

HCPO: Hierarchical Conductor-Based Policy Optimization in Multi-Agent Reinforcement Learning

Zejiao Liu^{1*}, Junqi Tu^{2*}, Yitian Hong^{2*}, Luolin Xiong², Yaochu Jin^{3†}, Yang Tang^{2†}, Fangfei Li^{1†}

¹The School of Mathematics, East China University of Science and Technology, Shanghai, China

²The Key Laboratory of Smart Manufacturing in Energy Chemical Process, Ministry of Education, East China University of Science and Technology, Shanghai, China

³The School of Engineering, Westlake University, Hangzhou, China

{liujejiao, 23012389, y20200105}@mail.ecust.edu.cn, xiongluolin@gmail.com, jinyaochu@westlake.edu.cn, {yangtang, lifangfei}@ecust.edu.cn

Abstract

In cooperative Multi-Agent Reinforcement Learning (MARL), efficient exploration is crucial for optimizing the performance of joint policy. However, existing methods often update joint policies via independent agent exploration, without coordination among agents, which inherently constrains the expressive capacity and exploration of joint policies. To address this issue, we propose a conductor-based joint policy framework that directly enhances the expressive capacity of joint policies and coordinates exploration. In addition, we develop a Hierarchical Conductor-based Policy Optimization (HCPO) algorithm that instructs policy updates for the conductor and agents in a direction aligned with performance improvement. A rigorous theoretical guarantee further establishes the monotonicity of the joint policy optimization process. By deploying local conductors, HCPO retains centralized training benefits while eliminating inter-agent communication during execution. Finally, we evaluate HCPO on three challenging benchmarks: StarCraft II Multi-agent Challenge, Multi-agent MuJoCo, and Multi-agent Particle Environment. The results indicate that HCPO outperforms competitive MARL baselines regarding cooperative efficiency and stability.

Introduction

Cooperative Multi-Agent Reinforcement Learning (MARL) methods have driven significant progress across various fields, including autonomous driving (Chen et al. 2025), robot cooperative control (Gu et al. 2023), and smart grid (Zhang et al. 2022). However, the increasing number of agents in the environment leads to the exponential growth of the state space and the joint action space, which brings the scalability challenge in MARL. A widely adopted solution to tackle this challenge is the Centralized Training with Decentralized Execution (CTDE) paradigm (Feng et al. 2024; Na and Moon 2024). It updates agents' policies with global information during training, while ensuring that agents make

decisions only based on their own local information during execution. Typical CTDE algorithms such as MADDPG (Lowe et al. 2017), QMIX (Rashid et al. 2020), and MAPPO (Yu et al. 2022) attract widespread attention for their enhanced coordination and overall effectiveness.

Under the CTDE paradigm, efficient exploration is important in MARL (Zheng et al. 2021; Xu, Zhang, and Huang 2023; Zhang et al. 2023b). Since parameter sharing restricts the behavioral diversity among agents, consequently impairing their exploration capabilities and impeding task completion (Li, Pan, and Zhang 2024; Li and Zhu 2025), researchers have developed heterogeneous MARL algorithms (Kuba et al. 2022; Li, Pan, and Zhang 2024). Specifically, to tackle the non-stationarity problem arising from simultaneous agent decisions in heterogeneous settings, sequential update algorithms such as Heterogeneous-Agent Trust Region Policy Optimisation (HATRPO) (Kuba et al. 2022) and Agent-by-agent Policy Optimization (A2PO) (Wang et al. 2023) have been proposed. This technique allows subsequent agents to integrate the action and policy information of previous agents within each training iteration. However, most of the existing CTDE algorithms, such as HATRPO and A2PO, presume that joint policy is expressed as the product of individual policies. This limits the expressive capacity of joint policy, making it difficult for multi-agent systems to explore the optimal joint policy during training. Therefore, we propose a hierarchical joint policy framework to address the limitation in expression and enhance exploration for better performance.

Inspired by soccer training and competition, we propose a conductor-based method to enhance the ability of joint policy expression and exploration, as shown in Figure 1. A centralized conductor provides instructions to the entire team by considering the on-field status of both sides. Specifically, the team can be instructed to adopt offensive instructions such as high pressing, flank attack, or defensive instructions such as zonal marking and low block. Players make decisions according to the instruction of the conductor and their own observations. For example, after the conductor chooses and broadcasts the “high pressing” instruction to all blue

*These authors contributed equally.

†Corresponding authors.

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

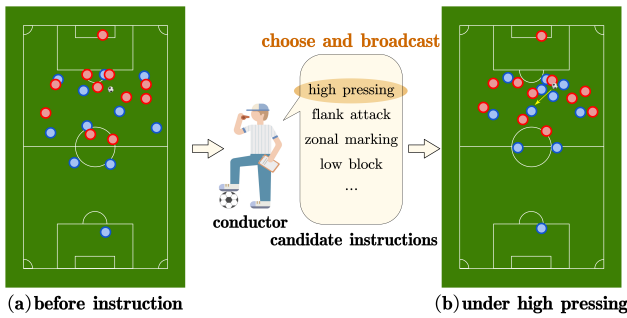


Figure 1: Visualization of conductor-based instructional impact in multi-agent learning: Blue and red dots represent players from opposing teams, with the blue team conductor providing strategic instructions.

teammates, each agent produces its own action by synthesizing this global instruction with its local observation. Consequently, some players drop back to consolidate defensive coverage, while others advance to compress space, jointly generating a coordinated pressing pattern.

In this work, we develop a heterogeneous sequential MARL framework, termed **Hierarchical Conductor-based Policy Optimization (HCPO)**. Firstly, a joint policy model based on the conductor’s instructions is constructed for strong expressive capacity, which is inspired by a Gaussian mixture model. Subsequently, we establish a theoretically grounded update mechanism with strict monotonic improvement guarantees. By decomposing the joint policy optimization into conductor-level instruction and agent-level execution, we derive dual trust region constraints that ensure robust policy evolution. To enable practical decentralized execution, we further deploy local conductors to all agents and adapt their policies through the cross-entropy method with the centralized conductor, eliminating dependency on inter-agent communication. In addition, we perform extensive empirical validation across standardized benchmarks, including StarCraft II Multi-Agent Challenge (SMAC), Multi-agent MuJoCo (MA-MuJoCo), and Multi-agent Particle Environment (MPE) environments. The results show that HCPO outperforms existing competitive MARL algorithms. The main contributions are summarized as follows.

- **Hierarchical Conductor-based Policy Expression:** To enhance policy expressive capacity and guide multi-agent exploration, we propose a conductor-based joint policy framework: $\pi_{\text{mar}}(\mathbf{a}|s) \triangleq \mathbb{E}_{M \sim w(\cdot|s)} \pi(\mathbf{a}|s, M)$.
- **HCPO Algorithm and Monotonic Improvement Guarantee:** With the above conductor-based joint policy expression, we derive a new decomposition approach that breaks down the joint policy’s KL divergence into two components: the conductor policies’ KL divergence and the agent policies’ KL divergence. Then, we present a policy improvement inequality and design a two-level policy update mechanism for the conductor and agents. Finally, we prove that HCPO can ensure a monotonic improvement of the joint policy.

