

# MARPO: A Reflective Policy Optimization for Multi-Agent Reinforcement Learning

Cuiling Wu<sup>1,2,\*</sup>, Yaozhong Gan<sup>2,\*</sup>, Junliang Xing<sup>2,†</sup>, Ying Fu<sup>1,†</sup>

<sup>1</sup> School of Computer Science and Technology, Beijing Institute of Technology

<sup>2</sup> QiYuan Lab

wucuilin@bit.edu.cn, yzgan@163.com, xingjunliang@qiyanlab.com, fuying@bit.edu.cn

## Abstract

We propose **Multi-Agent Reflective Policy Optimization (MARPO)** to alleviate the issue of sample inefficiency in multi-agent reinforcement learning. MARPO consists of two key components: a **reflection mechanism** that leverages subsequent trajectories to enhance sample efficiency, and an **asymmetric clipping mechanism** that is derived from the KL divergence and dynamically adjusts the clipping range to improve training stability. We evaluate MARPO in classic multi-agent environments, where it consistently outperforms other methods.

**Code** — <https://github.com/wucuilin12345-blip/MARPO>

## 1 Introduction

Multi-Agent Reinforcement Learning (MARL) has become a vital tool for complex decision-making tasks across areas such as autonomous driving, robotics, and cooperative games. Despite its potential, MARL faces challenges, particularly in achieving sample efficiency in high-dimensional environments. Unlike supervised learning, which uses static datasets, reinforcement learning (RL) relies on agent-environment interactions to collect samples—a costly process in both time and resources.

Early RL advancements, such as Trust Region Policy Optimization (TRPO) (Schulman et al. 2015) and Proximal Policy Optimization (PPO) (Schulman et al. 2017), have improved stability; however, they still face scalability and adaptability issues in real-world environments. Many MARL algorithms, including IPPO (De Witt et al. 2020), MAPPO (Yu et al. 2022), QMIX (Rashid et al. 2020), HATRPO, and HAPPO (Kuba et al. 2021), rely on frequent interactions with the environment and optimize policies based on individual state-action pairs, often underutilizing trajectory-level information that can enhance policy stability.

Recent research explores methods to improve sample efficiency. Trajectory Reuse Optimization, such as R2D2 (Kapturovski et al. 2019), accelerates convergence by prioritiz-

ing experience replay and reusing long sequences. Model-based methods, such as MuZero (Silver et al. 2017), reduce real-world interactions by simulating environment dynamics internally. Additionally, innovations in Sampling Policy and Reward mechanisms, such as Prioritized Experience Replay (PER) (Schaul et al. 2015) and Inverse Reinforcement Learning (IRL) (Ng, Russell et al. 2000), prioritize high-value experiences to improve learning efficiency. Frameworks such as Efficient Episodic Memory Utilization (EMU) (Na, Seo, and Moon 2024) and Episodic Multi-agent Reinforcement Learning with Curiosity-driven exploration (EMC) (Zheng et al. 2021) further enhance exploration and policy convergence.

While advancements in MARL have been made, challenges in sample efficiency, scalability, and trajectory-level utilization remain. To address these, we introduce **Multi-Agent Reflective Policy Optimization (MARPO)**, a framework that leverages trajectory feedback to improve policy optimization efficiency. Unlike methods that rely on auxiliary components such as value functions or human feedback (Deng, Wang, and Zhang 2024; Haarnoja et al. 2018; He et al. 2025), MARPO directly optimizes policies using trajectory signals, thereby enhancing learning efficiency and decision-making. We propose an asymmetric KL-divergence-based clipping mechanism with a dynamic adjustment strategy for the clipping range, further improving training stability and flexibility. We validate MARPO on the StarCraft II Multi-Agent Challenge (SMAC) (Samvelyan et al. 2019), including its more complex SMAC-Hard variants (Deng et al. 2024), SMACv2 (Ellis et al. 2023), and Google Research Football (GRF) (Kurach et al. 2020), demonstrating superior performance over MAPPO in terms of sample efficiency and policy optimization.

