

Backpropagation Through Agents

Zhiyuan Li¹, Wenshuai Zhao², Lijun Wu¹, Joni Pajarinen²

¹School of Computer Science and Engineering, University of Electronic Science and Technology of China

²Department of Electrical Engineering and Automation, Aalto University

zhiyuanli@std.uestc.edu.cn, {wenshuai.zhao,joni.pajarinen}@aalto.fi, wljuestc@sina.com

Abstract

A fundamental challenge in multi-agent reinforcement learning (MARL) is to learn the joint policy in an extremely large search space, which grows exponentially with the number of agents. Moreover, fully decentralized policy factorization significantly restricts the search space, which may lead to sub-optimal policies. In contrast, the auto-regressive joint policy can represent a much richer class of joint policies by factorizing the joint policy into the product of a series of conditional individual policies. While such factorization introduces the action dependency among agents explicitly in sequential execution, it does not take full advantage of the dependency during learning. In particular, the subsequent agents do not give the preceding agents feedback about their decisions. In this paper, we propose a new framework Back-Propagation Through Agents (BPTA) that directly accounts for both agents' own policy updates and the learning of their dependent counterparts. This is achieved by propagating the feedback through action chains. With the proposed framework, our Bidirectional Proximal Policy Optimisation (BPPO) outperforms the state-of-the-art methods. Extensive experiments on matrix games, StarCraftII v2, Multi-agent MuJoCo, and Google Research Football demonstrate the effectiveness of the proposed method.

Introduction

Multi-agent reinforcement learning (MARL) is a promising approach to many real-world applications that naturally comprise multiple decision-makers interacting at the same time, such as cooperative robotics (Yu et al. 2023), traffic management (Ma and Wu 2020), and autonomous driving (Shalev-Shwartz, Shammah, and Shashua 2016). Although reinforcement learning (RL) has recorded sublime success in various single-agent domains, trivially applying single-agent RL algorithms in this setting brings about the curse of dimensionality. In multi-agent settings, agents need to explore an extremely large policy space, which grows exponentially with the team size, to learn the optimal joint policy.

Existing popular multi-agent policy gradient (MAPG) frameworks (Lowe et al. 2017; Foerster et al. 2018; Yu et al. 2022; Wang et al. 2023; Zhang et al. 2021) often directly represent the joint policy as the Cartesian Product of each

agent's fully independent policy. However, this factorization ignores the coordination between agents and severely limits the complexity of the joint policy, causing the learning algorithm to converge to a *Pareto-dominant equilibrium* (Christianos, Papoudakis, and Albrecht 2022). This phenomenon is commonly referred to as *relative overgeneralization* (Wei et al. 2018; Wang et al. 2021b) and can occur even in simple scenarios (Fu et al. 2022; Ye et al. 2023). To tackle these issues, some recent works (Wang, Ye, and Lu 2023; Wen et al. 2022; Fu et al. 2022) present the joint policy in an *auto-regressive* form based on the chain rule (Box et al. 2015). The auto-regressive model specifies that an agent's policy depends on its preceding agents' actions. In this way, the dependency among agents' policies is explicitly considered and the expressive limitations of the joint policy can be significantly relaxed. However, they only take into account the preceding agents' actions during decision-making, i.e., the forward process, while disregarding reactions from subsequent agents during policy improvement, i.e., the backward process (Li et al. 2023). This may lead to conflicting directions in policy updates for individual agents, where their local improvements may jointly result in worse outcomes. In contrast, the neural circuits in the central nervous system responsible for the sensorimotor loop consist of two internal models (Müller, Ohström, and Lindenberger 2021): 1) the forward model, which builds the causal flow by integrating the joint actions, and 2) the backward model, which maps the relation between an action and its consequence to invert the causal flow. Such two bidirectional models internally interact in order to enhance learning mechanisms.

In this paper, we aim to augment the existing MAPG framework with bidirectional dependency (Li et al. 2023), i.e. forward and backward process, to provide richer peer feedback and align the policy improvement directions of individual agents with that of the joint policy. To this end, we propose *Back-Propagation Through Agents* (BPTA), a multi-agent reinforcement learning framework that follows the Back-Propagation Through Time (BPTT) used for training recursive neural networks (RNN) (Cho et al. 2014). Specifically, BPTA begins by unfolding the execution sequence in agents. The actions passed to subsequent agents during the forward process will be integrated with their own actions and serve as latent variables (Kingma and Welling 2022) in the backward process, the reactions from the subse-

quent agents are propagated to the preceding agents through the variables using the reparameterization trick (Kingma, Salimans, and Welling 2015). Taking the feedback from subsequent agents into account allows each agent to learn from the consequences of the collective actions and adapt to the changing behavior of the team. Furthermore, dependent on such rich feedback agents can complete the causality loop: cyclic interaction between the forward and backward process. As a result, BPTA enables individuals to function as a whole and find a consistent improvement direction. We incorporate PPO with auto-regressive policy and BPTA and propose Bidirectional Proximal Policy Optimisation (BPPO). Empirically, in several tasks, including matrix game (Claus and Boutilier 1998a), Google Research Football (GRF) (Kurach et al. 2020), StarCraft Multi-Agent Challenge Version 2 (SMACv2) (Ellis et al. 2022), and Multi-agent MuJoCo (MA-MuJoCo), BPPO achieves better performance than baselines.

Specifically, our contribution is summarized as follows.

- We propose a novel framework BPTA that, for the first time, explicitly models feedback from action-dependent peer agents. In particular, BPTA allows derivatives to pass across agents during learning.
- Our proposed framework can be naturally integrated with existing conditional policy-gradient methods. We augment PPO with the auto-regressive policy under the BPTA framework and propose Bidirectional Proximal Policy Optimisation (BPPO).
- Finally, the effectiveness of the proposed method is verified in four cooperative environments, and empirical results show that the proposed method outperforms the state-of-the-art algorithms.

Related Work

Various works on MARL have been proposed to tackle cooperative tasks, including algorithms in which agents make decisions simultaneously and algorithms that coordinate agents' actions based on static or dynamic execution orders.