- **Extensive Experimental Validation:** Comprehensive evaluations on MARL benchmarks demonstrate the superior performance of HCPO over strong MARL baselines. The results show improvements in cooperative efficiency, policy stability, and exploration.

Related Work

Exploration: In MARL, the exponential expansion of joint state and action spaces as the number of agents increases severely challenges their ability to efficiently identify high-value states and actions (Liu et al. 2021b). To improve the performance of MARL algorithm, research on the exploration of state and policy spaces is crucial. In terms of state space exploration, the classical work Multi-Agent Variational Exploration (MAVEN) (Mahajan et al. 2019) guides multi-agent systems to learn diverse exploration patterns by maximizing the mutual information between agent trajectories and latent variables. To further improve exploration efficiency, methods like (Jo et al. 2024) and (Li and Zhu 2025) are proposed. Regarding policy space exploration, the current research focuses on constructing policy diversity incentive mechanisms (Dou et al. 2024). For example, the paper (Xu, Zhang, and Huang 2023) proposes an exploration method based on joint policy diversity for sparse-reward multi-agent tasks. It drives agents to explore new policies by maximizing the cross-entropy between the current joint policy and previous joint policies. However, most existing methods, either explicitly or implicitly, assume that the joint policy is the product of individual policies (Kuba et al. 2022; Dou et al. 2024; Jo et al. 2024). This decoupled method limits the expressive capacity of joint policy, and sometimes hampers agents’ exploration during training. To address this, we design a conductor-based framework that enhances the joint policy’s expressive capacity. Specifically, our method provides instructions (latent variables (Mahajan et al. 2019; Ibrahim and Fayad 2022)) to guide agents’ exploration in the policy space during training, enabling them to explore new and potentially high-value policies that traditional methods might overlook.

Hierarchical mechanism: In recent years, hierarchical mechanisms have been increasingly adopted (Vezhnevets et al. 2017; Ahilan and Dayan 2019; Paolo et al. 2025). MAVEN (Mahajan et al. 2019) embeds a latent space for hierarchical control within the CTDE framework, which alleviates the expressive limitations introduced by the monotonicity hypothesis in QMIX (Rashid et al. 2020). Through MAVEN, agents condition their behavior on the shared latent variable to enhance exploration and mitigate issues related to suboptimal policies. Furthermore, HAVEN (Xu et al. 2023) proposes a QMIX-style policy optimization framework with a dual coordination mechanism between layers and agents. Skill discovery represents another key direction in hierarchical MARL (He, Shao, and Ji 2020; Zhang et al. 2023a), enabling agents to autonomously learn diverse teamwork skills without requiring manually designed rewards (Liu et al. 2022, 2025). For example, hierarchical learning with skill discovery method (Yang, Borovikov, and Zha 2020) is a two-level MARL algorithm that uses latent skill variables and intrinsic rewards for unsupervised

skill discovery. The hierarchical multi-agent skill discovery (Yang et al. 2023) further extends the research on skill discovery by introducing team and individual skills. By employing the probabilistic graphical model, it formulates multi-agent skill discovery as an inference problem and leverages transformer structure to assign skills for coordination. To handle dynamic team composition, COPA (Liu et al. 2021a) proposes a coach-player hierarchy where a centralized coach periodically broadcasts strategies derived from global information. This design, however, retains the monotonicity constraint of QMIX and relies on communication during execution. Consequently, existing methods mainly use the monotonicity hypothesis imposed by QMIX and optimize hierarchical policies through variational inference. In contrast, our HCPO provides a theoretical guarantee of monotonic improvement in joint policy performance, without requiring any hypothesis regarding the joint action-value function. Additionally, our approach combines a conductor-based framework with trust region and sequential update methods, diverging from the mutual information and variational inference approaches commonly employed.

Background

Cooperative MARL Problem Formulation

In this paper, we consider fully cooperative multi-agent task as a Decentralized Markov Decision Process (DEC-MDP) (Bernstein et al. 2002; Kuba et al. 2022; Wang et al. 2023), which is usually modeled as a tuple $G = \langle \mathcal{N}, \mathcal{S}, \mathcal{A}, \mathcal{P}, r, \gamma \rangle$. Here, $\mathcal{N} = \{1, 2, \dots, N\}$ is the set of agents. \mathcal{S} and \mathcal{A} represent the state space of the environment and the whole joint action space, respectively. Each agent i takes action $a^i \in \mathcal{A}^i$, and $\mathcal{A} = \prod_{i=1}^N \mathcal{A}^i$. $\mathcal{P} : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow [0, 1]$ is the transition probability function. $r : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$ is the reward function shared by all agents and $\gamma \in [0, 1)$ is the discount factor. To improve coordinated exploration, we propose a hierarchical conductor-based framework that adapts agent policies with the instructions from the conductor. Specifically, let M denote the conductor’s instructional decision, sampled from a *centralized* instruction preference distribution $w(\cdot|s)$. The multi-agent joint policy π_{mar} is formulated as the expectation over all possible instructions M : $\pi_{\text{mar}}(\mathbf{a}|s) \triangleq \mathbb{E}_{M \sim w(\cdot|s)} \pi(\mathbf{a}|s, M)$, where $\mathbf{a} = (a^1, a^2, \dots, a^N) \in \mathcal{A}$ is the joint action. To facilitate decentralized execution, we equip each agent with an independent *local* conductor $w^i(\cdot|o^i)$, sharing the same instruction space as the *centralized* conductor. For clarity, the term “conductor” is primarily used to denote the centralized conductor unless otherwise specified. Hereafter, we assume that the conductor has K discrete instructions available to choose at each time. For any given instruction M , the corresponding instruction-conditional joint policy is defined as the product of individual agent policies conditioned on M , i.e. $\pi(\mathbf{a}|s, M) = \prod_{i=1}^N \pi^i(a^i|s, M)$. Therefore, under the conductor-based framework, our goal is to maximize the expected cumulative reward:

$$J(\pi_{\text{mar}}) \triangleq \mathbb{E}_{\substack{\mathbf{s}_{0:\infty} \\ \rho_{\pi_{\text{mar}}}, \mathfrak{M}_w^{0:\infty}, \mathbf{a}_{0:\infty}^{0:\infty}}} \left[\sum_{t=0}^{\infty} \gamma^t r(s_t, \mathbf{a}_t) \right]. \quad (1)$$

