic programming*. John Wiley & Sons, 2014.
- Chao Qu, Shie Mannor, Huan Xu, Yuan Qi, Le Song, and Junwu Xiong. Value propagation for decentralized networked deep multi-agent reinforcement learning. In *NeurIPS*, 2019.
- Guannan Qu and Na Li. Harnessing smoothness to accelerate distributed optimization. *IEEE Transactions on Control of Network Systems*, 5(3):1245–1260, 2017.
- Guannan Qu, Adam Wierman, and Na Li. Scalable reinforcement learning for multiagent networked systems. *Operations Research*, 2022.
- Neil C Rabinowitz, Frank Perbet, H Francis Song, Chiyuan Zhang, SM Eslami, and Matthew Botvinick. Machine theory of mind. In *ICML*, 2018.

- Antonin Raffin, Ashley Hill, Adam Gleave, Anssi Kanervisto, Maximilian Ernestus, and Noah Dormann. Stable-baselines3: Reliable reinforcement learning implementations. *Journal of Machine Learning Research*, 22(268):1–8, 2021.
- Tabish Rashid, Mikayel Samvelyan, Christian Schroeder de Witt, Gregory Farquhar, Jakob Foerster, and Shimon Whiteson. QMIX: Monotonic value function factorisation for deep multi-agent reinforcement learning. In *ICML*, 2018.
- Amir Rasouli, Iuliia Kotseruba, and John K Tsotsos. Agreeing to cross: How drivers and pedestrians communicate. In *Intelligent Vehicles Symposium*, 2017.
- D Reddy, Amrita Saha, Srikanth G Tamilselvam, et al. Risk averse reinforcement learning for mixed multi-agent environments. In *AAMAS*, 2019.
- Stéphane Ross, Brahim Chaib-draa, and Joelle Pineau. Bayes-adaptive pomdps. In *NeurIPS*, 2007.
- Heechang Ryu, H. Shin, and Jinkyoo Park. Multi-agent actor-critic with generative cooperative policy network. *ArXiv*, abs/1810.09206, 2018.
- Mikayel Samvelyan, Tabish Rashid, Christian Schroeder de Witt, Gregory Farquhar, Nantas Nardelli, Tim GJ Rudner, Chia-Man Hung, Philip HS Torr, Jakob Foerster, and Shimon Whiteson. The Starcraft multi-agent challenge. In *AAMAS*, 2019.
- Ali H Sayed. Adaptive networks. *Proceedings of the IEEE*, 102(4):460–497, 2014.
- Franco Scarselli, Marco Gori, Ah Chung Tsoi, Markus Hagenbuchner, and Gabriele Monfardini. The graph neural network model. *IEEE Transactions on Neural Networks*, 20(1):61–80, 2008.
- Mark Schmidt, Nicolas Le Roux, and Francis Bach. Minimizing finite sums with the stochastic average gradient. *Mathematical Programming*, 162(1-2):83–112, 2017.
- Christian Schroeder de Witt, Jakob Foerster, Gregory Farquhar, Philip Torr, Wendelin Boehmer, and Shimon Whiteson. Multi-agent common knowledge reinforcement learning. In *NeurIPS*, 2019.
- John Schulman, Sergey Levine, Pieter Abbeel, Michael Jordan, and Philipp Moritz. Trust region policy optimization. In *ICML*, 2015.
- John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- Junjie Sheng, Xiangfeng Wang, Bo Jin, Junchi Yan, Wenhao Li, Tsung-Hui Chang, Jun Wang, and Hongyuan Zha. Learning structured communication for multi-agent reinforcement learning. *arXiv preprint arXiv:2002.04235*, 2020.
- Satinder P Singh, Tommi Jaakkola, and Michael I Jordan. Learning without state-estimation in partially observable markovian decision processes. In *ICML*. 1994.

- Leslie N Smith and Nicholay Topin. Super-convergence: Very fast training of neural networks using large learning rates. In *SPIE*, 2019.
- Kyunghwan Son, Daewoo Kim, Wan Ju Kang, David Earl Hostallero, and Yung Yi. QTRAN: Learning to factorize with transformation for cooperative MARL. In *ICML*, 2019.
- Yuhang Song, Jianyi Wang, Thomas Lukasiewicz, Zhenghua Xu, Mai Xu, Zihan Ding, and Lianlong Wu. Arena: A general evaluation platform and building toolkit for multi-agent intelligence. In *AAAI*, 2020.
- Miloš S Stanković and Srdjan S Stanković. Multi-agent temporal-difference learning with linear function approximation: Weak convergence under time-varying network topologies. In *ACC*, 2016.
- Jayakumar Subramanian, Amit Sinha, Raihan Seraj, and Aditya Mahajan. Approximate information state for approximate planning and reinforcement learning in partially observed systems. *Journal of Machine Learning Research*, 23:12–1, 2022.
- Sainbayar Sukhbaatar, Rob Fergus, et al. Learning multiagent communication with back-propagation. In *NeurIPS*, 2016.
- Wesley Suttle, Zhuoran Yang, Kaiqing Zhang, Zhaoran Wang, Tamer Basar, and Ji Liu. A multi-agent off-policy actor-critic algorithm for distributed reinforcement learning. *arXiv preprint arXiv:1903.06372*, 2019.
- Richard S Sutton and Andrew G Barto. *Reinforcement Learning: An Introduction*. MIT press, 2018.
- Richard S Sutton, David A McAllester, Satinder P Singh, and Yishay Mansour. Policy gradient methods for reinforcement learning with function approximation. In *NeurIPS*, 2000.
- Richard S Sutton, Hamid Reza Maei, Doina Precup, Shalabh Bhatnagar, David Silver, Csaba Szepesvári, and Eric Wiewiora. Fast gradient-descent methods for temporal-difference learning with linear function approximation. In *ICML*, 2009.
- Ming Tan. Multi-agent reinforcement learning: Independent vs. cooperative agents. In *ICML*, 1993.
- Hamidreza Tavaafoghi, Yi Ouyang, and Demosthenis Teneketzis. A unified approach to dynamic decision problems with asymmetric information: Nonstrategic agents. *IEEE Transactions on Automatic Control*, 67(3):1105–1119, 2021.
- Justin K. Terry, Nathaniel Grammel, Ananth Hari, Luis Santos, and Benjamin Black. Revisiting parameter sharing in multi-agent deep reinforcement learning. *arXiv: Learning*, 2020.
- G. Tesauro. Extending q-learning to general adaptive multi-agent systems. In *NIPS*, 2003.