Simultaneous decision scheme. Most algorithms tend to follow a simultaneous decision scheme, where agents' policies are only conditioned on their individual observations. One line of research extends PG from RL to MARL (Lowe et al. 2017; Foerster et al. 2018; Wang et al. 2021b; Yu et al. 2022; Wang et al. 2023; Zhang et al. 2021), adopting the Actor-Critic (AC) approach, where each actor explicitly represents the independent policy, and the estimated centralized value function is known as the critic. Under this scheme, in contrast to independent updates, some recent methods sequentially execute agent-by-agent updates, such as Rollout and Policy Iteration for a Single Agent (RPISA) (Bertsekas 2021), Heterogeneous PPO (HAPPO) (Kuba et al. 2022), and Agent-by-agent Policy Optimization (A2PO) (Wang et al. 2023). Another line is value-based methods, where the joint Q-function is decomposed into individual utility functions following different interpretations of Individual-Global-Maximum (IGM) (Sunehag et al. 2018; Rashid et al. 2020; Son et al. 2019; Wang et al. 2021a; Wan et al. 2022). VDN (Sunehag et al. 2018) and QMIX (Rashid et al. 2020)

provide sufficient conditions for IGM, however, suffer from structural constraints. QTRAN (Son et al. 2019) and QPLEX (Wang et al. 2021a) complete the representation capacity of the joint Q-function through optimization constraints and a dueling mixing network respectively, while it is impractical in complicated tasks. Wan et al. introduce Optimal consistency and True-Global-Max (TGM), then propose GVR to ensure the optimality. A special case is SeCA (Zang et al. 2023), which factorizes the joint policy evaluation into a sequence of successive evaluations.

Sequential decision scheme. In this scheme, algorithms explicitly model the coordination among agents via actions. One perspective is the auto-regressive paradigm, where agents make decisions sequentially (Wen et al. 2022; Fu et al. 2022; Ye et al. 2023; Wang, Ye, and Lu 2023; Li et al. 2023). MAT (Wen et al. 2022) transform MARL into a sequence modeling problem, and introduce *Transformer* (Vaswani et al. 2017) to generate solutions. However, MAT may fail to achieve the monotonic improvement guarantee as it does not follow the sequential update scheme. Wang, Ye, and Lu derives the multi-agent conditional factorized soft policy iteration theorem by incorporating auto-regressive policy into SAC (Haarnoja et al. 2019). ACE (Li et al. 2023) and TAD (Ye et al. 2023) first cast the Multi-agent Markov decision process (MMDP) (Littman 1994) as an equivalent single-agent Markov decision process (MDP), and solve the single-agent MDP with Q-learning and PPO, respectively. However, only ACE considers the reactions from subsequent agents by calculating the maximum Q-value over the possible actions of the successors. In another perspective, the interactions between agents are modeled by a coordination graph (Ruan et al. 2022; Yang et al. 2022). However, the updates of the agents in the graph are independent of the subsequent agents.

In contrast, our proposed BPTA augmented auto-regressive method lies in the second category and is the first bidirectional PG-based MARL method.

Background

Problem Formulation

In MARL, a *decentralized partially observable Markov decision process* (Dec-POMDP) (Oliehoek and Amato 2016) is commonly applied to model the interaction among agents within a shared environment under partial observability. A Dec-POMDP is defined by a tuple $G = \langle \mathcal{N}, \mathcal{S}, \mathcal{A}, \mathcal{P}, \Omega, O, \mathcal{R}, \rho_0, \gamma \rangle$, where $\mathcal{N} = \{1, \dots, n\}$ is a set of agents, $s \in \mathcal{S}$ denotes the state of the environment, $\mathcal{A} = \prod_{i=1}^n A^i$ is the product of the agents' action spaces, namely the joint action space, $\Omega = \prod_{i=1}^n \Omega^i$ is the set of joint observations, and ρ_0 is the distribution of the initial state. At time step $t \in \mathbb{N}$, each agent $i \in \mathcal{N}$ takes an action a_t^i according to its policy $\pi^i(\cdot | o_t^i)$, where o_t^i is drawn from the observation function $O(s_t, i)$. With the joint observation $\mathbf{o}_t = \{o_t^1, \dots, o_t^n\}$ and the joint action of agents $\mathbf{a}_t = \{a_t^1, \dots, a_t^n\}$ drawn from the joint policy $\pi(\cdot | \mathbf{o}_t)$, the environment moves to a state s' with probability $\mathcal{P}(s' | s, \mathbf{a}_t)$, and each agent receives a joint reward $r_t = \mathcal{R}(s_t, a_t) \in \mathbb{R}$. The state value function, the

state-action value function, and the advantage function are defined as $V_{\pi}(s) \triangleq \mathbb{E}_{\mathbf{a}_{0:\infty \sim \pi}, s_{1:\infty \sim \pi}} [\sum_{t=0}^{\infty} \gamma^t r_t | s_0 = s]$, $Q_{\pi}(s, \mathbf{a}) \triangleq \mathbb{E}_{\mathbf{a}_{1:\infty \sim \pi}, s_{1:\infty \sim \pi}} [\sum_{t=0}^{\infty} \gamma^t r_t | s_0 = s, \mathbf{a}_0 = \mathbf{a}]$, and $A_{\pi}(s, \mathbf{a}) \triangleq Q_{\pi}(s, \mathbf{a}) - V_{\pi}(s)$. The agents aim to maximize the expected total reward:

$$\mathcal{J}(\pi) \triangleq \mathbb{E}_{s_0, \mathbf{a}_0, \dots} \left[\sum_{t=0}^{\infty} \gamma^t r_t \right], \quad (1)$$

where $s_0 \sim \rho_0(s_0)$, $\mathbf{a}_t \sim \pi(\mathbf{a}_t | s_t)$. In order to keep the notation concise, we will use state s in the subsequent equations.

Independent Multi-Agent Stochastic Policy Gradient

In cooperative MARL tasks, popular PG methods follow fully independent factorization: $\pi(\mathbf{a}|s) = \prod_{i=1}^n \pi_{\theta_i}(a^i|s)$. With such a form, following along the standard Stochastic Policy Gradient Theorem, Wei et al. derive the independent multi-agent policy gradient estimator for the cooperative MARL:

$$\begin{aligned} \nabla_{\theta_i} \mathcal{J}(\theta) = & \int_{\mathcal{S}} \rho^{\pi}(s) \int_{\mathcal{A}^i} \nabla_{\theta_i} \pi_{\theta_i}(a^i|s) \\ & \int_{\mathcal{A}^{-i}} \prod_{j \neq i}^n \pi_{\theta_j}(a^j|s) Q_{\pi}(s, \mathbf{a}) d\mathbf{a}^{-i} da^i ds, \end{aligned} \quad (2)$$

where the notation $-i$ indicates all other agents except agent i , $\mathcal{P}(s \rightarrow s', t, \pi)$ denotes the density at state s' after transitioning for t time steps from state s , and $\rho^{\pi}(s) = \int_{\mathcal{S}} \sum_{t=1}^{\infty} \gamma^{t-1} \rho_0(s) \mathcal{P}(s \rightarrow s', t, \pi)$ is the (unnormalized) discounted distribution over states induced by the joint policy π .