In the above equation, we denote “ $s_{0:\infty} \sim \rho_{\pi_{\text{mar}}}^{0:\infty}$ ” as “ $\mathbf{s}_{0:\infty}$ ”, “ $M_{0:\infty} \sim w_{0:\infty}$ ” as “ $\mathfrak{M}_w^{0:\infty}$ ”, and “ $\mathbf{a}_{0:\infty} \sim \pi_{0:\infty}(M_{0:\infty})$ ” as “ $\mathbf{a}_{0:\infty}$ ” for the sake of brevity. Hereafter, we use this notation wherever no ambiguity arises. Consequently, the state value function $V_{\pi_{\text{mar}}}(s)$ is defined as the expected cumulative return under the multi-agent joint policy π_{mar} :

$$V_{\pi_{\text{mar}}}(s) \triangleq \mathbb{E}_{\substack{\mathfrak{M}_w^{0:\infty}, \mathbf{a}_{0:\infty}^{0:\infty}, \mathbf{s}_{0:\infty}^{1:\infty} \\ \rho_{\pi_{\text{mar}}}}} \left[\sum_{t=0}^{\infty} \gamma^t r_t | s_0 = s \right]. \quad (2)$$

Similarly, we define the state-action value function $Q_{\pi_{\text{mar}}}(s, \mathbf{a}) \triangleq \mathbb{E}_{\substack{\mathbf{s}_{0:\infty}^{1:\infty}, \mathfrak{M}_w^{1:\infty}, \mathbf{a}_{0:\infty}^{1:\infty}}} \left[\sum_{t=0}^{\infty} \gamma^t r_t | s_0 = s, \mathbf{a}_0 = \mathbf{a} \right]$. The joint advantage function is written as: $A_{\pi_{\text{mar}}}(s, \mathbf{a}) \triangleq Q_{\pi_{\text{mar}}}(s, \mathbf{a}) - V_{\pi_{\text{mar}}}(s)$.

Sequential Policy Update Mechanism

To address non-stationarity issues in MARL, the sequential update for agents has been widely investigated (Kuba et al. 2022; Wang et al. 2023; Dou et al. 2024; Wan et al. 2025). HATRPO extends the single-agent Trust Region Policy Optimization (TRPO) to multi-agent domain, utilizing multi-agent advantage decomposition to facilitate the sequential updates of agent policies. In HATRPO, each agent updates its policy parameters through the following protocol:

$$\begin{aligned} \theta_{k+1}^{i_n} &= \arg \max_{\theta^{i_n}} \mathbb{E}_{s, \mathbf{a}^{i_{1:n-1}}, a^{i_n}} \left[A_{\pi_{\theta_k}^{i_n}}(s, \mathbf{a}^{i_{1:n-1}}, a^{i_n}) \right], \\ &\text{subject to } \mathbb{E}_{s \sim \rho_{\pi_{\theta_k}}} \left[\text{D}_{\text{KL}} \left(\pi_{\theta_k^{i_n}}^{i_n}(\cdot|s), \pi_{\theta_{k+1}^{i_n}}^{i_n}(\cdot|s) \right) \right] \leq \delta. \end{aligned} \quad (3)$$

In (3), the compact notation $\mathbb{E}_{s, \mathbf{a}^{i_{1:n-1}}, a^{i_n}}$ stands for the full expectation $\mathbb{E}_{s \sim \rho_{\pi_{\theta_k}}, \mathbf{a}^{i_{1:n-1}} \sim \pi_{\theta_k^{i_{1:n-1}}}^{i_{1:n-1}}, a^{i_n} \sim \pi_{\theta_k^{i_n}}^{i_n}}$. Here, $\theta_{k+1}^{i_n}$

denotes the policy parameter for agent i_n in episode $k+1$. $\text{D}_{\text{KL}}(\cdot, \cdot)$ is the KL-divergence between two policies, and δ is a hyperparameter. It is important to note that during the $k+1$ -th episode, policy $\theta_{k+1}^{i_n}$ leverages the updated policies $\pi_{\theta_{k+1}^{i_{1:n-1}}}^{i_{1:n-1}}$ from the preceding agents $i_{1:n-1}$. This is the core idea behind the sequential update mechanism.

Methods

Value Functions

In the previous section, we presented a conductor-based framework that enhances learning through instruction guidance. To evaluate this framework, we now need quantitative metrics to assess both the learning process and policy performance. First, we define the value function for the conductor’s instruction M as: $Q_{\pi_{\text{mar}}}(M|s) \triangleq \mathbb{E}_{\substack{\mathbf{a}_{0:\infty}^{0:\infty}, \mathbf{s}_{0:\infty}^{1:\infty} \\ \rho_{\pi_{\text{mar}}}, \mathfrak{M}_w^{1:\infty}}} \left[\sum_{t=0}^{\infty} \gamma^t r_t | s_0 = s, M_0 = M \right]$. The instruction advantage function is written as:

$$A_{\pi_{\text{mar}}}(M|s) \triangleq Q_{\pi_{\text{mar}}}(M|s) - V_{\pi_{\text{mar}}}(s). \quad (5)$$

Here, $A_{\pi_{\text{mar}}}(M|s)$ measures the relative benefit of instruction M compared to all other instructions $M' \sim w(\cdot|s)$.

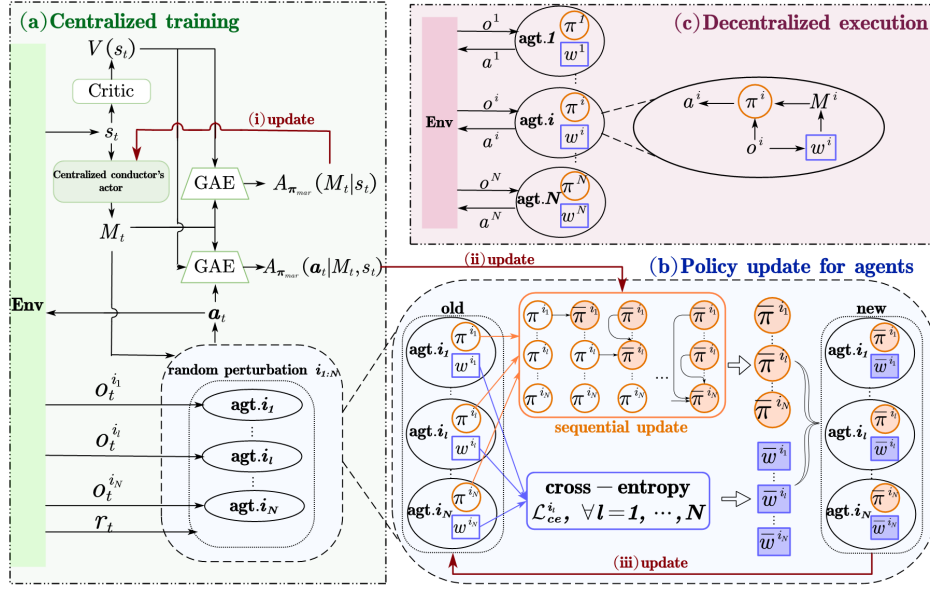


Figure 2: The overall framework of HCPO. **(a) Centralized training:** A two-level policy update mechanism with a virtual centralized conductor is proposed, leveraging well-designed advantage functions. **(b) Policy update for agents:** Here, local agents’ policies denoted by orange ellipses are optimized through sequential updates, and local conductors’ policies denoted by blue rectangles are optimized through the cross-entropy method. During this iteration, policies with shaded outlines represent updated versions, while those without shading indicate unmodified ones. **(c) Decentralized execution:** HCPO enables agents to make decisions based only on local information.

