

Cooperative and Competitive Biases for Multi-Agent Reinforcement Learning

Heechang Ryu
KAIST
Daejeon, Republic of Korea
rhc93@kaist.ac.kr

Hayong Shin
KAIST
Daejeon, Republic of Korea
hyshin@kaist.ac.kr

Jinkyoo Park*
KAIST
Daejeon, Republic of Korea
jinkyoo.park@kaist.ac.kr

ABSTRACT

Training a multi-agent reinforcement learning (MARL) algorithm is more challenging than training a single-agent reinforcement learning algorithm, because the result of a multi-agent task strongly depends on the complex interactions among agents and their interactions with a stochastic and dynamic environment. We propose an algorithm that boosts MARL training using the biased action information of other agents based on a friend-or-foe concept. For a cooperative and competitive environment, there are generally two groups of agents: cooperative-agents and competitive-agents. In the proposed algorithm, each agent updates its value function using its own action and the biased action information of other agents in the two groups. The biased joint action of cooperative agents is computed as the sum of their actual joint action and the imaginary cooperative joint action, by assuming all the cooperative agents jointly maximize the target agent's value function. The biased joint action of competitive agents can be computed similarly. Each agent then updates its own value function using the biased action information, resulting in a biased value function and corresponding biased policy. Subsequently, the biased policy of each agent is inevitably subjected to recommend an action to cooperate and compete with other agents, thereby introducing more active interactions among agents and enhancing the MARL policy learning. We empirically demonstrate that our algorithm outperforms existing algorithms in various mixed cooperative-competitive environments. Furthermore, the introduced biases gradually decrease as the training proceeds and the correction based on the imaginary assumption vanishes.

KEYWORDS

Multi-Agent Reinforcement Learning; Cooperation; Competition; Bias

ACM Reference Format:

Heechang Ryu, Hayong Shin, and Jinkyoo Park*. 2021. Cooperative and Competitive Biases for Multi-Agent Reinforcement Learning. In *Proc. of the 20th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2021)*, Online, May 3–7, 2021, IFAAMAS, 9 pages.

1 INTRODUCTION

Reinforcement learning (RL) algorithms solve sequential decision-making problems using experiences obtained by a single agent (decision maker) dynamically interacting with an environment. The RL algorithms typically estimate an action-value function (Q -function)

*Corresponding author

Proc. of the 20th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2021), U. Endriss, A. Nowé, F. Dignum, A. Lomuscio (eds.), May 3–7, 2021, Online. © 2021 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

or a decision-making policy, using various function approximators (e.g., deep neural networks), to model how a particular action (decision) affects future outcomes. From this model, the optimal action in the current state for completing a target task can be deduced [17]. The success of RL algorithms for solving various tasks depends on how effectively they learn such temporal interactions between an action and a future outcome.

In particular, when the RL algorithm is trained with a sparse or delayed reward, *i.e.*, a reward signal is infrequently realized after an agent executes action, it becomes difficult to estimate the Q -function or policy. This is because it is challenging for the RL agent to learn the dynamic causal effect of its action on future outcomes, owing to the sparse and delayed reward signal. A representative example of such a task with a sparse reward is a goal-oriented task, in which a binary reward is given only when an agent reaches a goal. To solve this task, HER [1] has been proposed to learn the Q -function or policy using the reward signals obtained from failed tasks (episodes), while considering these reward signals as being obtained from the successful tasks that are different from the original task. This strategy of HER transforms the sparse-reward environment into a dense-reward environment, enabling the RL agent to easily learn the Q -function or policy. In addition, HPG [13] extends the concept of HER to efficiently generalize learning about different goals using information obtained by the current policy for a specific goal.

Recently, to enhance the performance of multi-agent reinforcement learning (MARL) algorithms, which are extensions of the RL algorithms to multi-agent settings, various strategies have been proposed. Some researchers have proposed an intrinsic reward to induce certain collective behaviors of multiple agents that are believed to help achieve the objective of a multi-agent task. For example, intrinsic reward is designed to promote the agents to execute actions influencing other agents' state transitions [2, 18] or visit unexplored state spaces more frequently [6]. However, in general, designing a good intrinsic reward is difficult because it requires prior knowledge of the types of interactions that help solve the multi-agent task. Moreover, designing an effective intrinsic reward often requires an iterative reward-shaping procedure until satisfactory performance is reached. Consequently, for multi-agent tasks where a sparse reward is given and prior knowledge is not available, another effective method for MARL should be developed.

To address these difficulties, we herein propose an algorithm, called Friend-or-Foe multi-agent Deep Deterministic Policy Gradient (F2DDPG), which boosts MARL training using the biased action information of other agents based on a friend-or-foe concept [11]. F2DDPG adopts an actor-critic algorithm in a centralized training

and decentralized execution (CTDE) framework [12]. For a cooperative and competitive environment, there are generally two groups of agents: an ally group composed of cooperative agents to a target agent and an enemy group composed of competitive agents to a target agent. In the proposed algorithm, each agent updates its Q -function (critic) using its own action and the biased action information of other agents in the two groups. The biased joint action of cooperative agents is computed as the sum of their actual joint action and the imaginary cooperative joint action obtained by assuming that all the cooperative agents jointly maximize the target agent’s critic. The biased joint action of competitive agents can be computed similarly. Each agent then updates its own critic using biased action information, resulting in a biased critic and corresponding biased policy. Thereafter, the biased policy of each agent is inevitably subjected to recommend an action to cooperate and compete with other agents, which introduces more active interactions among agents, and thus, enhances the MARL policy learning.

Using biased actions in estimating a critic can be viewed as a different version of using the biased reward (intrinsic reward) in the estimation to induce the intended interactions among agents. However, we do not compare our method with MARL approaches using the intrinsic reward because they require a certain type of prior knowledge about the environment (game). Instead, we compare F2DDPG with M3DDPG [8], the latter of which uses modified (biased) action information of other agents when training the critic; it modifies the other agents’ actions in an adversarial manner to induce a robust policy. Empirically, we demonstrate that our algorithm outperforms existing algorithms in four mixed cooperative-competitive game scenarios, in which the agents have cooperative and competitive interactions [12]. Furthermore, we empirically show that the introduced biases gradually decrease as the training proceeds and that the correction based on the imaginary assumption vanishes.

2 RELATED WORK

2.1 Friend-or-Foe Q-Learning

In general-sum games, such as mixed cooperative-competitive games, friend-or-foe Q-learning (FFQ) [11] has been proposed to provide strong convergence guarantees compared to the existing Nash-equilibrium-based learning rule [4]. It requires that other agents be identified as either ‘friend’ (ally) or ‘foe’ (enemy). FFQ then assumes that agent i ’s friends are working together to maximize agent i ’s value, while agent i ’s foes are working together to minimize agent i ’s value. Thus, n -player FFQ considers any game as a two-player zero-sum game with an extended action set, and is easy to implement for multiple agents.