Method

This section considers an auto-regressive joint policy with fixed execution order $\{1, 2, \dots, n\}$:

$$\pi(\mathbf{a}|s) = \prod_{i=1}^n \pi_{\theta_i}(a^i|s, a^1, \dots, a^{i-1}) \quad (3)$$

Although such factorization introduces forward dependency among agents, it ignores the reaction of subsequent policy updates on the preceding actions. To achieve bidirectional dependency, we propose Back-Propagation Through Agents (BPTA) to pass gradients across agents. Specifically, we leverage the reparameterization trick and devise a new *multi-agent conditional policy gradient theorem* that exploits the action dependency among agents. To cover any action-dependent policy, the relationship between the joint policy and individual policies can be stated as:

$$\pi(\mathbf{a}|s) = \prod_{i=1}^n \pi_{\theta_i}(a^i|s, a^{\mathcal{F}_i}), \quad (4)$$

where \mathcal{F}_i denotes the set of agents on which agent i has a forward dependency, and $a^{\mathcal{F}_i}$ are the actions taken by those agents.

Back-Propagation Through Agents

In the social context, joint action usually requires people to actively modify their own actions to reach a common action goal. Accordingly, joint action demands not only the integration of one's own and others' actions but also the corresponding consequences. However, most previous approaches assume that each agent only needs to account for its own learning process and completely disregarded the evaluation of its dependent actions result. In this section, we will show that our conditional gradient explicitly accounts for the effect of an agent's actions on the policies of its backward-dependent peer agents by additionally including agent feedback passed through the action dependency.

Theorem 1 (Conditional Multi-Agent Stochastic Policy Gradient Theorem). *For any episodic cooperative stochastic game with n agents, the gradient of the expected total reward for agent i , who has a backward dependency on some other peer agents \mathcal{B}_i using parameters $\theta_{\mathcal{B}_i}$, with respect to current policy parameters θ_i is:*

$$\begin{aligned} \nabla_{\theta_i} \mathcal{J}(\theta) = & \int_{\mathcal{S}} \rho^{\pi}(s) \left[\int_{\mathcal{A}^i} \underbrace{\nabla_{\theta_i} \pi_{\theta_i}(a^i|s, a^{\mathcal{F}_i})}_{\text{Own Learning}} \right. \\ & \int_{\mathcal{A}^{-i}} \pi_{\theta_{-i}}(a^{-i}|s', a^{\mathcal{F}_{-i}}) + \int_{\mathcal{A}^i} \pi_{\theta_i}(a^i|s, a^{\mathcal{F}_i}) \\ & \left. \int_{\mathcal{A}^{\mathcal{F}_i}} \pi_{\theta_{\mathcal{F}_i}}(a^{\mathcal{F}_i}|s, a^{\mathcal{F}_{\mathcal{F}_i}}) \right] \\ & \underbrace{\int_{\mathcal{A}^{\mathcal{B}_i}} \nabla_{a^i} \pi_{\theta_{\mathcal{B}_i}}(a^{\mathcal{B}_i}|s, a^i, a^{\mathcal{F}_{\mathcal{B}_i} \setminus \{i\}}) \nabla_{\theta_i} g(\theta_i, \varepsilon)}_{\text{Peer Learning}} \\ & Q_{\pi}(s, \mathbf{a}) d\mathbf{a}^{-i} da^i ds, \end{aligned} \quad (5)$$

where $\mathcal{F}_{\mathcal{B}_i}$ indicates the set of agents on which \mathcal{B}_i have forward dependencies.

Proof. See the Appendix for detailed proof. \square

From Theorem 1, we note that the policy gradient for agent i at each state has two primary terms. The first term $\nabla_{\theta_i} \pi_{\theta_i}(a^i|s, a^{\mathcal{F}_i})$ corresponds to the independent multi-agent policy gradient which explicitly differentiates through π_{θ_i} with respect to the current parameters θ_i . This enables agent i to model its own learning. By contrast, the second term $\nabla_{a^i} \pi_{\theta_{\mathcal{B}_i}}(a^{\mathcal{B}_i}|s, a^i, a^{\mathcal{F}_{\mathcal{B}_i} \setminus \{i\}}) \nabla_{\theta_i} g(\theta_i, \varepsilon)$ aims to additionally account for how the consequences of the corresponding action on its backward dependent agents' policies influence its direction of performance improvement. As a result, the peer learning term enables agents to adjust their own policies to those of action partners, which facilitates fast and accurate inter-agent coordination. Interestingly, the peer learning term, which evaluates the impact of an agent's actions on its peer agents, specifies auxiliary rewards for adapting its policy in accordance with these reactions.

With Theorem 1 and an auto-regressive joint policy, we are ready to present the learning framework of our BPTA-augmented auto-regressive policy gradient algorithm. As illustrated in Figure 1, in the forward process, direct connections and skip connections (He et al. 2015) connect the

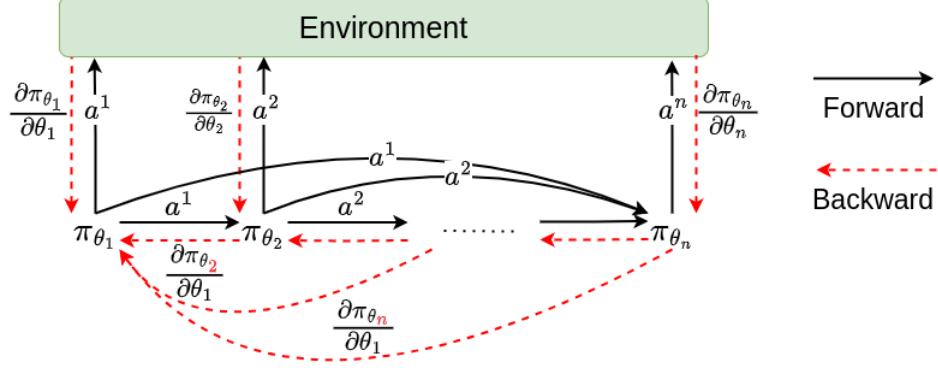


Figure 1: Learning framework of BPTA-augmented auto-regressive policy gradient algorithm. BPTA internally completes the causality loop by two processes: (1) the forward process, which generates the causal flow by action; and (2) the backward process, which inverts the causal flow by propagating the feedback.

action of one predecessor agent to the input of subsequent agents, even those are not adjacent to it in execution order. As for the backward process described by dashed lines, in addition to the interactive feedback from the environment, there are alternative pathways provided by direct and skip connections, which allows successors to provide feedback to predecessors through gradients. Furthermore, these two types of processes are interleaved to allow for a causal flow loop within and across agents.