**Our contributions can be summarized as follows:**

- We propose MARPO, the first framework to integrate reflection mechanisms into multi-agent policy optimization, improving sample efficiency by leveraging trajectory feedback.
- We derive a KL-based asymmetric clipping mechanism that enables more accurate and flexible policy updates.
- We introduce a dynamic adaptation strategy for the clipping range, which enhances the exploration capability compared to fixed-boundary approaches such as PPO.

\*These authors contributed equally.

†Co-corresponding author.

## 2 Related Work

### 2.1 KL-based Asymmetric Clipping Mechanism

Trust Region Policy Optimization (TRPO) (Schulman et al. 2015) ensures stable policy updates by using a trust-region constraint that prevents substantial updates that could destabilize the learning process. It optimizes the policy using a second-order approximation, specifically the Fisher information matrix, to account for the impact of policy updates on the value function (Prokopenko et al. 2011). The goal is to maximize expected return while keeping the KL divergence between the old and new policies within a predefined threshold, ensuring steady progress and stability during training.

Proximal Policy Optimization (PPO) (Schulman et al. 2017), an extension of TRPO, simplifies computation by using first-order approximations and a clipping mechanism to limit deviations between the new and old policies. This approach improves sample efficiency and computational feasibility, making PPO scalable to large-scale problems. However, the fixed clipping range in PPO is a limitation in dynamic MARL environments, where continuous interactions between agents require more adaptive mechanisms (Massaoudi and Davis 2025; Liu and Liu 2024). This has led to research on dynamic clipping mechanisms, such as Trust Region-Guided Proximal Policy Optimization (TRGPPO) (Wang et al. 2019), which adjusts the clipping range based on KL divergence. This approach enables larger updates in regions with lower KL divergence, promoting exploration, while restricting updates in regions with higher KL divergence to maintain stability. Although it improves sample efficiency and exploration, it relies on careful tuning of the KL threshold.

### 2.2 Reflective Policy Optimization

Reflective Policy Optimization (RPO) (Gan et al. 2024) extends on-policy reinforcement learning methods such as TRPO and PPO by integrating future-trajectory feedback into policy updates. Although PPO is widely adopted for its simplicity and empirical robustness, it remains limited in sample efficiency (Xiong et al. 2023). By leveraging information from the whole trajectory, RPO stabilizes learning and promotes monotonic improvement, enabling agents to adjust actions with a more introspective, trajectory-aware update rule. Proposed initially for single-agent RL, RPO naturally extends to MARL, where agents jointly reflect on both current and future trajectories to improve coordination and accelerate convergence. Similar ideas have also been explored for multi-agent systems under communication delays (Qin et al. 2024), further demonstrating the stability advantages of future-aware updates.

In MARL, MAPPO (Yu et al. 2022) builds on PPO under the CTDE paradigm: agents train with global information while acting from local observations. This mitigates non-stationarity (Hernandez-Leal et al. 2017; Canese et al. 2021) via centralized value functions and stabilized importance-sampling updates. However, its reliance on current state-action pairs restricts the exploitation of future-trajectory structure, limiting stability and scalability in complex environments (Hernandez-Leal et al. 2017; Canese

et al. 2021; Zhang et al. 2025; Jin et al. 2025).

Our work unifies these lines by retaining PPO-style practicality, adding KL-guided adaptive clipping to avoid brittle fixed thresholds, and injecting future-trajectory information into the update. This yields steadier improvement under multi-agent non-stationarity while remaining computationally lightweight. Beyond our on-policy MARL focus, proximal principles also appear in federated offline RL—e.g., proximal evaluation under distributed data constraints (Yue et al. 2024; Zhang, Fu, and Zhang 2024)—and in continuous-control RL via DDPG-family proximal variants (Long et al. 2024), pointing to scalable pathways for deployment.