2.2 Biases in RL and MARL

In RL and MARL, various forms of inductive bias have been used to improve learning. The most straightforward inductive biases entail designing network structures for the critic or policy, such as attention networks [5], graph neural networks [15], and implicit communication structures [14]. However, biases in game information, such as state, reward, and action have also been used in an

attempt to boost training. We review the biases in the information in this subsection.

2.2.1 Biases in States. Bias has been reported to help train RL by injecting a biased belief regarding the state at the initialization stage of the Q -table [3]. For example, in a goal-oriented task, if the goal is known in advance, biased information about the state, near and far from the goal, is injected into the Q -table before training. In addition, a distributed Q -learning algorithm for a cooperative multi-agent setting has been proposed based on the optimistic assumption [7]. Under this assumption, the algorithm biasedly updates the Q -table only when the new value for Q is greater than the current value in the current state.

2.2.2 Biases in Rewards. Intrinsic rewards have been proposed as a bias for multi-agent exploration to induce certain collective behaviors of agents, based on prior knowledge of the types of interactions that help solve the multi-agent task. The intrinsic rewards are provided when one agent’s action affects the state transitions of other agents [2, 18] and when all the agents explore only different or the same areas for the task of collecting scattered treasures [6].

2.2.3 Biases in Actions. M3DDPG [8] has been proposed to learn a robust policy using other agents’ action information corrupted with adversarial noise. In this approach, each agent assumes that other agents provide the adversarial actions to the target agent, and updates its critic using such adversarial action information. To compute the adversarial action (biased action), each agent modifies the actions of other agents in the direction that minimizes the target agent’s critic. M3DDPG asserts that the trained policy using this biased action information outperforms its baseline algorithm, MADDPG [12]. However, the limitation of this approach is that it does not consider the relationships among agents. It assumes that all the agents are adversarial to the target agent, regardless of whether they are allies or enemies in the cooperative and competitive game; this assumption is inconsistent with the actual situation.

Notably, M3DDPG is similar to our method as it uses the biased information of other agents’ actions. However, F2DDPG explicitly considers the roles of agents in a cooperative and competitive environment, in which the cooperative and competitive agents are known a priori. In addition, our method can be justified by the well-known FFQ [11]. We mainly compare the performance of our method with that of M3DDPG, while not addressing other approaches using different types of information bias. In this study, we assume that it is possible to identify other agents as being either cooperative (allies) or competitive (enemies) in mixed cooperative-competitive environments, to fully implement and compare the proposed method with the baseline methods.

3 BACKGROUND

3.1 Partially Observable Markov Game

We consider a partially observable Markov game [10], which is an extension of the partially observable Markov decision process to a game with multiple agents. A partially observable Markov game for N agents is defined as follows: $s \in \mathcal{S}$ denotes the global state of the game; $o_i \in \mathcal{S} \mapsto \mathcal{O}_i$ denotes a local observation that agent i can acquire correlated with the state; and $a_i \in \mathcal{A}_i$ is an action of

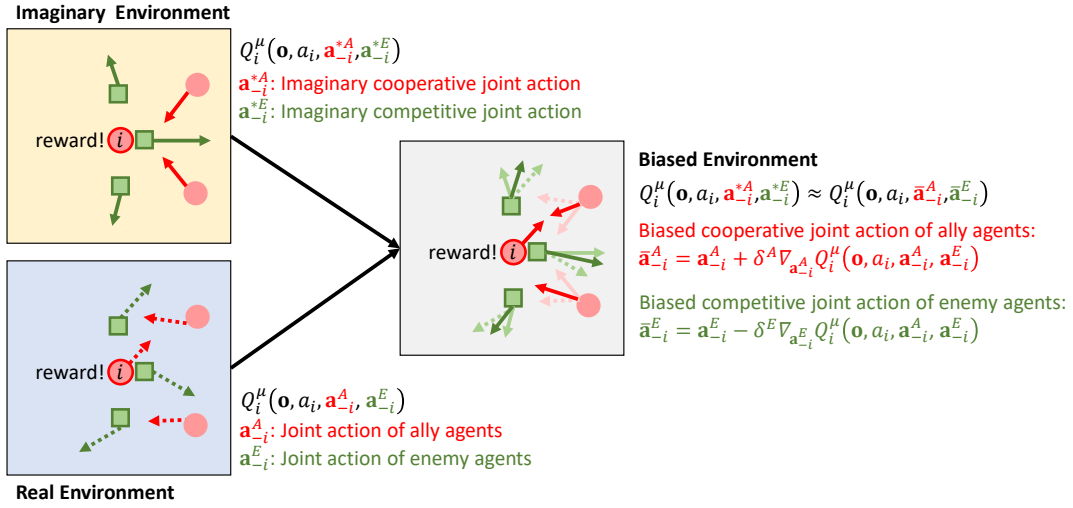


Figure 1: Overview of F2DDPG.

agent i . The reward for agent i is obtained as a function of state s and joint action \mathbf{a} as $r_i : \mathcal{S} \times \mathcal{A}_1 \times \dots \times \mathcal{A}_N \mapsto \mathbb{R}$. The state evolves to the next state according to the state transition function $\mathcal{T} : \mathcal{S} \times \mathcal{A}_1 \times \dots \times \mathcal{A}_N \mapsto \mathcal{S}$. The initial state is determined by the initial state distribution $\rho : \mathcal{S} \mapsto [0, 1]$. Agent i aims to maximize its discounted return $R_i = \sum_{t=0}^T \gamma^t r_{i,t}$, where $\gamma \in [0, 1]$ is a discount factor.