By choosing M to maximize $A_{\pi_{mar}}(M|s)$ with generalized advantage estimation (GAE) (Schulman et al. 2018), we optimize the conductor’s policy w , as illustrated in Figure 2(i). Additionally, we define a joint action advantage function for agents’ joint actions as follows:

$$A_{\pi_{mar}}(\mathbf{a}|s, M) \triangleq Q_{\pi_{mar}}(s, \mathbf{a}) - Q_{\pi_{mar}}(M|s). \quad (6)$$

The advantage function $A_{\pi_{mar}}(\mathbf{a}|s, M)$ evaluates the joint action \mathbf{a} over all other possible joint actions $\mathbf{a}' \sim \pi(\cdot|s, M)$. As illustrated in Figure 2(ii), maximizing this advantage function enables the optimization of the instruction-conditional joint policy $\pi(\cdot|s, M)$, thereby favoring actions that yield superior expected returns. Based on the above definitions, we can find that:

$$A_{\pi_{mar}}(s, \mathbf{a}) = A_{\pi_{mar}}(M|s) + A_{\pi_{mar}}(\mathbf{a}|s, M). \quad (7)$$

After the conductor chooses an instruction $M^j, j \in \{1, 2, \dots, K\}$, we update each individual agent’s policy $\pi^{i_l}(a^{i_l}|s, M^j)$ in the order determined by the random perturbation set $i_{1:N} = \{i_1, i_2, \dots, i_N\}$. In addition, at the state s , based on the conductor’s any instruction M^j , after the previous agents $i_{1:l}$ take joint action $\mathbf{a}^{i_{1:l}}$, we define the expected value of $\mathbf{a}^{i_{1:l}}$ as:

$$Q_{\pi_{mar}}^{i_{1:l}}(\mathbf{a}^{i_{1:l}}|s, M^j) \triangleq \mathbb{E}_{\mathbf{a}_0^{-i_{1:l}} \sim \pi_0^{-i_{1:l}}(\cdot|s, M^j), \mathcal{M}_{\rho_{\pi_{mar}}}^{s^1:\infty}, \mathcal{M}_w^{s^1:\infty}, \mathcal{A}_{\pi}^{s^1:\infty}} \left[\sum_{t=0}^{\infty} \gamma^t r_t | s_0 = s, M_0 = M^j, \mathbf{a}_0^{i_{1:l}} = \mathbf{a}^{i_{1:l}} \right], \quad (8)$$

where $-i_{1:l}$ is the complement of $i_{1:l}$.

Lemma 1. For any instruction $M^j, j \in \{1, 2, \dots, K\}$ chosen by the conductor, the conditional Q -function for agents $i_{1:l}$ satisfies:

$$Q_{\pi_{mar}}^{i_{1:l}}(\mathbf{a}^{i_{1:l}}|s, M^j) = \mathbb{E}_{\mathbf{a}^{-i_{1:l}} \sim \pi^{-i_{1:l}}(\cdot|s, M^j)} [Q_{\pi_{mar}}(s, \mathbf{a})], \quad (9)$$

where $\mathbf{a} = (\mathbf{a}^{i_{1:l}}, \mathbf{a}^{-i_{1:l}})$.

The proof is proposed in Appendix A.1. This lemma provides the basis for the subsequent definitions of advantage functions, which are crucial for evaluating the relative advantages of specific actions compared to average actions. Specifically, at the state s , based on the conductor’s any action M^j , after the agents $i_{1:l}$ first take joint action $\mathbf{a}^{i_{1:l}}$, and agents $-i_{1:l}$ take joint action $\mathbf{a}^{-i_{1:l}} \sim \pi^{-i_{1:l}}(\cdot|s, M^j)$, we define the advantage function for $\mathbf{a}^{i_{1:l}}$ as:

$$A_{\pi_{mar}}^{i_{1:l}}(\mathbf{a}^{i_{1:l}}|s, M^j) \triangleq Q_{\pi_{mar}}^{i_{1:l}}(\mathbf{a}^{i_{1:l}}|s, M^j) - Q_{\pi_{mar}}(M^j|s). \quad (10)$$

It is noted that when $l = 0$, we have $A_{\pi_{mar}}^{i_{1:l}}(\mathbf{a}^{i_{1:l}}|s, M^j) = 0$. Besides, at the state s , based on the conductor’s any instruction M^j , for any individual agent i_l , we define the advantage of its individual action a^{i_l} over all actions $a^{i_l'} \sim \pi^{i_l'}(\cdot|s, M^j)$:

$$A_{\pi_{mar}}^{i_l}(\mathbf{a}^{i_{1:l-1}}, a^{i_l}|s, M^j) \triangleq Q_{\pi_{mar}}^{i_{1:l}}(\mathbf{a}^{i_{1:l}}|s, M^j) - Q_{\pi_{mar}}^{i_{1:l-1}}(\mathbf{a}^{i_{1:l-1}}|s, M^j). \quad (11)$$

In the subsequent part, we introduce the conductor-based multi-agent advantage function decomposition lemma. As shown in Figure 2(ii), the lemma is designed to facilitate the transition from updating the instruction-conditional

joint policy $\pi(\mathbf{a}|s, M^j)$ to updating individual agents' policies $\pi^{i_l}(\mathbf{a}^{i_l}|s, M^j)$, $l \in \mathcal{N}$.

Lemma 2. (Conditional Advantage Decomposition) Consider a cooperative Markov game with a joint policy π_{mar} . For any state s , any instruction M^j , and any subset of agents $i_{1:n} = \{i_1, i_2, \dots, i_n\} \subseteq \mathcal{N}$, the following equation holds for all states s , joint actions $\mathbf{a}^{i_{1:n}}$, and $M^j \sim w$:

$$A_{\pi_{\text{mar}}}^{i_{1:n}}(\mathbf{a}^{i_{1:n}}|s, M^j) = \sum_{l=1}^n A_{\pi_{\text{mar}}}^{i_l}(\mathbf{a}^{i_{1:l-1}}, a^{i_l}|s, M^j). \quad (12)$$

The proof is proposed in Appendix A.1.

Quantifying Policy Updates for HCPO

TRPO (Schulman et al. 2015) is a reinforcement learning algorithm that improves learning stability by constraining policy update magnitude. In this section, we explore a conductor-based mechanism that incorporates TRPO to measure the difference in expected returns between the new policy and its predecessor. This analysis serves as the foundation for designing effective policy update algorithms.