- Gerald Tesauro. Temporal difference learning and TD-Gammon. *Communications of the ACM*, 38(3):58–68, 1995.
- Zheng Tian, Shihao Zou, Tim Warr, Lisheng Wu, Haitham Bou Ammar, and Jun Wang. Learning to communicate implicitly by actions. *arXiv preprint arXiv:1810.04444*, 2018.
- Nikos Vlassis, Michael L Littman, and David Barber. On the computational complexity of stochastic controller optimization in pomdps. *ACM Transactions on Computation Theory*, 4(4):1–8, 2012.
- Hoi-To Wai, Zhuoran Yang, Zhaoran Wang, and Mingyi Hong. Multi-agent reinforcement learning via double averaging primal-dual optimization. In *NeurIPS*, 2018.
- L. Wang, Zhuoran Yang, and Zhaoran Wang. Breaking the curse of many agents: Provable mean embedding q-iteration for mean-field reinforcement learning. In *ICML*, 2020.
- X. Wang and T. Sandholm. Reinforcement learning to play an optimal nash equilibrium in team markov games. In *NIPS*, 2002.
- Ermo Wei, Drew Wicke, David Freelan, and Sean Luke. Multiagent soft Q-learning. In *AAAI Spring Symposium Series*, 2018.
- Ying Wen, Yaodong Yang, and Jun Wang. Modelling bounded rationality in multi-agent interactions by generalized recursive reasoning. In *IJCAI*, 2021.
- David H Wolpert and Kagan Tumer. Optimal payoff functions for members of collectives. In *Modeling complexity in economic and social systems*, pages 355–369. World Scientific, 2002.
- Y. Yang, R. Luo, Minne Li, M. Zhou, W. Zhang, and J. Wang. Mean field multi-agent reinforcement learning. In *ICML*, 2018a.
- Yaodong Yang, Rui Luo, Minne Li, Ming Zhou, Weinan Zhang, and Jun Wang. Mean field multi-agent reinforcement learning. In *ICML*, 2018b.
- Zhuoran Yang, Zuyue Fu, K. Zhang, and Zhaoran Wang. Convergent reinforcement learning with function approximation: A bilevel optimization perspective. In *Manuscript*, 2018c.
- Zhuoran Yang, Yongxin Chen, M. Hong, and Zhaoran Wang. Provably global convergence of actor-critic: A case for linear quadratic regulator with ergodic cost. In *NeurIPS*, 2019.
- Bora Yongacoglu, Gürdal Arslan, and S. Yüksel. Decentralized learning for optimality in stochastic dynamic teams and games with local control and global state information. In *CDC*, 2019.
- Shmuel Zamir. Bayesian games: Games with incomplete information. In *Encyclopedia of Complexity and Systems Science*, 2008.
- Sihan Zeng, Thinh T. Doan, and J. Romberg. Finite-time analysis of decentralized stochastic approximation with applications in multi-agent and multi-task learning. *ArXiv*, abs/2010.15088, 2020.

- K. Zhang, Zhuoran Yang, H. Liu, T. Zhang, and Tamer Başar. Finite-sample analysis for decentralized batch multi-agent reinforcement learning with networked agents. *arXiv: Learning*, 2018a.
- K. Zhang, Zhuoran Yang, and T. Başar. Multi-agent reinforcement learning: A selective overview of theories and algorithms. *ArXiv*, abs/1911.10635, 2019.
- Kaiqing Zhang, Zhuoran Yang, and Tamer Basar. Networked multi-agent reinforcement learning in continuous spaces. In *CDC*, 2018b.
- Kaiqing Zhang, Zhuoran Yang, Han Liu, Tong Zhang, and Tamer Basar. Fully decentralized multi-agent reinforcement learning with networked agents. In *ICML*, 2018c.
- Yan Zhang and Michael M Zavlanos. Distributed off-policy actor-critic reinforcement learning with policy consensus. *arXiv preprint arXiv:1903.09255*, 2019.
- Lianmin Zheng, Jiacheng Yang, Han Cai, Ming Zhou, Weinan Zhang, et al. MAgent: A many-agent reinforcement learning platform for artificial collective intelligence. In *AAAI*, 2018.
- Vladimir Mikhailovich Zolotarev. Probability metrics. *Teoriya Veroyatnostei i ee Prime-neniya*, 28(2):264–287, 1983.