3.2 Multi-Agent Deep Deterministic Policy Gradient (MADDPG)

While the policy can be deterministic, $a = \mu(s)$, or stochastic, $a \sim \pi(\cdot|s)$, deterministic policy gradient (DPG) [16] for RL adopts a deterministic policy. DPG aims to directly derive the deterministic policy, $a = \mu(s; \theta)$, that maximizes the expected return or objective $\mathcal{J}(\theta) = \mathbb{E}_{s \sim \rho^\mu, a \sim \mu_\theta} [R] \approx \mathbb{E}_{s \sim \rho^\mu, a \sim \mu_\theta} [Q^\mu(s, a; \phi)]$, where $Q^\mu(s, a; \phi) = \mathbb{E}_{s' [r + \gamma \mathbb{E}_{a' \sim \mu_\theta} [Q^\mu(s', a'; \phi)]]$. Parameter θ of $\mu(s; \theta)$ is subsequently optimized by the gradient of $\mathcal{J}(\theta)$ as $\nabla_\theta \mathcal{J}(\theta) = \mathbb{E}_{s \sim \mathcal{D}} [\nabla_\theta \mu(s; \theta) \nabla_a Q^\mu(s, a; \phi)|_{a=\mu(s; \theta)}]$. \mathcal{D} is an experience replay buffer that stores (s, a, r, s') samples obtained from the training episodes. Deep deterministic policy gradient (DDPG) [9], an actor-critic algorithm based on DPG, uses deep neural networks to approximate critic $Q^\mu(s, a; \phi)$ and actor $\mu(s; \theta)$ of the agent.

MADDPG is a multi-agent extension of DDPG for deriving decentralized policies in the CTDE framework. In MADDPG, each agent learns an individual policy that maps the observation to its action to maximize its expected return, which is approximated by the Q -network. MADDPG comprises individual Q -networks and policy networks for each agent. MADDPG lets the Q -network (centralized critic) of agent i be trained by minimizing the loss with the target Q -value, y_i , as follows:

$$\begin{aligned} \mathcal{L}(\phi_i) &= \mathbb{E}_{\mathbf{o}, \mathbf{a}, r, \mathbf{o}' \sim \mathcal{D}} [(Q_i^\mu(\mathbf{o}, \mathbf{a}; \phi_i) - y_i)^2], \\ y_i &= r_i + \gamma Q_i^{\mu'}(\mathbf{o}', \mathbf{a}'; \phi_i')|_{a'_j = \mu'_j(\mathbf{o}'; \theta'_j)}, \end{aligned} \quad (1)$$

where $\mathbf{o} = (o_1, \dots, o_N)$ and $\mathbf{a} = (a_1, \dots, a_N)$ represent the joint observation and joint action of all agents, respectively. \mathcal{D} is an

experience replay buffer that stores $(\mathbf{o}, \mathbf{a}, r, \mathbf{o}')$ samples obtained from the training episodes. $Q^{\mu'}$ and μ' are target networks for the stable learning of Q and policy networks. The policy network (actor), $\mu_i(o_i; \theta_i)$, of agent i is optimized using the gradient computed as

$$\nabla_{\theta_i} \mathcal{J}(\theta_i) = \mathbb{E}_{\mathbf{o}, \mathbf{a} \sim \mathcal{D}} [\nabla_{\theta_i} \mu_i(o_i; \theta_i) \nabla_{a_i} Q_i^\mu(\mathbf{o}, \mathbf{a}; \phi_i)|_{a_i = \mu_i(o_i; \theta_i)}]. \quad (2)$$

4 METHODS

F2DDPG learns the critic and actor of each agent using the biased action information of other agents based on a friend-or-foe concept [11]. In F2DDPG, as shown in Figure 1, each agent has two perceptions on the environment:

- The real environment, where agent i 's true critic can be estimated using its own action and the realized (actual) actions, $(a_i, \mathbf{a}_{-i}^A, \mathbf{a}_{-i}^E)$;
- An imaginary environment, where agent i 's imaginary critic can be estimated using its own action and the imaginary cooperative and competitive joint actions, $(a_i, \mathbf{a}_{-i}^{*A}, \mathbf{a}_{-i}^{*E})$. These are computed based on the assumption that all the ally agents execute the cooperative joint action to agent i , while all the enemy agents execute the competitive joint action to agent i .

In F2DDPG, each agent learns the decentralized actor (policy) by applying the following three iterative procedures: (1) computing the biased actions, $\bar{\mathbf{a}}_{-i}^A$ and $\bar{\mathbf{a}}_{-i}^E$, by combining the real joint action and imaginary cooperative and competitive joint actions; (2) updating the biased critic using the biased actions; and (3) updating the actor using the biased critic.

Thus, the updated policy with biased cooperative-competitive joint actions is more likely to recommend such cooperative and competitive actions to other agents, which introduces meaningful interactions among agents and enhances policy learning. Therefore, using the biased actions, F2DDPG can learn the level of cooperation and competition among agents in a sample-efficient manner. The biased actions become increasingly closer to the actually executed

actions as the training proceeds, implying that the biases in the actors vanish.

Figure 1 illustrates how agent i (predator) updates its critic in F2DDPG, while playing the 3 vs. 3 predator-prey game. In this game, the three predators (red circles) attempt to capture the three prey (green squares) together. Because the prey are faster than the predators, the predators (ally group) must cooperate strategically to capture the prey (enemy group). In the early stages of learning, however, it is difficult for the predators to achieve a reward. Whenever the reward is realized, whether by chance or strategic moves, predator i computes imaginary actions by assuming that the reward is realized when the other two predators choose the optimal cooperative joint action for predator i , while the other three prey execute the competitive joint action to predator i . Agent i updates its critic using its own action and the one-step biased joint cooperative-competitive actions of other agents.

4.1 Computing Biased Cooperative-Competitive Actions

Considering relationships to agent i , in the mixed cooperative-competitive environment, we categorize all agents, except i , as cooperative and competitive to agent i as follows:

- $A(i)$: set of agents cooperative to agent i (ally group);
- $E(i)$: set of agents competitive to agent i (enemy group);
- \mathbf{a}_{-i}^A : joint action of agents in $A(i)$;
- \mathbf{a}_{-i}^E : joint action of agents in $E(i)$.

If one assumes that the agents in $A(i)$ and $E(i)$ jointly maximize and minimize the critic of agent i , respectively, then agent i can estimate its critic as follows:

$$\bar{Q}_i^H(\mathbf{o}, a_i; \phi_i) = \max_{\mathbf{a}_{-i}^A} \min_{\mathbf{a}_{-i}^E} Q_i^H(\mathbf{o}, a_i, \mathbf{a}_{-i}^A, \mathbf{a}_{-i}^E; \phi_i). \quad (3)$$

This estimated critic is biased because each agent in $A(i)$ or $E(i)$ executes its action to only maximize its own individual critic similar to the process in MADDPG.