Proposition 1. As defined in Equation (1), the relationship between the expected return of the new policy $\bar{\pi}_{\text{mar}}$ and the old policy π_{mar} is expressed as:

$$J(\bar{\pi}_{\text{mar}}) = J(\pi_{\text{mar}}) + \mathbb{E}_{\tau \sim \bar{\pi}_{\text{mar}}} \left[\sum_{t=0}^{\infty} \gamma^t A_{\pi_{\text{mar}}}(s_t, \mathbf{a}_t) \right], \quad (13)$$

where $\tau := (s_0, M_0, \mathbf{a}_0, s_1, M_1, \mathbf{a}_1, \dots)$.

The proof is proposed in Appendix A.1. Similar to HATRPO, we introduce the approximation function $L_{\pi_{\text{mar}}}(\bar{\pi}_{\text{mar}})$, which serves as an alternative objective function for the new policy's performance function $J(\bar{\pi}_{\text{mar}})$:

$$L_{\pi_{\text{mar}}}(\bar{\pi}_{\text{mar}}) = J(\pi_{\text{mar}}) + \mathbb{E}_{s \sim \rho_{\pi_{\text{mar}}}, M \sim \bar{w}, \mathbf{a} \sim \bar{\pi}} [A_{\pi_{\text{mar}}}(s, \mathbf{a})], \quad (14)$$

where the compact notation $\mathbb{E}_{s \sim \rho_{\pi_{\text{mar}}}, M \sim \bar{w}, \mathbf{a} \sim \bar{\pi}}$ stands for the full expectation $\mathbb{E}_{s \sim \rho_{\pi_{\text{mar}}}, M \sim \bar{w}(\cdot|s), \mathbf{a} \sim \bar{\pi}(\cdot|s, M)}$. Therefore, we can derive the following theorem.

Theorem 1. Under the proposed conductor-based framework, a significant policy improvement inequality holds for the joint policy π_{mar} :

$$\begin{aligned} J(\bar{\pi}_{\text{mar}}) &\geq J(\pi_{\text{mar}}) + \mathbb{E}_{s \sim \rho_{\pi_{\text{mar}}}} \left[\mathbb{E}_{M \sim \bar{w}(\cdot|s)} A_{\pi_{\text{mar}}}(M|s) \right. \\ &\quad - \text{CD}_{\text{KL}}^{\max}(w, \bar{w}) + \mathbb{E}_{M \sim \bar{w}(\cdot|s), \mathbf{a} \sim \bar{\pi}(\cdot|s, M)} A_{\pi_{\text{mar}}}(\mathbf{a}|s, M) \\ &\quad \left. - \max_C \sum_{j=1}^K w(M^j|s) \text{D}_{\text{KL}}(\pi(\cdot|s, M^j), \bar{\pi}(\cdot|s, M^j)) \right], \end{aligned} \quad (15)$$

where $C = \frac{4\gamma \max_{s, \mathbf{a}} |A_{\pi_{\text{mar}}}(s, \mathbf{a})|}{(1-\gamma)^2}$, $\text{D}_{\text{KL}}^{\max}(w_k, \bar{w}) = \max_s \text{D}_{\text{KL}}(w_k(\cdot|s), \bar{w}(\cdot|s))$.

For proof see Appendix A.2. The inequality quantifies the expected return difference between the new policy $\bar{\pi}_{\text{mar}}$ and the existing policy π_{mar} under the conductor-based mechanism. This finding offers precise guidance for the subsequent policy update process.

Guaranteed Monotonic Joint Policy Optimization

In this section, we formulate the policy update mechanisms for both the centralized conductor's policy and individual agents' policies. Our theoretical analysis aims to establish the monotonic improvement guarantee for the conductor-based joint policy π_{mar} through a two-level optimization framework. This guarantee constitutes a critical theoretical foundation, ensuring progressive performance improvement through successive policy iterations while validating the conductor's instructions. The specific two-level policy update mechanisms in episode $k+1$ are as follows:

(i) The conductor's policy $w(\cdot|s) = w_k(\cdot|s)$ is updated first according to the rule:

$$w_{k+1} = \arg \max_{\bar{w}} \left[\mathbb{E}_{s \sim \rho_{\pi_{\text{mar}, k}}, M \sim \bar{w}} A_{\pi_{\text{mar}, k}}(M|s) - \text{CD}_{\text{KL}}^{\max}(w_k, \bar{w}) \right]. \quad (16)$$

(ii) For each M^j where $j \in \{1, 2, \dots, K\}$, the agents sequentially update their policies in accordance with the order $i_{1:N}$, following the update rule:

$$\begin{aligned} &\pi_{k+1}^{i_l}(\cdot|s, M^j) \\ &= \arg \max_{\bar{\pi}^{i_l}(\cdot|s, M^j)} \left[w_{k+1}(M^j|s) L_{\pi_{\text{mar}, k}}^{i_{1:l}} \left(\pi_{k+1}^{i_{1:l-1}}, \bar{\pi}^{i_l}|s, M^j \right) \right. \\ &\quad \left. - \max_C w_k(M^j|s) \text{D}_{\text{KL}}(\pi_k^{i_l}(\cdot|s, M^j), \bar{\pi}^{i_l}(\cdot|s, M^j)) \right], \end{aligned} \quad (17)$$

where $L_{\pi_{\text{mar}, k}}^{i_{1:l}}(\pi_{k+1}^{i_{1:l-1}}, \bar{\pi}^{i_l}|s, M^j) \triangleq$

$\mathbb{E}_{\mathbf{a}^{i_{1:l-1}} \sim \pi_{k+1}^{i_{1:l-1}}, a^{i_l} \sim \bar{\pi}^{i_l}} \left[A_{\pi_{\text{mar}, k}}(\mathbf{a}^{i_{1:l-1}}, a^{i_l}|s, M^j) \right]$. For proof see Appendix A.4. Equation (17) presents a policy update mechanism using w_k to regulate the update magnitude between the agents' new and existing policies. This approach optimizes the exploration-exploitation trade-off, helps avoid local optima, and enhances the adaptability and effectiveness of the policy update process. In HCPO's practical implementation, we employ the CTDE framework to mitigate the limitations posed by communication interference, as shown in Figure 2(c). During training, we incorporate a virtual *centralized* conductor w parameterized by Ψ , and each individual agent is equipped with a *local* conductor network w^i (parameterized by ψ^i) and an actor network π^i (parameterized by θ^i). To maintain the theoretical assumption that agents update their policies using only local observations and the broadcast instruction M , we employ a two-stage training protocol. First, we collect experience with the centralized conductor and optimize its policy parameter Ψ . Then, every agent's actor network θ^i is sequentially updated to improve the policy. Second, we fix Ψ and distill its policy into ψ^i via cross-entropy loss (Chen et al. 2024), enabling fully decentralized execution at evaluation. We summarize our HCPO as Algorithm 1 of Appendix B.