## 3 Preliminaries

### 3.1 Cooperative MARL Problem Formulation

In the context of MARL, the challenge of partially observable environments is modeled using the Decentralized Partially Observable Markov Decision Process (Dec-POMDP) (Oliehoek and Amato 2016). In the Dec-POMDP framework for cooperative multi-agent tasks, the environment is modeled by a tuple  $G = \langle \mathcal{N}, S, \mathcal{A}, P, O, r, \gamma \rangle$ , where  $S$  is the finite state space, each agent  $i \in \mathcal{N} \equiv \{1, \dots, n\}$  chooses an action  $a_i \in \mathcal{A}$ ,  $\mathcal{A}$  is the individual action space, and  $P(s'|s, A)$  defines the state transition dynamics,  $r : S \times \mathcal{A} \rightarrow \mathbb{R}$  is the reward function, and  $\gamma \in [0, 1]$  is the discount factor. Agents use a policy  $\pi_\theta(a_i|o_i)$  parameterized by  $\theta$  to produce an action  $a_i$  from the local observation  $o_i$ , and jointly optimize the discounted accumulated reward  $J(\theta) = \mathbb{E}_{A^t, s^t} [\sum_t \gamma^t R(s^t, A^t)]$  where  $A^t = (a_1^t, \dots, a_n^t)$  is the joint action at time step  $t$ .

### 3.2 Reflection Mechanism

Reflective Policy Optimization (RPO) extends the clipped surrogate objective to leverage interactions over multiple steps, aiming to address the substantial data requirements per update common to TRPO and PPO, which often lead to sample inefficiency. The surrogate objective function of RPO is defined as:

$$L(\pi, \pi_{\text{old}}) = L_0^{\text{clip}}(\pi, \pi_{\text{old}}) + \alpha L_1^{\text{clip}}(\pi, \pi_{\text{old}}), \quad (1)$$

where

$$L_0^{\text{clip}}(\pi, \pi_{\text{old}}) = \mathbb{E}_{(s,a)} [\min(\rho(a|s)A^{\pi_{\text{old}}}(s,a), \text{clip}(\rho(a|s), 1-\epsilon, 1+\epsilon)A^{\pi_{\text{old}}}(s,a))], \quad (2)$$

$$L_1^{\text{clip}}(\pi, \pi_{\text{old}}) = \mathbb{E}_{(s,a,s',a')} [\min(\rho(a|s)\rho(a'|s')A^{\pi_{\text{old}}}(s',a'), C(\rho, \rho')A^{\pi_{\text{old}}}(s',a'))], \quad (3)$$

and  $\rho(a|s) = \frac{\pi(a|s)}{\pi_{\text{old}}(a|s)}$ ,  $C(\rho, \rho') = \text{clip}(\rho(a|s), 1-\epsilon, 1+\epsilon) \cdot \text{clip}(\rho(a'|s'), 1-\epsilon_1, 1+\epsilon_1)$ ,  $\epsilon, \epsilon_1$  and  $\alpha$  are the hyperparameter. The essence of the reflection mechanism is captured in Eqn. (3), and the theoretical foundation of RPO ensures monotonic improvements in policy performance. As a result, RPO enhances sample efficiency while maintaining stability, addressing the computational and performance challenges in policy optimization.

---

**Algorithm 1: MARPO**

---

**Input:** Initial policy parameters  $\theta$ ; hyperparameter  $\alpha$ **Parameter:** Number of iterations  $n$ , number of epochs  $K$ **Output:** Updated parameters  $\theta$ 

- 1: **for** iteration = 1 to  $n$  **do**
  - 2:   Collect trajectories for all agents and store them in dataset  $D$
  - 3:   **for** epoch = 1 to  $K$  **do**
  - 4:     Sample mini-batch  $b$  from dataset  $D$
  - 5:     Compute clipping bounds  $x_1$  and  $x_2$  by Eqn. (12)
  - 6:     Compute policy loss  $L_\theta(\pi, \pi_{\text{old}})$  by Eqn. (5)
  - 7:     Update policy parameters:  $\theta \leftarrow \theta - \alpha \nabla_\theta L_\theta$
  - 8:   **end for**
  - 9: **end for**
- 