The two optimal joint actions (\mathbf{a}_{-i}^{*A} , \mathbf{a}_{-i}^{*E}) that achieve the *maxmin* value in Equation 4 are called a saddle-point equilibrium strategy, which is equivalent to the Nash equilibrium strategy, which satisfies

$$\begin{aligned} \mathbf{a}_{-i}^{*A} &= \operatorname{argmax}_{\mathbf{a}_{-i}^A} Q_i^H(\mathbf{o}, a_i, \mathbf{a}_{-i}^A, \mathbf{a}_{-i}^{*E}; \phi_i), \\ \mathbf{a}_{-i}^{*E} &= \operatorname{argmin}_{\mathbf{a}_{-i}^E} Q_i^H(\mathbf{o}, a_i, \mathbf{a}_{-i}^{*A}, \mathbf{a}_{-i}^E; \phi_i). \end{aligned} \quad (4)$$

We refer to these two actions in Equation 4 as the imaginary cooperative and competitive joint actions because these two joint actions rarely occur in the real environment.

The estimated critic of agent i using Equation 4 can be used in a decentralized training setting, where each agent cannot observe other agents' actions during training, and thus, has to infer them. However, the current study focuses on developing an efficient MARL algorithm in the CTDE framework that allows each agent to observe other agents' actions during training. To help stabilize training without introducing any bias, other agents' actions are explicitly used when learning the critic and associated actor.

In this study, we propose to combine these two different learning paradigms to: (1) achieve reliable and stable learning using the true

action information in the CTDE framework and (2) infuse the desirable behaviors (information biases) into agents using imaginary joint cooperative-competitive actions (biased action information) computed in the decentralized learning framework. To achieve both objectives, the proposed method computes the biased actions by combining the actual and imaginary actions as follows:

$$\begin{aligned} \bar{\mathbf{a}}_{-i}^A &= \mathbf{a}_{-i}^A + \delta^A \nabla_{\mathbf{a}_{-i}^A} Q_i^H(\mathbf{o}, a_i, \mathbf{a}_{-i}^A, \mathbf{a}_{-i}^E; \phi_i), \\ \bar{\mathbf{a}}_{-i}^E &= \mathbf{a}_{-i}^E - \delta^E \nabla_{\mathbf{a}_{-i}^E} Q_i^H(\mathbf{o}, a_i, \mathbf{a}_{-i}^A, \mathbf{a}_{-i}^E; \phi_i). \end{aligned} \quad (5)$$

In Equation 5, we compute the one-step-biased cooperative-competitive joint actions, which approximate the imaginary cooperative-competitive joint actions in Equation 4, using the partial gradient of critic Q_i^H . In addition, δ^A and δ^E are the step sizes for the biased cooperative and competitive joint actions, respectively.

Note that computing $\bar{\mathbf{a}}_{-i}^A$ and $\bar{\mathbf{a}}_{-i}^E$ is computationally tractable because it only requires computing the partial gradient of the Q -network with respect to the joint action variables. In addition, δ^A and δ^E adjust the level of biases (injected cooperative and competitive biases). The optimal δ^A and δ^E can be empirically determined during training; however, these values do not significantly affect the learning performance because the partial gradient eventually becomes small, making the amount of action modification negligible.

4.2 Learning the Biased Critic Using Biased Actions

The biased actions in Equation 5 are then used to update the Q -network as:

$$\begin{aligned} \mathcal{L}(\phi_i) &= \mathbb{E}_{\mathbf{o}, a, r, \sigma' \sim \mathcal{D}} [(Q_i^H(\mathbf{o}, a; \phi_i) - y_i)^2], \\ y_i &= r_i + \gamma Q_i^H(\sigma', a', \bar{\mathbf{a}}_{-i}^A, \bar{\mathbf{a}}_{-i}^E; \phi_i') |_{a'_i = \mu'_i(\sigma'_i, \theta'_i)}, \end{aligned} \quad (6)$$

where $\bar{\mathbf{a}}_{-i}^A$ and $\bar{\mathbf{a}}_{-i}^E$ are, respectively, the biased actions computed using Equation 5 with the target Q -network, Q_i^H . The target network is designed to stabilize the learning of the Q -network by slowly changing the parameters ϕ' of the target Q -network [12].

4.3 Learning the Actor from the Biased Critic

The biased critic Q_i^H is used to update the decentralized actor as a deterministic policy. The deterministic policy network μ_i is updated using the gradient computed as

$$\begin{aligned} \nabla_{\theta_i} \mathcal{J}(\theta_i) &= \\ \mathbb{E}_{\mathbf{o}, a \sim \mathcal{D}} [\nabla_{\theta_i} \mu_i(o_i; \theta_i) \nabla_{a_i} Q_i^H(\mathbf{o}, a_i, \bar{\mathbf{a}}_{-i}^A, \bar{\mathbf{a}}_{-i}^E; \phi_i) |_{a_i = \mu_i(o_i; \theta_i)}], \end{aligned} \quad (7)$$

where $\bar{\mathbf{a}}_{-i}^A$ and $\bar{\mathbf{a}}_{-i}^E$ are the biased joint actions of the ally and enemy agents, respectively, each of which is computed using Equation 5 with the Q -network, Q_i^H . We estimate the gradient reliably with the biased actions computed using the sampled joint actions (executed true actions) from the experience replay buffer.

Each agent then updates the parameters of its own policy network using the computed policy gradient in Equation 7. Owing to the biased action information, the trained policy is also biased such that it is inevitably subjected to recommend an action to cooperate and compete with other agents. Consequently, the biased policy introduces more active interactions among agents, and thus,