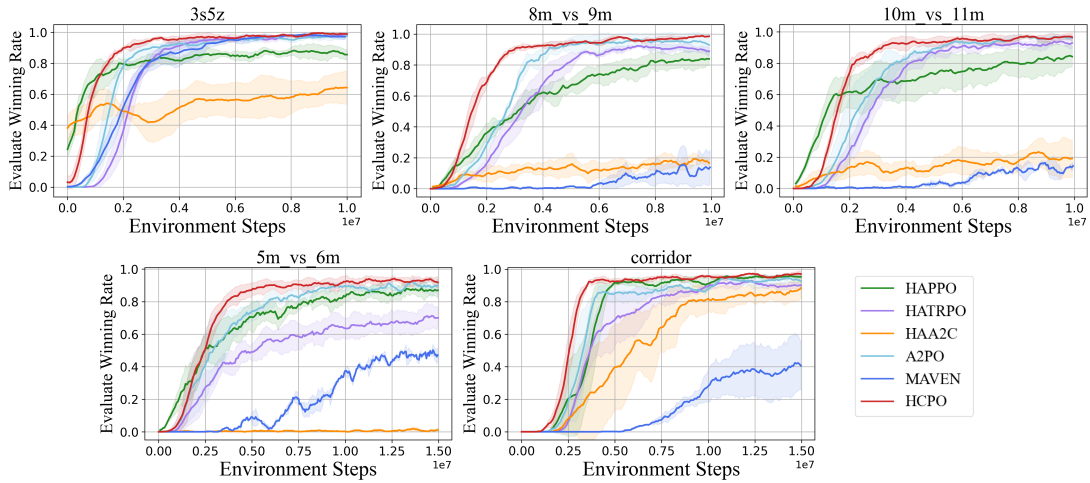


Figure 3: Performance comparison on SMAC. With the conductor-based joint policy enhancing learning efficiency, HCPO reliably outperforms all baselines.

Experiments

In this section, we evaluate HCPO on standard MARL benchmarks including SMAC (Samvelyan et al. 2019), MA-MuJoCo (Li, Pan, and Zhang 2024), and MPE (Lowe et al. 2017). It is compared against strong baselines: HATRPO (Kuba et al. 2022), HAPPO (Kuba et al. 2022), A2PO (Wang et al. 2023), HAA2C (Zhong et al. 2024). Detailed experimental setup and results are presented below.



Figure 4: Effective coordination in SMAC on the 3s5z map: A visual analysis of agent strategies.

Settings and Performance

SMAC: We evaluate algorithms on five SMAC maps, including the widespread hierarchical algorithm MAVEN (Mahajan et al. 2019). Each algorithm is tested with five different random seeds to ensure the robustness and reliability. The results demonstrate that our HCPO achieves outstanding performance in all test maps, as illustrated in Figure 3, with shadows showing standard deviation across different runs. Specifically, HCPO is the first to achieve a 90% winning rate on all maps, enhancing exploration and learning efficiency. Furthermore, HCPO exhibits the lowest standard deviation, demonstrating its high stability. In addition, we visualize the gameplay scenarios on the 3s5z map in Figure 4. Figure 4(a) shows the early game strategy where our team uses one stalker to draw enemy fire, while two other stalkers attack from

behind the zealots. Figure 4(b) presents the overall force division into two groups, with two stalkers providing support. These coordinated actions highlight the agents’ effective collaboration, driven by the latent instructions of HCPO, which ultimately leads to victory. Detailed analysis and the evaluation method are presented in Appendix C.1.

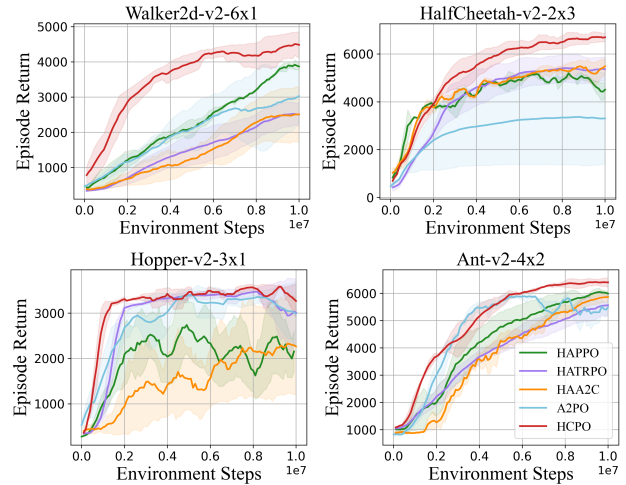


Figure 5: Comparative evaluation on MA-MuJoCo.

MA-MuJoCo: We compare HCPO with four advanced on-policy MARL algorithms using three different random seeds. As shown in Figure 5, HCPO not only achieves significantly higher final returns but also exhibits a lower standard deviation, indicating superior exploration and stability. Specifically, in *HalfCheetah-v2-2x3*, HCPO achieves around 23.42% higher final returns than the next best algorithm HAA2C. In Figure 6, we employ t-SNE (t-Distributed Stochastic Neighbor Embedding) technique to project the states explored by HCPO, HATRPO, and A2PO algorithms during the early stage of training in the *Walker2d-v2-6x1*

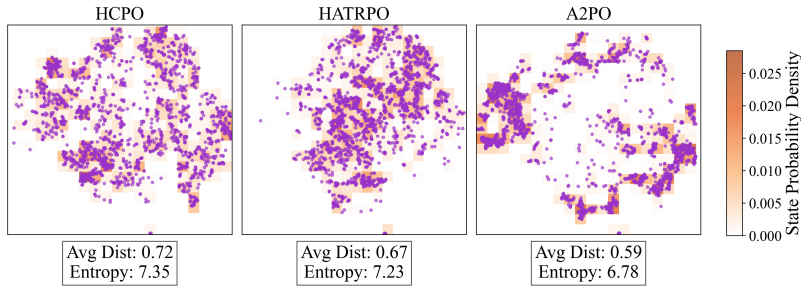


Figure 6: Exploration comparison: t-SNE visualization and entropy analysis in *Walker2d-v2-6x1*.

task onto a 2-D plane. Through entropy analysis and calculation of the average nearest neighbor distance, we can observe that HCPO demonstrates superior exploration. Detailed analysis is presented in Appendix C.2.

MPE: The results on MPE under three random seeds are shown in Figure 7. HCPO exhibits rapid policy improvement in the early stage (0-2 million steps) of training, indicating that HCPO has a high cooperative efficiency and sufficient exploration. Furthermore, compared with HATRPO and A2PO, HCPO shows significant stability and robustness. More experiments are illustrated in Appendix C.3.




Figure 7: Performance comparison of HCPO and other strong MARL algorithms across different MPE tasks.