Our proposed algorithm can be conveniently integrated into most PG-based methods. Given the empirical performance and monotonic policy improvement of PPO (Schulman et al. 2017), we propose *Bidirectional Proximal Policy Optimisation* (BPPO) to incorporate the proposed theorem with PPO. Following the sequential decision scheme, it is intuitive for BPPO to adopt the sequential update scheme (Wang et al. 2023; Kuba et al. 2022), where the updates are performed in reverse order of the execution sequence. We provide comparisons of the simultaneous update scheme and the sequential update scheme in Appendix.

Corollary 1.1 (Clipping Objective of BPPO). *Let π be an auto-regressive joint policy with fixed execution order $\{1, 2, \dots, n\}$, and $\tilde{\pi}^{i+1:n}$ be the updated joint policy of agents set $\{i+1, \dots, n\}$. For brevity, we omit the preceding actions in the policy. Then the clipping objective of BPPO is:*

$$\begin{aligned} & \mathbb{E}_{s \sim \rho^\pi(s), a \sim \pi, \varepsilon \sim p(\varepsilon)} \left[\min \left(\frac{\pi_{\theta_i}(a^i|s)}{\pi_{\theta_i}^{old}(a^i|s)} M^{i+1:n} + \right. \right. \\ & \quad \text{detach} \left(\frac{\pi_{\theta_i}(a^i|s)}{\pi_{\theta_i}^{old}(a^i|s)} \right) \nabla_{a^i} M^{i+1:n} g(\theta_i, \varepsilon) \hat{A}(s, a), \\ & \quad \left. \left. \text{clip} \left(\frac{\pi_{\theta_i}(a^i|s)}{\pi_{\theta_i}^{old}(a^i|s)}, 1 \pm \epsilon \right) M^{i+1:n} + \right) \right. \\ & \quad \text{detach} \left(\text{clip} \left(\frac{\pi_{\theta_i}(a^i|s)}{\pi_{\theta_i}^{old}(a^i|s)}, 1 \pm \epsilon \right) \right) \\ & \quad \left. \left. \nabla_{a^i} M^{i+1:n} g(\theta_i, \varepsilon) \hat{A}(s, a) \right) \right], \end{aligned} \quad (6)$$

$$\text{where } M^{i+1:n} = \frac{\tilde{\pi}^{i+1:n}(a^{i+1:n}|s)}{\pi^{i+1:n}(a^{i+1:n}|s)}, \quad \nabla_{a^i} M^{i+1:n} =$$

$\nabla_{a^i} \tilde{\pi}^{i+1:n}(a^{i+1:n}|s)$, and $\text{detach}()$ represents detaching the input from the computation graph, meaning that the input will not contain gradients. $\hat{A}(s, a)$ is an estimator of the advantage function $A(s, a)$ computed by GAE (Schulman et al. 2018).

We provide the pseudo-code for BPPO in Algorithm 1.

Algorithm 1 Bidirectional Proximal Policy Optimisation

Initialize: The auto-regressive joint policy $\pi = \{\pi_{\theta_1}, \dots, \pi_{\theta_n}\}$, the global value function V , replay buffer B , and the execution order $\{1, \dots, n\}$.

- 1: **for** episode $k = 0, 1, \dots$ **do**
- 2: Collect a set of trajectories by sequentially executing policies according to the execution order;
- 3: Push data into B .
- 4: Compute the advantage approximation $\hat{A}(s, a)$ with GAE.
- 5: Compute the value target $v(s) = \hat{A}(s, a) + V(s)$.
- 6: Set agent i 's gradient w.r.t. agent j 's action $\{c_i^j = 0 \mid i \in \mathcal{N}, j \in \mathcal{N}\}$ and $M^{n+1} = 1$.
- 7: **for** agent $i = n, n-1, \dots, 1$ **do**
- 8: Generate $g(\theta_i, \varepsilon)$ based on the reparameterization trick.
- 9: Compute $\nabla_{a^i} M^{i+1:n}$ based on $\{c_j^i = 0 \mid j \in \{i+1, \dots, n\}\}$ and the chain rule.
- 10: Optimize Eq. 6 w.r.t. θ_i .
- 11: **for** agent $j = 1, \dots, i-1$ **do**
- 12: Compute the gradient c of $\frac{\pi_{\theta_i}(a^i|s)}{\pi_{\theta_i}^{old}(a^i|s)}$ w.r.t. a^j .
- 13: Set $c_i^j = c$.
- 14: **end for**
- 15: Compute $M^{i:n} = \frac{\pi_{\theta_i}(a^i|s)}{\pi_{\theta_i}^{old}(a^i|s)} M^{i+1:n}$.
- 16: **end for**
- 17: Update the value function by the following formula:
- 18: $V = \text{argmin}_V \mathbb{E}_{s \sim \rho^\pi(s)} [\|v(s) - V(s)\|^2]$.
- 19: **end for**