Algorithm 1 Friend-or-Foe Multi-Agent Deep Deterministic Policy Gradient Algorithm for N agents

```

1: Initialize actor networks  $\mu$ , critic networks  $Q$ , target networks  $\mu'$  and  $Q'$ , experience replay buffer  $\mathcal{D}$ 
2: for episode = 1 to  $M$  do
3:   Initialize a random process  $\mathcal{N}$  for action exploration
4:   Receive initial observation  $\mathbf{o}$ 
5:   for  $t = 1$  to  $T$  do
6:     For each agent  $i$ , select action  $a_i = \mu_i(o_i; \theta_i) + \mathcal{N}_t$ 
7:     Execute actions  $\mathbf{a}$  and receive reward  $\mathbf{r} = (r_1, \dots, r_N)$  and new observation  $\mathbf{o}'$ 
8:     Store  $(\mathbf{o}, \mathbf{a}, \mathbf{r}, \mathbf{o}')$  in experience replay buffer  $\mathcal{D}$ 
9:      $\mathbf{o} \leftarrow \mathbf{o}'$ 
10:    for agent  $i = 1$  to  $N$  do
11:      Sample a random minibatch of  $\mathcal{B}$  samples  $(\mathbf{o}^b, \mathbf{a}^b, \mathbf{r}^b, \mathbf{o}'^b)$  from  $\mathcal{D}$ 
12:      Set  $y_i^b = r_i^b + \gamma Q_i^{\mu'}(\mathbf{o}^b, a_i^b, \bar{\mathbf{a}}_{-i}^A, \bar{\mathbf{a}}_{-i}^E; \phi_i) |_{a_i' = \mu_i'(o_i^b; \theta_i')}$ ,
13:       $\bar{\mathbf{a}}_{-i}^A = \mathbf{a}_{-i}^A + \delta^A \nabla_{\mathbf{a}_{-i}^A} Q_i^{\mu'}(\mathbf{o}^b, a_i^b, \mathbf{a}_{-i}^A, \mathbf{a}_{-i}^E; \phi_i) |_{a_k' = \mu_k'(o_k^b; \theta_k')}$ ,
14:       $\bar{\mathbf{a}}_{-i}^E = \mathbf{a}_{-i}^E - \delta^E \nabla_{\mathbf{a}_{-i}^E} Q_i^{\mu'}(\mathbf{o}^b, a_i^b, \mathbf{a}_{-i}^A, \mathbf{a}_{-i}^E; \phi_i) |_{a_k' = \mu_k'(o_k^b; \theta_k')}$ 
15:      Update critic by minimizing the loss:  $\mathcal{L}(\phi_i) = \frac{1}{\mathcal{B}} \sum_b (Q_i^{\mu}(\mathbf{o}^b, \mathbf{a}^b; \phi_i) - y_i^b)^2$ 
16:      Update actor using the sampled policy gradient:  $\nabla_{\theta_i} \mathcal{J}(\theta_i) \approx \frac{1}{\mathcal{B}} \sum_b \nabla_{\theta_i} \mu_i(o_i^b; \theta_i) \nabla_{a_i} Q_i^{\mu}(\mathbf{o}^b, a_i, \bar{\mathbf{a}}_{-i}^A, \bar{\mathbf{a}}_{-i}^E; \phi_i) |_{a_i = \mu_i(o_i^b; \theta_i)}$ ,
17:       $\bar{\mathbf{a}}_{-i}^A = \mathbf{a}_{-i}^{bA} + \delta^A \nabla_{\mathbf{a}_{-i}^{bA}} Q_i^{\mu}(\mathbf{o}^b, a_i^b, \mathbf{a}_{-i}^{bA}, \mathbf{a}_{-i}^{bE}; \phi_i)$ ,
18:       $\bar{\mathbf{a}}_{-i}^E = \mathbf{a}_{-i}^{bE} - \delta^E \nabla_{\mathbf{a}_{-i}^{bE}} Q_i^{\mu}(\mathbf{o}^b, a_i^b, \mathbf{a}_{-i}^{bA}, \mathbf{a}_{-i}^{bE}; \phi_i)$ 
19:    end for
20:    Update target network parameters for each agent  $i$ :  $\phi_i' \leftarrow \tau \phi_i + (1 - \tau) \phi_i'$  and  $\theta_i' \leftarrow \tau \theta_i + (1 - \tau) \theta_i'$ 
21:  end for
22: end for

```

enhances the MARL policy learning. In addition, if the critic's gradient (bias) in the biased action and actual action from the learned policy gradually become similar, the introduced bias then naturally decreases, and the training eventually becomes unbiased, as the biased and actual actions become the same. For completeness, we provide the F2DDPG algorithm in Algorithm 1.

In some cases of training F2DDPG, the critic's gradient (bias) dominates the actual action in the biased action in learning the Q -network and policy network. Thus, the magnitude of the gradient is made equal to the magnitude of the actual action to prevent the biased action from being extremely biased and destabilizing the training, as follows:

$$\begin{aligned}
g &= \nabla_{a_i} Q_i^{\mu}(\mathbf{o}, \mathbf{a}; \phi_i), \\
\bar{a}_{-i} &= a_{-i} \pm \delta \|a_{-i}\|_2 \frac{g}{\|g\|_2}.
\end{aligned} \tag{8}$$

The trick of Equation 8 is utilized in Equation 5.

5 EXPERIMENTS

Figure 2 shows the environments (games) used to evaluate the performances of the proposed and baseline algorithms. The environments are those used in previous studies [8, 12] and those designed to make MARL training more difficult. We assume that the agents in the environments observe the relative positions and velocities of all agents. In Figure 2, games (a) and (b) correspond to cooperative environments with only agents in a cooperative relationship, and games (c) and (d) correspond to mixed cooperative-competitive environments with both cooperative and competitive agents. In

games (b) and (d), the agents need to communicate with other agents depending on the purpose of the games.

For the experiments in these games, we compare the performance of F2DDPG to the following baseline algorithms:

- **MADDPG** [12] is the algorithm that learns the critics and actors using only the actual action information in the CTDE framework.
- **M3DDPG** [8] is the algorithm based on MADDPG. Instead of using the actual action information, this algorithm uses the noisy actions of other agents to update the critics and actors. In particular, this algorithm computes the adversarial noise for other agents' actions such that the noisy actions collectively minimize the target agent's critic. Owing to the use of adversarial noise, this algorithm is referred to as robust MARL. In the view of F2DDPG, this algorithm can be interpreted as one in which the competitive roles are infused, as an information bias, to all the agents, regardless of their roles (cooperative or competitive) with respect to the target agent.
- **All Plus** is the algorithm with only cooperative biases employed in the actions of other agents; other agents are assumed to maximize the target agent's critic, regardless of their roles.
- **Random Sign** is the algorithm with random biases employed in the actions of other agents. Other randomly selected agents are assumed to maximize the target agent's critic, while the remaining agents are assumed to minimize the target agent's critic, regardless of their roles.

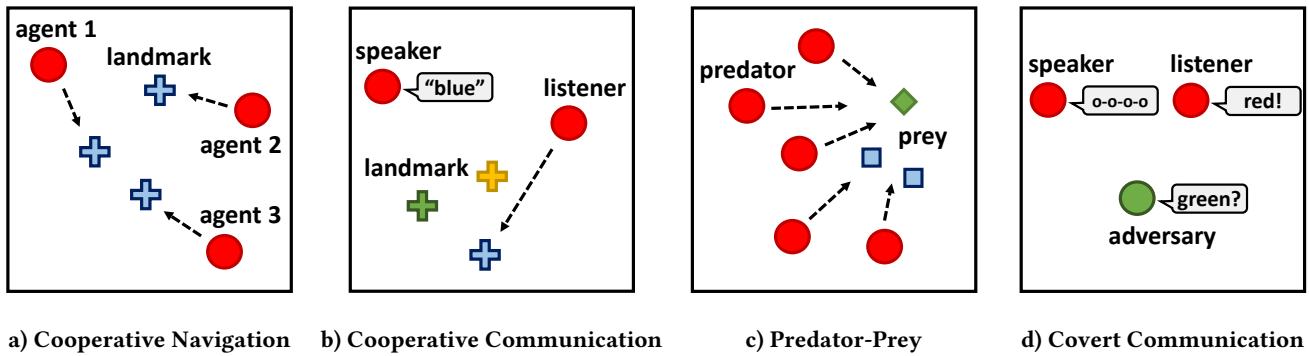


Figure 2: Illustrations of the experimental environments.