Ablation Studies

For several key components of HCPO, we conduct ablation studies with results shown in Figure 8. Figure 8(a) examines the impact of the conductor and varying the number of instructions (hyperparameter K) on performance. HCPO with the conductor shows a faster increase in winning rate and a higher final rate than without any conductor, demonstrating its effectiveness in boosting cooperation efficiency. The performance is also influenced by K . It is important to balance performance and resource consumption when selecting K . In Figure 8(b), we examine the hyperparameter δ_1 , which represents the KL-divergence constraint during the conductor’s policy update. In Figure 8(c), we evaluate HCPO under four conductor configurations: a centralized conductor, a random conductor (non-learning baseline), no conductor and local conductors (the core of our HCPO algorithm). Figure 8(d) presents the final episode returns in a boxplot format. The results show that HCPO with local conductors achieves a median return comparable to that of a centralized conductor, while substantially outperforming the variant without any conductor. Furthermore, replacing

the learned instruction preference distribution with a non-learning conductor that outputs uniformly random instructions leads to inferior performance. These findings validate the effectiveness of our proposed update mechanisms. We present detailed analysis in Appendix C.4.

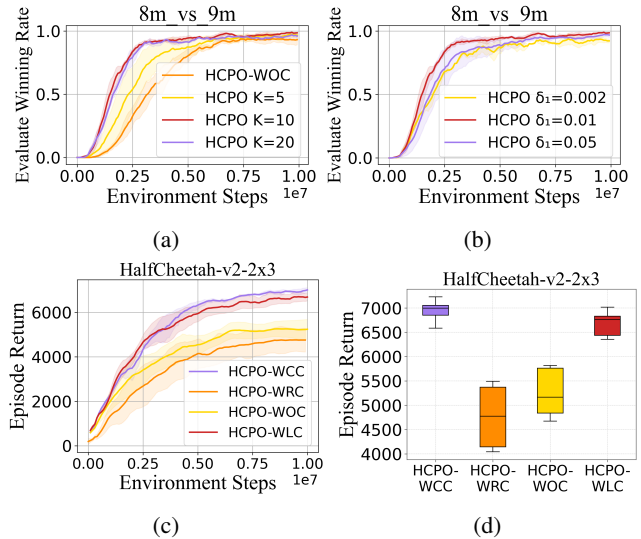


Figure 8: Ablation studies.

Conclusion and Future Work

In this paper, we introduce HCPO, a hierarchical MARL algorithm designed to enhance the expressive capacity of joint policies and improve exploration. By leveraging specific advantage functions, we propose a two-level policy update mechanism with monotonic improvement guarantees without a monotonicity hypothesis in QMIX. Based on the cross-entropy method, each agent is equipped with a local conductor. This not only improves cooperation but also avoids the limitations imposed by communication constraints. Through comprehensive experiments on SMAC, MA-MuJoCo, and MPE, HCPO demonstrates superior performance compared to competitive algorithms. The main limitation of this work is that we design an on-policy algorithm for our hierarchical framework. In the future, we plan to integrate the conductor-based mechanism of HCPO into off-policy algorithms to improve sample efficiency.

Acknowledgments

This work was supported by National Natural Science Foundation of China under Grants 62233005, U2441245, 62573198, 62173142, and in part by the Shanghai Institute for Mathematics and Interdisciplinary Sciences (SIMIS) under Grant SIMIS-ID-2025-SP.

References

- Ahilan, S.; and Dayan, P. 2019. Feudal Multi-Agent Hierarchies for Cooperative Reinforcement Learning. arXiv:1901.08492.
- Bernstein, D. S.; Givan, R.; Immerman, N.; and Zilberstein, S. 2002. The Complexity of Decentralized Control of Markov Decision Processes. *Mathematics of Operations Research*, 27(4): 819–840.
- Chen, X.; Wang, X.; Zhao, W.; Wang, C.; Cheng, S.; and Luan, Z. 2025. Hierarchical Deep Reinforcement Learning based Multi-Agent Game Control for Energy Consumption and Traffic Efficiency Improving of Autonomous Vehicles. *Energy*, 323: 135669.
- Chen, Y.; Mao, H.; Mao, J.; Wu, S.; Zhang, T.; Zhang, B.; Yang, W.; and Chang, H. 2024. PTDE: Personalized Training with Distilled Execution for Multi-Agent Reinforcement Learning. In *Proceedings of the 33rd International Joint Conference on Artificial Intelligence (IJCAI-24)*, 31–39. Jeju, South Korea: IJCAI Organization.
- Dou, H.; Dang, L.; Luan, Z.; and Chen, B. 2024. Measuring Mutual Policy Divergence for Multi-Agent Sequential Exploration. In *Advances in Neural Information Processing Systems (NeurIPS 2024)*, volume 37, 76265–76288.
- Feng, P.; Liang, J.; Wang, S.; Yu, X.; Ji, X.; Chen, Y.; Zhang, K.; Shi, R.; and Wu, W. 2024. Hierarchical Consensus-Based Multi-Agent Reinforcement Learning for Multi-Robot Cooperation Tasks. In *2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 642–649. Abu Dhabi, United Arab Emirates: IEEE.
- Gu, S.; Kuba, J. G.; Chen, Y.; Du, Y.; Yang, L.; Knoll, A.; and Yang, Y. 2023. Safe Multi-Agent Reinforcement Learning for Multi-Robot Control. *Artificial Intelligence*, 319: 103905.
- He, S.; Shao, J.; and Ji, X. 2020. Skill Discovery of Coordination in Multi-agent Reinforcement Learning. arXiv:2006.04021.
- Ibrahim, M.; and Fayad, A. 2022. Hierarchical Strategies for Cooperative Multi-Agent Reinforcement Learning. arXiv:2212.07397.
- Jo, Y.; Lee, S.; Yeom, J.; and Han, S. 2024. FoX: Formation-Aware Exploration in Multi-Agent Reinforcement Learning. In *Proceedings of the 38th AAAI Conference on Artificial Intelligence*, volume 38, 12985–12994. Vancouver, Canada: AAAI Press.
- 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. In *Proceedings of the 10th International Conference on Learning Representations (ICLR 2022)*, 1046. Virtual Only.
- Li, T.; and Zhu, K. 2025. Toward Efficient Multi-Agent Exploration with Trajectory Entropy Maximization. In *Proceedings of the 13th International Conference on Learning Representations (ICLR 2025)*. Singapore.
- Li, X.; Pan, L.; and Zhang, J. 2024. Kaleidoscope: Learnable Masks for Heterogeneous Multi-agent Reinforcement Learning. In *Advances in Neural Information Processing Systems (NeurIPS 2024)*.
- Liu, B.; Liu, Q.; Stone, P.; Garg, A.; Zhu, Y.; and Anandkumar, A. 2021a. Coach-Player Multi-Agent Reinforcement Learning for Dynamic Team Composition. In *Proceedings of the 38th International Conference on Machine Learning (ICML 2021)*, 6860–6870. Virtual Only: PMLR.
- Liu, I.-J.; Jain, U.; Yeh, R. A.; and Schwing, A. 2021b. Cooperative Exploration for Multi-Agent Deep Reinforcement Learning. In *Proceedings of the 38th International Conference on Machine Learning (ICML 2021)*, 6826–6836. Virtual Only: PMLR.
- Liu, S.; Shu, Y.; Guo, C.; and Yang, B. 2025. Learning Generalizable Skills from Offline Multi-Task Data for Multi-Agent Cooperation. In *Proceedings of the 13th International Conference on Learning Representations (ICLR 2025)*. Singapore.
- Liu, Y.; Li, Y.; Xu, X.; Dou, Y.; and Liu, D. 2022. Heterogeneous Skill Learning for Multi-Agent Tasks. In *Advances in Neural Information Processing Systems (NeurIPS 2022)*, volume 35, 37011–37023.
- Lowe, R.; Tamar, A.; Harb, J.; Pieter Abbeel, O.; and Mordatch, I. 2017. Multi-Agent Actor-Critic for Mixed Cooperative-Competitive Environments. In *Advances in Neural Information Processing Systems (NeurIPS 2017)*, volume 30, 6379–6390.
- Mahajan, A.; Rashid, T.; Samvelyan, M.; and Whiteson, S. 2019. MAVEN: Multi-Agent Variational Exploration. In *Advances in Neural Information Processing Systems (NeurIPS 2019)*, volume 32.
- Na, H.; and Moon, I.-C. 2024. LAGMA: Latent Goal-Guided Multi-Agent Reinforcement Learning. In *Proceedings of the 41st International Conference on Machine Learning (ICML 2024)*, 37122–37140. Vienna, Austria: PMLR.
- Paolo, G.; Benechehab, A.; Cherkaoui, H.; Thomas, A.; and Kégl, B. 2025. TAG: A Decentralized Framework for Multi-Agent Hierarchical Reinforcement Learning. arXiv:2502.15425.
- 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. *Journal of Machine Learning Research*, 21(178): 1–51.
- Samvelyan, M.; Rashid, T.; de Witt, C. S.; Farquhar, G.; Nardelli, N.; Rudner, T. G.; Hung, C.-M.; Torr, P. H.; Foerster, J.; and Whiteson, S. 2019. The StarCraft Multi-Agent Challenge. In *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems (AAMAS’19)*, 2186–2188. Montreal, Canada: IFAAMAS.
- Schulman, J.; Levine, S.; Abbeel, P.; Jordan, M.; and Moritz, P. 2015. Trust Region Policy Optimization. In *Proceedings*