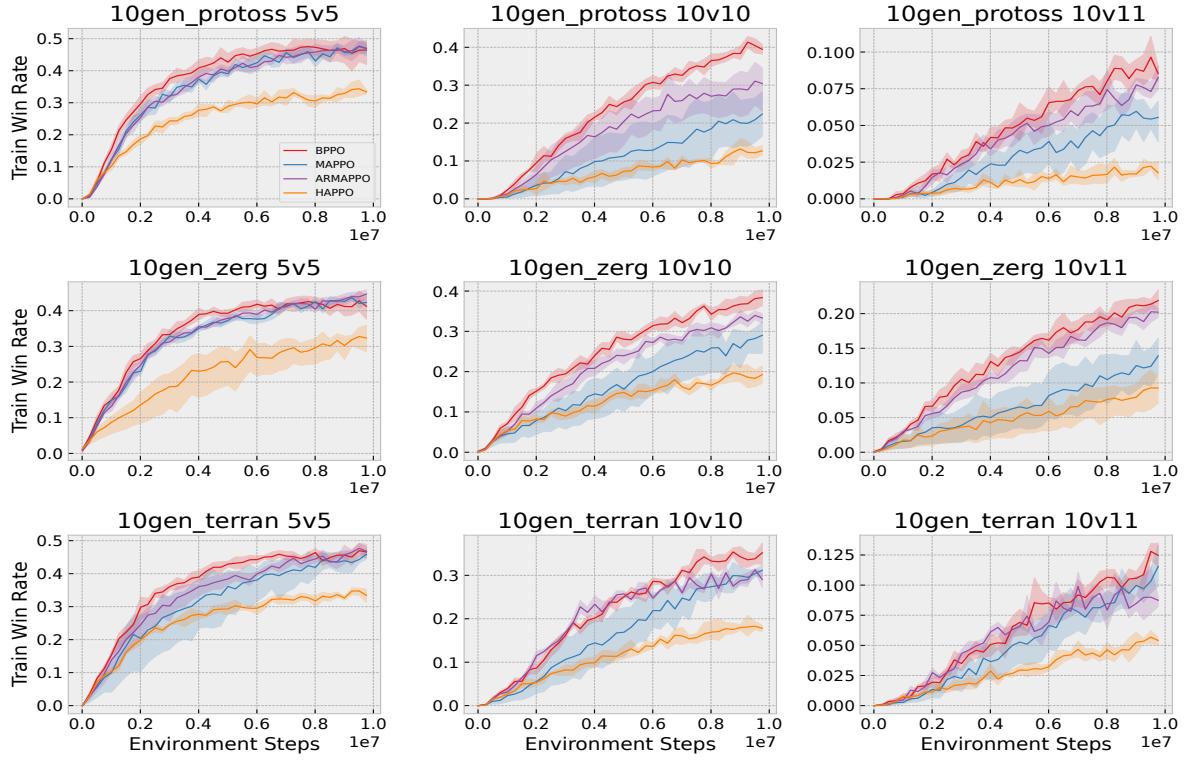


Figure 2: Comparison of training results on SMACv2.

Experiments

In this section, we experimentally evaluate BPPO on several multi-agent benchmarks, including two matrix games (Claus and Boutilier 1998b), the StarCraft Multi-Agent Challenge Version 2 (SMACv2) (Ellis et al. 2022), the Multi-agent MuJoCo (MA-MuJoCo) (Peng et al. 2021), and the Google Research Football (GRF) (Kurach et al. 2020), comparing them against MAPPO (Yu et al. 2022), HAPPO (Kuba et al. 2022), and Auto-Regressive MAPPO (ARMAPPO) (Fu et al. 2022). All results are presented using the mean and standard deviation of five random seeds. We fixed the execution order as sequential in all experiments. Additionally, we compared the effects of different execution orders in Appendix. More experimental details and results on these tasks are included in Appendix.

		Player 2		
		A	B	C
Player 1	A	11	-30	0
	B	-30	7	0
	C	0	6	5

Table 1: Payoff Matrix of the Climbing game.

Matrix Games

As presented in Table 1 and 2, the Climbing game and the Penalty game (Claus and Boutilier 1998b) are shared-reward

		Player 2		
		A	B	C
Player 1	A	-100	0	10
	B	0	2	0
	C	10	0	-100

Table 2: Payoff Matrix of the Penalty game.

multi-agent matrix games with two players in which each player has three actions at their disposal. The two matrix games have several Nash equilibria, but only one or two Pareto-optimal Nash equilibria (Christanios, Papoudakis, and Albrecht 2022). Although stateless and with simple action space, the matrix games are difficult to solve as the agents need to coordinate among two optimal joint actions. Figure 4 shows that the compared baselines will converge to a locally optimal policy while BPPO is the only method that converges to the Pareto-optimal equilibria in all games. This is because BPPO explicitly considers the dependency success to find the optimal joint policy. The gap between the proposed method and the baselines is possibly due to that agents are fully independent of each other when making decisions in those methods. Interestingly, we observe that even with an auto-regressive policy, ARMAPPO still fails to find the optima. However, when we project the preceding actions inputted to each agent in ARMAPPO to higher-dimensional vectors, ARMAPPO w/ PROJ successfully converges to the optimal policy (verified in Appendix).

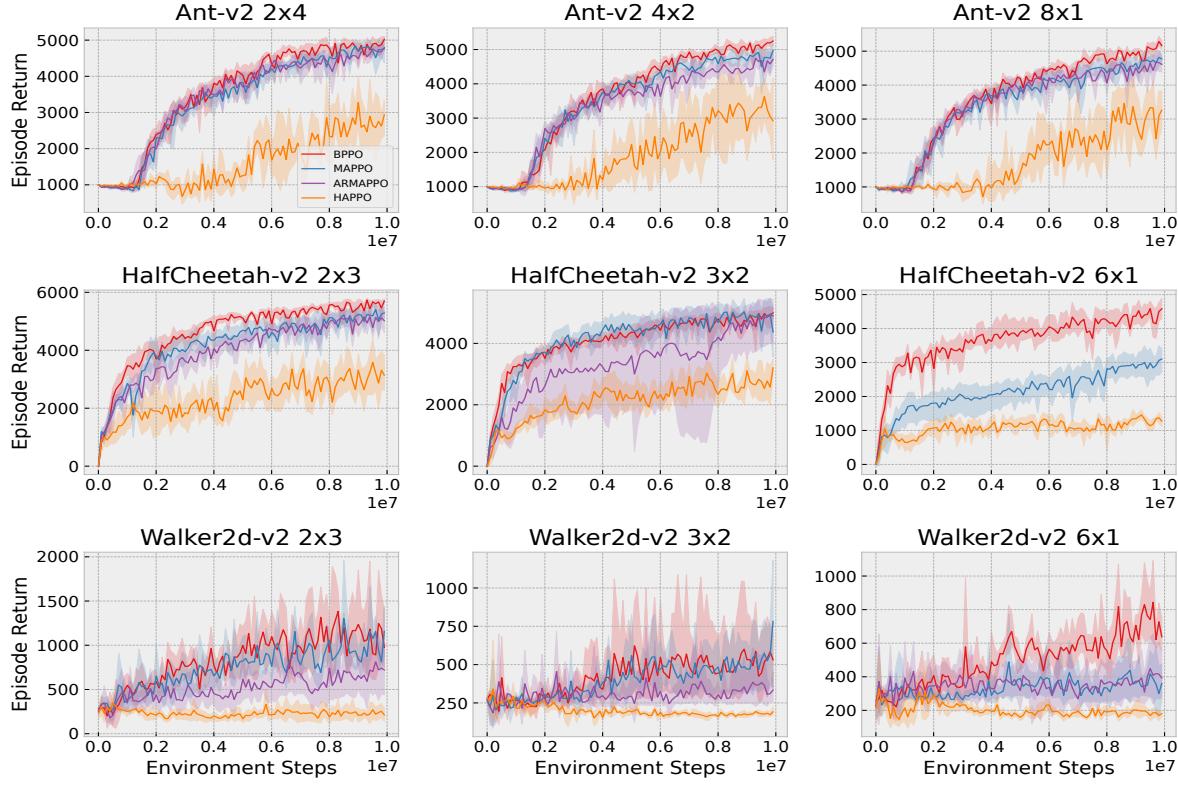


Figure 3: Performance comparison on multiple Multi-Agent MuJoCo tasks.