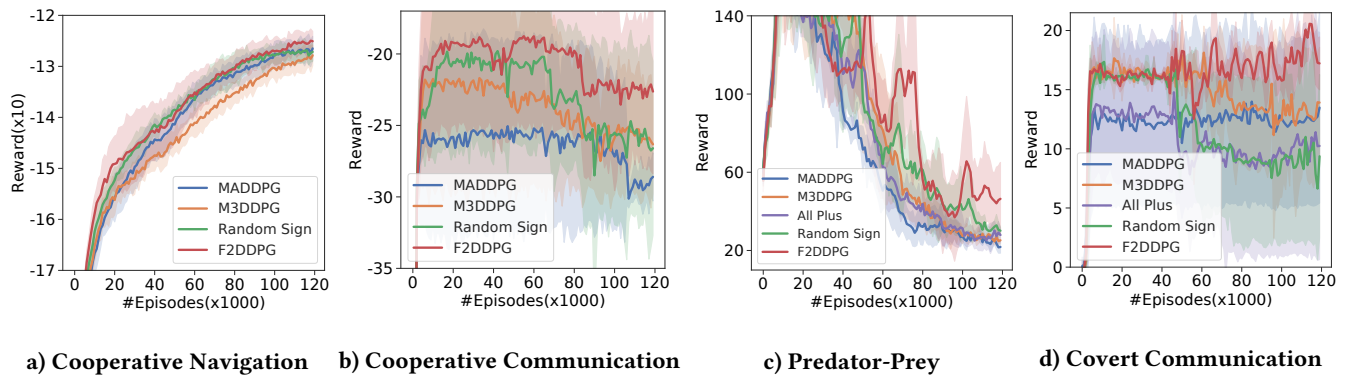


Figure 3: Rewards of the red agents in the experimental environments.

We can differentiate the proposed F2DDPG and other baseline algorithms depending on the type of information bias infused into other agents' actions. While the baseline algorithms either do not use any information bias (MADDPG) or use a certain information bias, regardless of the relationships among agents, F2DDPG is the only algorithm that aligns the information bias with the actual roles of agents. Note that we exclude the All Plus algorithm from the baseline algorithms in cooperative environments, such as games (a) and (b), because it has the same cooperative biases as F2DDPG for all the cooperative agents. In this study, all the performances are obtained by the trained policies with four different random seeds.

5.1 Cooperative Navigation

Cooperative navigation is a cooperative environment, as shown in Figure 2 (a), in which three cooperative agents (red circles) must reach three landmarks (blue crosses) without colliding with each other, while covering all of the landmarks. Every episode starts with randomly initialized positions for the agents and landmarks. The agents are collectively rewarded based on the distance of the nearest agent to each landmark and penalized for collisions with other agents during navigation. Thus, each agent must cooperate to occupy a distinct landmark without colliding with other agents.

As shown in Figure 3 (a), F2DDPG outperforms other baseline algorithms with faster training speed and higher converged rewards. We consider that the performance improvement of F2DDPG is a

result of the optimal use of information bias corresponding to the agents' roles. F2DDPG constantly updates the critic and actor of each agent using biased joint actions, which induces agents' policies to recommend more coherently exploratory actions, especially toward cooperation. We believe that such coherent exploration helps the MARL policy learning more than random exploration.

In contrast, M3DDPG exhibits slow training in inducing cooperation among the three agents by learning only the adversarial action information of other agents, although the agents are in a cooperative relationship. Training with adversarial action information may enable robust policy learning; however, it is believed to be unhelpful in inducing cooperation among the agents. The Random Sign algorithm randomly injects cooperative or competitive biases for biased actions at every step of training, which leads to random noise in training. This random noise allows the algorithm to train faster than MADDPG and M3DDPG; however, the algorithm exhibits slower training than F2DDPG.

5.2 Cooperative Communication

Cooperative communication is a cooperative environment, as shown in Figure 2 (b), with two cooperative agents, a speaker and listener (red circles), and three landmarks of differing colors. The listener must navigate to a landmark of a particular color. However, the listener does not know the landmark to which it must navigate, while observing the relative position and color of the landmarks. In

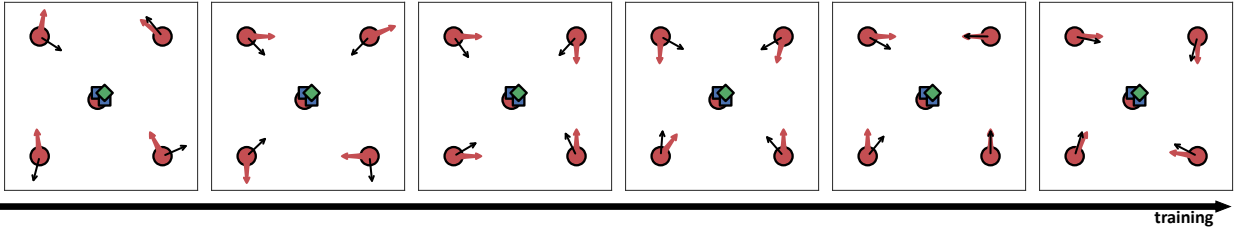


Figure 5: Actual actions (red arrows) and biases (black arrows) for the centered predator as training proceeds.

contrast, the speaker observes the correct color of the landmark to which the listener must navigate, and broadcasts a message (communication vector) at each time step, which is observed by the listener. For each episode, the positions of the listener and landmarks are randomly initialized. The listener and speaker are rewarded based on listener’s distance to the correct landmark. Thus, the speaker must learn to generate a message that optimally guides the listener to reach the correct landmark, and simultaneously, the listener must learn to decipher the message that is transmitted from the speaker and navigate to the correct landmark.

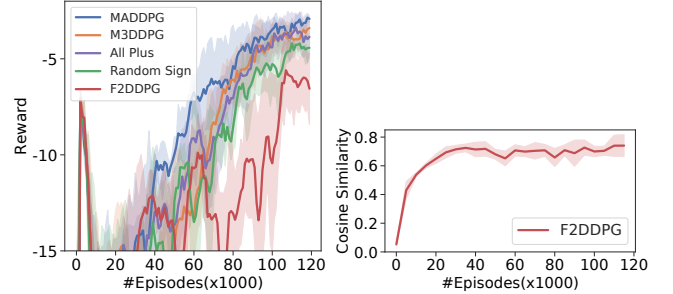
As shown in Figure 3 (b), F2DDPG outperforms other baseline algorithms with faster training speed and higher rewards. Additionally, M3DDPG exhibits better performance than MADDPG. Moreover, the Random Sign algorithm learns faster than MADDPG, possibly because of the enhanced exploration with random noises added to the observed actions.