of the 32nd International Conference on Machine Learning (ICML 2015), 1889–1897. Lille, France: PMLR.

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

Vezhnevets, A. S.; Osindero, S.; Schaul, T.; Heess, N.; Jaderberg, M.; Silver, D.; and Kavukcuoglu, K. 2017. Feudal Networks for Hierarchical Reinforcement Learning. In *Proceedings of the 34th International Conference on Machine Learning (ICML 2017)*, 3540–3549. Sydney, Australia: PMLR.

Wan, X.; Yang, C.; Song, J.; and Sun, M. 2025. SrSv: Integrating Sequential Rollouts with Sequential Value Estimation for Multi-agent Reinforcement Learning. In *Proceedings of the 39th AAAI Conference on Artificial Intelligence*, volume 39, 23333–23342. Philadelphia, Pennsylvania: AAAI Press.

Wang, X.; Tian, Z.; Wan, Z.; Wen, Y.; Wang, J.; and Zhang, W. 2023. Order Matters: Agent-by-agent Policy Optimization. In *Proceedings of the 11th International Conference on Learning Representations (ICLR 2023)*. Kigali, Rwanda.

Xu, P.; Zhang, J.; and Huang, K. 2023. Exploration via Joint Policy Diversity for Sparse-Reward Multi-Agent Tasks. In *Proceedings of the 32nd International Joint Conference on Artificial Intelligence (IJCAI-23)*, 326–334. Macao, China: IJCAI Organization.

Xu, Z.; Bai, Y.; Zhang, B.; Li, D.; and Fan, G. 2023. HAVEN: Hierarchical Cooperative Multi-Agent Reinforcement Learning with Dual Coordination Mechanism. In *Proceedings of the 37th AAAI Conference on Artificial Intelligence*, volume 37, 11735–11743. Washington, USA: AAAI Press.

Yang, J.; Borovikov, I.; and Zha, H. 2020. Hierarchical Cooperative Multi-Agent Reinforcement Learning with Skill Discovery. In *Proceedings of the 19th International Conference on Autonomous Agents and MultiAgent Systems (AAMAS’20)*. Virtual Only: IFAAMAS.

Yang, M.; Yang, Y.; Lu, Z.; Zhou, W.; and Li, H. 2023. Hierarchical Multi-Agent Skill Discovery. In *Advances in Neural Information Processing Systems (NeurIPS 2023)*, volume 36, 61759–61776.

Yu, C.; Velu, A.; Vinitzky, E.; Gao, J.; Wang, Y.; Bayen, A.; and Wu, Y. 2022. The Surprising Effectiveness of PPO in Cooperative Multi-Agent Games. In *Advances in Neural Information Processing Systems (NeurIPS 2022)*, volume 35, 24611–24624.

Zhang, F.; Jia, C.; Li, Y.-C.; Yuan, L.; Yu, Y.; and Zhang, Z. 2023a. Discovering Generalizable Multi-Agent Coordination Skills from Multi-Task Offline Data. In *Proceedings of the 11th International Conference on Learning Representations (ICLR 2023)*. Kigali, Rwanda.

Zhang, S.; Cao, J.; Yuan, L.; Yu, Y.; and Zhan, D.-C. 2023b. Self-Motivated Multi-Agent Exploration. In *Proceedings of the 22nd International Conference on Autonomous Agents and MultiAgent Systems (AAMAS’23)*, 476–484. London, United Kingdom: IFAAMAS.

Zhang, Y.; Yang, Q.; An, D.; Li, D.; and Wu, Z. 2022. Multistep Multiagent Reinforcement Learning for Optimal Energy Schedule Strategy of Charging Stations in Smart Grid. *IEEE Transactions on Cybernetics*, 53(7): 4292–4305.

Zheng, L.; Chen, J.; Wang, J.; He, J.; Hu, Y.; Chen, Y.; Fan, C.; Gao, Y.; and Zhang, C. 2021. Episodic Multi-Agent Reinforcement Learning with Curiosity-Driven Exploration. In *Advances in Neural Information Processing Systems (NeurIPS 2021)*, volume 34, 3757–3769.

Zhong, Y.; Kuba, J. G.; Feng, X.; Hu, S.; Ji, J.; and Yang, Y. 2024. Heterogeneous-Agent Reinforcement Learning. *Journal of Machine Learning Research*, 25(32): 1–67.