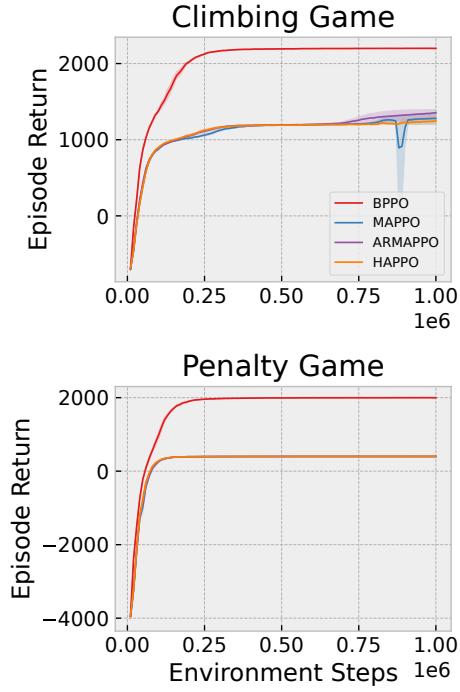


Figure 4: Performance comparison on the Climbing game and the Penalty game.

SMACv2

In SMAC, a group of learning agents aims to defeat the units of the enemy army controlled by the built-in heuristic AI. Despite its popularity, the SMAC is restricted to limited stochasticity (Ellis et al. 2022). Compared to the SMAC, we instead evaluate our method on the more challenging SMAC-v2 benchmark which is designed with higher randomness. We evaluate our method on 3 maps (Zerg, Terran, and Protoss) with symmetric (20-vs-20) and asymmetric (10-vs-11 and 20-vs-23) units. As shown in Figure 2, we generally observe that BPPO outperforms the baselines across most scenarios. In three 10 vs 10 scenarios, the margin between BPPO and the baselines becomes larger. Additionally, we also observe that BPPO has better stability as the variance shows.

MA-MuJoCo

Multi-Agent MuJoCo is a novel benchmark for decentralized cooperative continuous multi-agent robotic control in which single robots are decomposed into individual segments controlled by different agents. We show the performance comparison against the baselines in Figure 3. We can see that BPPO achieves comparable performance compared to the baselines in most tasks while superiorly outperforming others in certain scenarios. It is also worth noting that the observed performance gap between BPPO and ARMAPPO can be attributed to the effectiveness of backward depen-

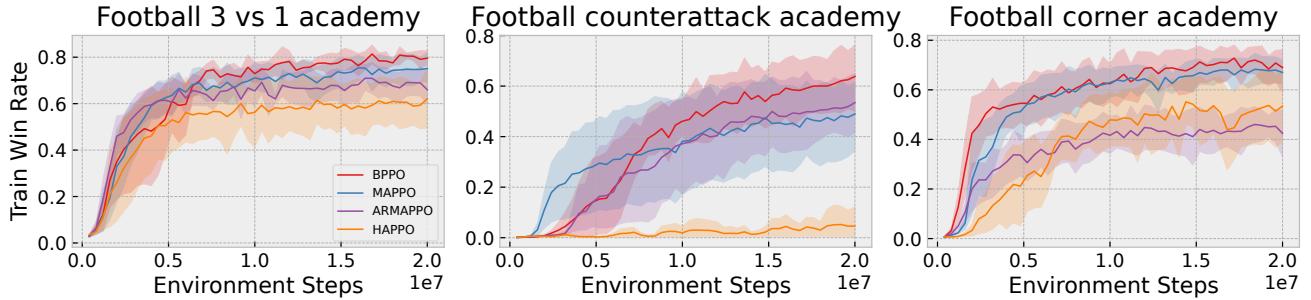


Figure 5: Averaged train win rate on the Google Research Football scenarios.

dency. Meanwhile, we can observe that the performance gap between BPPO and its rivals enlarges with the increasing number of agents. Especially in HalfCheetah-v2 6x1 and Walker2d-v2 6x1, when other algorithms fail to learn any meaningful joint policies or converge to suboptimal points, BPPO outperforms the baselines by a large margin. Interestingly, especially in HalfCheetah 6x1 task, the performance of ARMAPPO even drops to negative. These results show that BPPO enables agents to achieve consistent joint improvement.

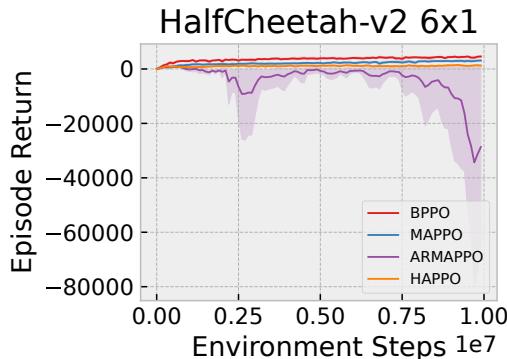


Figure 6: Performance comparison on the HalfCheetah-v2 6x1 Multi-Agent MuJoCo task. ARMAPPO performs poorly, with even negative rewards.

GRF

Google Research Football is a complex environment with large action space and sparse rewards where agents aim to score goals against fixed rule-based opponents. We evaluate our method in both GRF academy scenarios (3-vs-1 with keeper, corner, and counterattack hard) and full-game scenarios (5-vs-5). As can be seen in Figure 5, in the academy scenarios, only a minor difference can be observed between the proposed method and the baselines except for the counterattack task. Additionally, Figure 7 shows that BPPO gains the highest score in the complex 5 vs 5 full-game scenario, while the baselines barely learn anything. Despite the negative performance of all algorithms, BPPO has achieved im-

proved returns and is still learning compared to other algorithms that have consistently maintained their initial values.

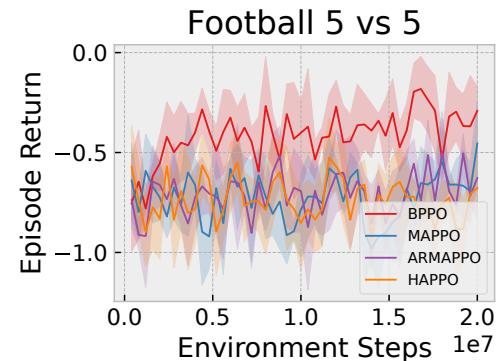


Figure 7: Comparisons of averaged return on the 5-vs-5 scenario.

Conclusion

In this paper, we propose Back-Propagation Through Agents (BPTA) to enable bidirectional dependency in any action-dependent multi-agent policy gradient (MAPG) methods. By conditional multi-agent stochastic policy gradient theorem, we can directly model both an agent’s own action effect and the feedback from its backward dependent agents. We evaluate the proposed Bidirectional Proximal Policy Optimisation (BPPO) based on BPTA and auto-regressive policy on several multiagent benchmarks. Results show that BPPO improves the performance against current state-of-the-art MARL methods. For future work, we plan to study the methods to learn the adaptive order.