5.3 Predator-Prey

Predator-prey is a mixed cooperative-competitive environment, as shown in Figure 2 (c), in which the five predator agents (red circles) seek to capture the three prey agents (blue squares and green diamond), which is called 5 vs. 3 predator-prey. If there are m predators and n prey, it is denoted as m vs. n predator-prey. Because prey can move at a higher speed and have greater acceleration than predators, predators must cooperate to capture the prey. In particular, the green prey (green diamond) can move faster with greater acceleration than the blue prey (blue squares); the green and blue prey are factors of 3 and 1.3 faster than the predators, respectively. The positions of the predators and prey are randomly initialized for every episode.

Each time the predators collide with (*i.e.*, capture) the prey, the predators are collectively rewarded, while the prey are penalized. When the predators capture green prey, they are rewarded with a factor of 10 more than when they capture blue prey. The predators can capture prey multiple times during an episode. In the predator-prey game, the prey are trained with MADDPG, while the predators are trained with F2DDPG and other baseline algorithms.

As shown in Figure 3 (c), F2DDPG outperforms other baseline algorithms with higher rewards. The reason why F2DDPG achieves higher rewards is hypothesized by investigating the reward lost by the green prey shown in Figure 4 (a). This figure compares the reward achieved by the green prey (trained with MADDPG) when playing against the predators trained with F2DDPG and baseline algorithms; when the reward of the prey is lower, it is more likely to be captured by the predators. As shown in the figure, the reward of



a) Rewards of green prey b) Cosine similarity between actual actions and biases

Figure 4: Results in 5 vs. 3 predator-prey.

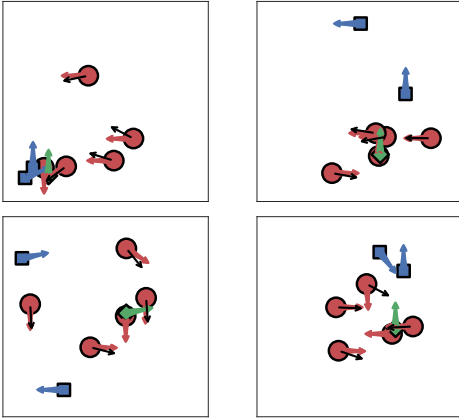
the green prey captured by the predators trained with F2DDPG is lower than that of the prey captured by the predators trained with other baseline algorithms. This indicates that the high rewards of the predators trained with F2DDPG (shown in Figure 3 (c)) are the results of capturing the green prey more frequently. Thus, it can be concluded that the predators trained with F2DDPG are more capable of cooperating strategically to capture the most rewardable, but fast, prey than the predators trained with other baseline algorithms.

We also verify that the infused bias induces desirable behaviors from the agents. Figure 4 (b) shows how the cosine similarity between the actual action \mathbf{a}_{-i}^A and bias $\nabla_{\mathbf{a}_{-i}^A} Q_i^H$ in the biased action in Equation 5 varies as the F2DDPG training proceeds. The cosine similarity between \mathbf{a}_{-i}^A and $\nabla_{\mathbf{a}_{-i}^A} Q_i^H$ is calculated as $\frac{\langle \mathbf{a}_{-i}^A, \nabla_{\mathbf{a}_{-i}^A} Q_i^H \rangle}{\|\mathbf{a}_{-i}^A\|_2 \|\nabla_{\mathbf{a}_{-i}^A} Q_i^H\|_2}$ for all agents in $A(i)$, where $\langle \cdot, \cdot \rangle$ represents the inner product operator. The similarity is used to judge the similarity of directions of the actual and biased actions. The similarity ranges from -1, indicating exactly opposite, to 1, indicating exactly the same, and 0 indicating orthogonality or decorrelation. As shown in Figure 4 (b), the similarity increases from 0 to close to 1 as the training proceeds, indicating that the actual and biased actions (desirable actions designed by the biases) become more similar as the training proceeds. Therefore, the agents eventually execute the actions designed by the biases, and accordingly, the biases between the actual and biased actions vanish.

To consider a particular fixed state where a centered predator captures the prey, we investigate how the centered predator’s trained policy induces other agents’ (predators’) actions. In Figure 5, the

Table 1: Fraction between the number of successful episodes and the total number of testing episodes in predator-prey games.

	3 vs. 1		5 vs. 3		7 vs. 3	
	$N_c \geq 1$	$N_c \geq 3$	$N_c \geq 1$	$N_c \geq 3$	$N_c \geq 1$	$N_c \geq 3$
MADDPG	2.75±0.83	0.25±0.43	14.50±1.50	1.25±1.09	56.25±6.46	17.00±6.09
M3DDPG	3.75±0.43	0.25±0.43	14.25±0.43	1.25±0.83	55.75±6.17	13.25±3.63
All Plus	3.75±1.09	0.25±0.43	15.75±2.28	1.00±1.00	62.75±7.39	14.24±4.14
Random Sign	1.50±0.50	0.00±0.00	18.75±6.98	1.75±0.83	62.50±11.71	19.25±6.37
F2DDPG	3.25±0.43	0.25±0.43	32.50±11.39	4.50±2.29	71.75±7.52	28.75±6.17

**Figure 6: Actual actions (red arrows) and biases (black arrows) for the predator capturing the green prey after the policies are learned with F2DDPG.**

figures on the left show the other agents’ biased joint action (black) induced at an early training stage, and those on the right show how the other agents’ actual joint action (red) concurrently and coherently alter as the training proceeds. Noticeable observations are that (1) the other agents are jointly heading toward the prey, which is a strategic movement to capture the prey, and (2) the actual (red) and biased (black) actions become extremely similar, which demonstrates that the biases vanish as the training proceeds, and the agents actually behave as intended by the biases.

In addition, we execute the policies trained with F2DDPG in the predator-prey game with a random state. Figure 6 shows four snapshots of the predator-prey game with a random state after the policies are trained with F2DDPG. In the figure, the predators tend to gather toward the predator capturing the green prey and attempt to capture it together. The differences between the actual (red) and biased (black) actions are negligible, meaning that the converged policies no longer carry the biases.

Table 1 compares the fraction between the number of successful episodes and the total number of testing episodes (100 episodes). We define two types of success: $N_c \geq 1$ is the case in which the predators capture the green prey at least once and $N_c \geq 3$ is the case in which the predators capture the green prey at least three times. To validate the scalability of the proposed algorithm, we compare these performance measures for F2DDPG and other baseline algorithms for different sizes of predator-prey games. In 3 vs.

1, the performances are not significantly differentiated because the number of predators is insufficient to capture the faster green prey, even if the three predators cooperate. However, as the number of predators increases, the predators capture the green prey through cooperation, and F2DDPG outperforms other baseline algorithms in 5 vs. 3, as shown in the table. In addition, F2DDPG outperforms other baseline algorithms in 7 vs. 3. Thus, it can be concluded that the proposed method improves the MARL training, even in cases with many agents, by utilizing biased action information.

5.4 Covert Communication

Covert communication is a mixed cooperative-competitive environment, as shown in Figure 2 (d), with two cooperative agents, a speaker and listener (red circles), and an adversary (green circle). The speaker must encode a message as a communication vector using a randomly generated key to output the communication vector. The listener must reconstruct the communication vector into the message using the key. However, the adversary also observes the communication vector and attempts to reconstruct the communication vector without the key. The speaker and listener are rewarded based on the listener’s reconstruction and penalized based on the adversary’s reconstruction. The adversary is rewarded based on its reconstruction. Therefore, the speaker must encrypt the message as the communication vector such that the adversary cannot decrypt the communication vector, and the listener must decrypt the communication vector as the message. In the covert communication, the adversary is trained with MADDPG for comparison, while the speaker and listener are trained with F2DDPG and other baseline algorithms.

As shown in Figure 3 (d), F2DDPG outperforms other baseline algorithms with higher rewards. At the early stage of training, M3DDPG, Random Sign, and F2DDPG outperform MADDPG and All Plus. However, as the training proceeds and the adversary becomes intelligent, the rewards of M3DDPG and Random Sign decrease, while F2DDPG still maintains high rewards.

6 CONCLUSIONS

We proposed F2DDPG, an algorithm that boosts MARL training using biased action information of other agents based on a friend-or-foe concept. Empirically, we demonstrated that F2DDPG outperforms existing algorithms in several mixed cooperative-competitive environments. We also demonstrated that F2DDPG learns the agents’ policies such that their actions become similar to the biased actions and that the biases decrease as the learning proceeds.

REFERENCES

- [1] Marcin Andrychowicz, Filip Wolski, Alex Ray, Jonas Schneider, Rachel Fong, Peter Welinder, Bob McGrew, Josh Tobin, OpenAI Pieter Abbeel, and Wojciech Zaremba. 2017. Hindsight Experience Replay. In *Advances in Neural Information Processing Systems*. 5048–5058.
- [2] Wendelin Böhrer, Tabish Rashid, and Shimon Whiteson. 2019. Exploration with unreliable intrinsic reward in multi-agent reinforcement learning. *arXiv preprint arXiv:1906.02138* (2019).
- [3] G Hailu and G Sommer. 1999. On amount and quality of bias in reinforcement learning. In *IEEE SMC'99 Conference Proceedings. 1999 IEEE International Conference on Systems, Man, and Cybernetics*, Vol. 2. IEEE, 728–733.
- [4] Junling Hu, Michael P Wellman, et al. 1998. Multiagent reinforcement learning: theoretical framework and an algorithm. In *Proceedings of the Fifteenth International Conference on Machine Learning*, Vol. 98. Citeseer, 242–250.
- [5] Shariq Iqbal and Fei Sha. 2019. Actor-attention-critic for multi-agent reinforcement learning. In *International Conference on Machine Learning*, Vol. 97. PMLR, 2961–2970.
- [6] Shariq Iqbal and Fei Sha. 2019. Coordinated Exploration via Intrinsic Rewards for Multi-Agent Reinforcement Learning. *arXiv preprint arXiv:1905.12127* (2019).
- [7] Martin Lauer and Martin Riedmiller. 2000. An algorithm for distributed reinforcement learning in cooperative multi-agent systems. In *Proceedings of the Seventeenth International Conference on Machine Learning*. Citeseer.
- [8] Shihui Li, Yi Wu, Xinyue Cui, Honghua Dong, Fei Fang, and Stuart Russell. 2019. Robust multi-agent reinforcement learning via minimax deep deterministic policy gradient. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 33. 4213–4220.
- [9] Timothy P Lillicrap, Jonathan J Hunt, Alexander Pritzel, Nicolas Heess, Tom Erez, Yuval Tassa, David Silver, and Daan Wierstra. 2016. Continuous control with deep reinforcement learning. *International Conference on Learning Representations* (2016).
- [10] Michael L Littman. 1994. Markov games as a framework for multi-agent reinforcement learning. In *Machine learning proceedings 1994*. Elsevier, 157–163.
- [11] Michael L. Littman. 2001. Friend-or-Foe Q-Learning in General-Sum Games. In *Proceedings of the Eighteenth International Conference on Machine Learning*, Vol. 1. Morgan Kaufmann Publishers Inc., 322–328.
- [12] Ryan Lowe, Yi I Wu, Aviv Tamar, Jean Harb, OpenAI Pieter Abbeel, and Igor Mordatch. 2017. Multi-agent actor-critic for mixed cooperative-competitive environments. In *Advances in Neural Information Processing Systems*. 6379–6390.
- [13] Paulo Rauber, Avinash Ummadisingu, Filipe Mutz, and Juergen Schmidhuber. 2019. Hindsight policy gradients. *International Conference on Learning Representations* (2019).
- [14] Julien Roy, Paul Barde, Félix G Harvey, Derek Nowrouzezahrai, and Christopher Pal. 2019. Promoting Coordination through Policy Regularization in Multi-Agent Deep Reinforcement Learning. *arXiv preprint arXiv:1908.02269* (2019).
- [15] Heechang Ryu, Hayong Shin, and Jinkyoo Park. 2020. Multi-Agent Actor-Critic with Hierarchical Graph Attention Network. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 34. 7236–7243.
- [16] David Silver, Guy Lever, Nicolas Heess, Thomas Degris, Daan Wierstra, and Martin Riedmiller. 2014. Deterministic Policy Gradient Algorithms. In *International Conference on Machine Learning*, Vol. 32. PMLR, 387–395.
- [17] Richard S Sutton and Andrew G Barto. 2018. *Reinforcement learning: An introduction*. MIT press.
- [18] Tonghan Wang, Jianhao Wang, Yi Wu, and Chongjie Zhang. 2020. Influence-Based Multi-Agent Exploration. *International Conference on Learning Representations* (2020).