Acknowledgments

Zhiyuan Li acknowledges the financial support from the China Scholarship Council (CSC).

References

Bertsekas, D. 2021. Multiagent Reinforcement Learning: Rollout and Policy Iteration. *IEEE/CAA Journal of Automatica Sinica*, 8(2): 249–272.

Box, G.; Jenkins, G.; Reinsel, G.; and Ljung, G. 2015. *Time Series Analysis: Forecasting and Control*. Wiley Series in Probability and Statistics. Wiley. ISBN 9781118674925.

Cho, K.; van Merriënboer, B.; Gülcühre, Ç.; Bahdanau, D.; Bougares, F.; Schwenk, H.; and Bengio, Y. 2014. Learning Phrase Representations using RNN Encoder-Decoder for Statistical Machine Translation. In Moschitti, A.; Pang, B.; and Daelemans, W., eds., *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing, EMNLP 2014, October 25-29, 2014, Doha, Qatar, A meeting of SIGDAT, a Special Interest Group of the ACL*, 1724–1734. ACL.

Christianos, F.; Papoudakis, G.; and Albrecht, S. V. 2022. Pareto Actor-Critic for Equilibrium Selection in Multi-Agent Reinforcement Learning. arXiv:2209.14344.

Claus, C.; and Boutilier, C. 1998a. The Dynamics of Reinforcement Learning in Cooperative Multiagent Systems. In *Proceedings of the Fifteenth National/Tenth Conference on Artificial Intelligence/Innovative Applications of Artificial Intelligence, AAAI '98/IAAI '98*, 746–752. USA: American Association for Artificial Intelligence. ISBN 0262510987.

Claus, C.; and Boutilier, C. 1998b. The dynamics of reinforcement learning in cooperative multiagent systems. *AAAI/IAAI*, 1998(746-752): 2.

Ellis, B.; Moalla, S.; Samvelyan, M.; Sun, M.; Mahajan, A.; Foerster, J. N.; and Whiteson, S. 2022. SMACv2: An Improved Benchmark for Cooperative Multi-Agent Reinforcement Learning. arXiv:2212.07489.

Foerster, J. N.; Farquhar, G.; Afouras, T.; Nardelli, N.; and Whiteson, S. 2018. Counterfactual Multi-Agent Policy Gradients. In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence and Thirtieth Innovative Applications of Artificial Intelligence Conference and Eighth AAAI Symposium on Educational Advances in Artificial Intelligence, AAAI'18/IAAI'18/EAAI'18*. AAAI Press. ISBN 978-1-57735-800-8.

Fu, W.; Yu, C.; Xu, Z.; Yang, J.; and Wu, Y. 2022. Revisiting Some Common Practices in Cooperative Multi-Agent Reinforcement Learning. In Chaudhuri, K.; Jegelka, S.; Song, L.; Szepesvári, C.; Niu, G.; and Sabato, S., eds., *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, 6863–6877. PMLR.

Haarnoja, T.; Zhou, A.; Hartikainen, K.; Tucker, G.; Ha, S.; Tan, J.; Kumar, V.; Zhu, H.; Gupta, A.; Abbeel, P.; and Levine, S. 2019. Soft Actor-Critic Algorithms and Applications. arXiv:1812.05905.

He, K.; Zhang, X.; Ren, S.; and Sun, J. 2015. Deep Residual Learning for Image Recognition. arXiv:1512.03385.

Kingma, D. P.; Salimans, T.; and Welling, M. 2015. Variational Dropout and the Local Reparameterization Trick. In Cortes, C.; Lawrence, N.; Lee, D.; Sugiyama, M.; and Garnett, R., eds., *Advances in Neural Information Processing Systems*, volume 28. Curran Associates, Inc.

Kingma, D. P.; and Welling, M. 2022. Auto-Encoding Variational Bayes. arXiv:1312.6114.

Kuba, J. G.; Chen, R.; Wen, M.; Wen, Y.; Sun, F.; Wang, J.; and Yang, Y. 2022. Trust Region Policy Optimisation in Multi-Agent Reinforcement Learning. arXiv:2109.11251.

Kurach, K.; Raichuk, A.; Stańczyk, P.; Zajac, M.; Bachem, O.; Espeholt, L.; Riquelme, C.; Vincent, D.; Michalski, M.; Bousquet, O.; and Gelly, S. 2020. Google Research Football: A Novel Reinforcement Learning Environment. *Proceedings of the AAAI Conference on Artificial Intelligence*, 34(04): 4501–4510.

Li, C.; Liu, J.; Zhang, Y.; Wei, Y.; Niu, Y.; Yang, Y.; Liu, Y.; and Ouyang, W. 2023. ACE: Cooperative Multi-agent Q-learning with Bidirectional Action-Dependency. In *Proceedings of the AAAI Conference on Artificial Intelligence*.

Littman, M. L. 1994. Markov Games as a Framework for Multi-Agent Reinforcement Learning. In *Proceedings of the Eleventh International Conference on International Conference on Machine Learning, ICML'94*, 157–163. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc. ISBN 1558603352.

Lowe, R.; Wu, Y.; Tamar, A.; Harb, J.; Abbeel, P.; and Mordatch, I. 2017. Multi-Agent Actor-Critic for Mixed Cooperative-Competitive Environments. In *Proceedings of the 31st International Conference on Neural Information Processing Systems, NIPS'17*, 6382–6393. Red Hook, NY, USA: Curran Associates Inc. ISBN 9781510860964.

Ma, J.; and Wu, F. 2020. Feudal Multi-Agent Deep Reinforcement Learning for Traffic Signal Control. In Seghrouchni, A. E. F.; Sukthankar, G.; An, B.; and Yorke-Smith, N., eds., *Proceedings of the 19th International Conference on Autonomous Agents and Multiagent Systems, AAMAS '20, Auckland, New Zealand, May 9-13, 2020*, 816–824. International Foundation for Autonomous Agents and Multiagent Systems.

Müller, V.; Ohström, K.-R. P.; and Lindenberger, U. 2021. Interactive brains, social minds: Neural and physiological mechanisms of interpersonal action coordination. *Neuroscience & Biobehavioral Reviews*, 128: 661–677.

Oliehoek, F. A.; and Amato, C. 2016. *A Concise Introduction to Decentralized POMDPs*. Springer Publishing Company, Incorporated, 1st edition. ISBN 3319289276.

Peng, B.; Rashid, T.; Schroeder de Witt, C.; Kamienny, P.-A.; Torr, P.; Böhmer, W.; and Whiteson, S. 2021. Facmac: Factored multi-agent centralised policy gradients. *Advances in Neural Information Processing Systems*, 34: 12208–12221.

Rashid, T.; Samvelyan, M.; De Witt, C. S.; Farquhar, G.; Foerster, J.; and Whiteson, S. 2020. Monotonic Value Function Factorisation for Deep Multi-Agent Reinforcement Learning. *J. Mach. Learn. Res.*, 21(1).

Ruan, J.; Du, Y.; Xiong, X.; Xing, D.; Li, X.; Meng, L.; Zhang, H.; Wang, J.; and Xu, B. 2022. GCS: Graph-

Based Coordination Strategy for Multi-Agent Reinforcement Learning. In *Proceedings of the 21st International Conference on Autonomous Agents and Multiagent Systems*, AAMAS '22, 1128–1136. Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems. ISBN 9781450392136.

Schulman, J.; Moritz, P.; Levine, S.; Jordan, M.; and Abbeel, P. 2018. High-Dimensional Continuous Control Using Generalized Advantage Estimation. arXiv:1506.02438.

Schulman, J.; Wolski, F.; Dhariwal, P.; Radford, A.; and Klimov, O. 2017. Proximal Policy Optimization Algorithms. arXiv:1707.06347.

Shalev-Shwartz, S.; Shammah, S.; and Shashua, A. 2016. Safe, Multi-Agent, Reinforcement Learning for Autonomous Driving. *CoRR*, abs/1610.03295.

Son, K.; Kim, D.; Kang, W. J.; Hostallero, D. E.; and Yi, Y. 2019. QTRAN: Learning to Factorize with Transformation for Cooperative Multi-Agent Reinforcement Learning. In Chaudhuri, K.; and Salakhutdinov, R., eds., *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, 5887–5896. PMLR.

Sunehag, P.; Lever, G.; Gruslys, A.; Czarnecki, W. M.; Zambaldi, V.; Jaderberg, M.; Lanctot, M.; Sonnerat, N.; Leibo, J. Z.; Tuyls, K.; and Graepel, T. 2018. Value-Decomposition Networks For Cooperative Multi-Agent Learning Based On Team Reward. In *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems*, AAMAS '18, 2085–2087. Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems.

Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, L. u.; and Polosukhin, I. 2017. Attention is All you Need. In Guyon, I.; Luxburg, U. V.; Bengio, S.; Wallach, H.; Fergus, R.; Vishwanathan, S.; and Garnett, R., eds., *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.

Wan, L.; Liu, Z.; Chen, X.; Lan, X.; and Zheng, N. 2022. Greedy based Value Representation for Optimal Coordination in Multi-agent Reinforcement Learning. In Chaudhuri, K.; Jegelka, S.; Song, L.; Szepesvari, C.; Niu, G.; and Sabato, S., eds., *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, 22512–22535. PMLR.

Wang, J.; Ren, Z.; Liu, T.; Yu, Y.; and Zhang, C. 2021a. {QPLEX}: Duplex Dueling Multi-Agent Q-Learning. In *International Conference on Learning Representations*.

Wang, J.; Ye, D.; and Lu, Z. 2023. More Centralized Training, Still Decentralized Execution: Multi-Agent Conditional Policy Factorization. In *The Eleventh International Conference on Learning Representations*.

Wang, X.; Tian, Z.; Wan, Z.; Wen, Y.; Wang, J.; and Zhang, W. 2023. Order Matters: Agent-by-agent Policy Optimization. In *The Eleventh International Conference on Learning Representations*.

Wang, Y.; Han, B.; Wang, T.; Dong, H.; and Zhang, C. 2021b. {DOP}: Off-Policy Multi-Agent Decomposed Policy Gradients. In *International Conference on Learning Representations*.

Wei, E.; Wicke, D.; Freelan, D.; and Luke, S. 2018. Multiagent Soft Q-Learning. In *2018 AAAI Spring Symposia, Stanford University, Palo Alto, California, USA, March 26-28, 2018*. AAAI Press.

Wen, M.; Kuba, J.; Lin, R.; Zhang, W.; Wen, Y.; Wang, J.; and Yang, Y. 2022. Multi-Agent Reinforcement Learning is a Sequence Modeling Problem. In Koyejo, S.; Mohamed, S.; Agarwal, A.; Belgrave, D.; Cho, K.; and Oh, A., eds., *Advances in Neural Information Processing Systems*, volume 35, 16509–16521. Curran Associates, Inc.

Yang, Q.; Dong, W.; Ren, Z.; Wang, J.; Wang, T.; and Zhang, C. 2022. Self-Organized Polynomial-Time Coordination Graphs. In Chaudhuri, K.; Jegelka, S.; Song, L.; Szepesvari, C.; Niu, G.; and Sabato, S., eds., *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, 24963–24979. PMLR.

Ye, J.; Li, C.; Wang, J.; and Zhang, C. 2023. Towards Global Optimality in Cooperative MARL with the Transformation And Distillation Framework. arXiv:2207.11143.

Yu, C.; Velu, A.; Vinitsky, E.; Gao, J.; Wang, Y.; Bayen, A.; and Wu, Y. 2022. The Surprising Effectiveness of PPO in Cooperative Multi-Agent Games. In Koyejo, S.; Mohamed, S.; Agarwal, A.; Belgrave, D.; Cho, K.; and Oh, A., eds., *Advances in Neural Information Processing Systems*, volume 35, 24611–24624. Curran Associates, Inc.

Yu, C.; Yang, X.; Gao, J.; Chen, J.; Li, Y.; Liu, J.; Xiang, Y.; Huang, R.; Yang, H.; Wu, Y.; and Wang, Y. 2023. Asynchronous Multi-Agent Reinforcement Learning for Efficient Real-Time Multi-Robot Cooperative Exploration. In *Proceedings of the 2023 International Conference on Autonomous Agents and Multiagent Systems*, AAMAS '23, 1107–1115. Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems. ISBN 9781450394321.

Zang, Y.; He, J.; Li, K.; Fu, H.; Fu, Q.; and Xing, J. 2023. Sequential Cooperative Multi-Agent Reinforcement Learning. In *Proceedings of the 2023 International Conference on Autonomous Agents and Multiagent Systems*, AAMAS '23, 485–493. Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems. ISBN 9781450394321.

Zhang, T.; Li, Y.; Wang, C.; Xie, G.; and Lu, Z. 2021. FOP: Factorizing Optimal Joint Policy of Maximum-Entropy Multi-Agent Reinforcement Learning. In Meila, M.; and Zhang, T., eds., *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, 12491–12500. PMLR.