### 3.3 Challenges of Current Policy Gradient Approaches in MARL

Extending Policy Gradient methods to MARL presents significant challenges. A simple approach is to use a shared parameter set for all agents and aggregate their trajectories for policy optimization. This strategy, implemented in MAPPO, optimizes the policy parameter  $\theta$  by maximizing an objective function that combines the benefits of centralized training with decentralized execution, ensuring stability and efficiency in multi-agent environments. The objective function is defined as:

$$L(\theta) = \frac{1}{n} \sum_{i=1}^n \mathbb{E}_{(o_i^k, a_i^k)} \left[ \min \left( \rho_{\theta, i}^k A_i^k, \text{clip}(\rho_{\theta, i}^k, 1 - \epsilon, 1 + \epsilon) A_i^k \right) \right] + \sigma \cdot \frac{1}{n} \sum_{i=1}^n \mathbb{E}_{o_i^k} [S[\pi_\theta^k(o_i^k)]], \quad (4)$$

where  $\rho_{\theta, i}^k = \pi_\theta^k(a_i^k | o_i^k) / \pi_{\theta_{\text{old}}}^k(a_i^k | o_i^k)$  is the policy ratio for agent  $i$ , and  $A_i^k$  is the advantage function computed using the Generalized Advantage Estimation (GAE) method (Schulman et al. 2016). The clip term ensures that policy updates remain within a trust region, thereby mitigating large, destabilizing changes.  $S$  represents the policy entropy, and  $\sigma$  is the entropy coefficient hyper parameter.

## 4 Method

In light of the challenges associated with sample inefficiency, we introduce the Multi-Agent Reflective Policy Optimization (MARPO) framework. Unlike MAPPO, which represents a straightforward extension of PPO to multi-agent environments, MARPO constitutes a natural and principled evolution of RPO. It not only enhances sample efficiency by incorporating a joint reflection mechanism but also improves exploration capabilities via an adaptive clipping mechanism. The details of the algorithm are outlined in Algorithm 1.

### 4.1 Description of the MARPO Method

We propose a multi-agent framework that combines reflection and clipping to enhance sample efficiency. For clarity of exposition, we omit explicit time-step subscripts when not ambiguous. The overall objective function is defined as:

$$L(\pi, \pi_{\text{old}}) = L_0^{\text{clip}}(\pi, \pi_{\text{old}}) + \alpha L_1^{\text{clip}}(\pi, \pi_{\text{old}}), \quad (5)$$

where

$$L_0^{\text{clip}}(\pi, \pi_{\text{old}}) = \frac{1}{n} \sum_{i=1}^n \mathbb{E}_{(o_i^k, a_i^k)} \left[ \min \left( \rho_i^k(a_i^k | o_i^k) \cdot A^{\pi_{\text{old}}^k}, \text{clip}(\rho_i^k(a_i^k | o_i^k), x_1, x_2) \cdot A^{\pi_{\text{old}}^k} \right) \right], \quad (6)$$

and

$$L_1^{\text{clip}}(\pi, \pi_{\text{old}}) = \frac{1}{n} \sum_{i=1}^n \mathbb{E}_{(o_i^k, a_i^k, o_i^{k+1}, a_i^{k+1})} \left[ \min \left( \rho_i^k(a_i^k | o_i^k) \rho_i^{k+1}(a_i^{k+1} | o_i^{k+1}) A^{\pi_{\text{old}}^k}, c(\rho_i^k, \rho_i^{k+1}) A^{\pi_{\text{old}}^{k+1}} \right) \right], \quad (7)$$

where  $\rho_i^k$  denotes the probability ratio between the new policy  $\pi$  and the old policy  $\pi_{\text{old}}$  for agent  $i$  at time step  $k$ , given the local observation  $o_i^k$  and the action  $a_i^k$  taken. Similarly,  $\rho_i^{k+1}(a_i^{k+1} | o_i^{k+1}) = \pi^{k+1}(a_i^{k+1} | o_i^{k+1}) / \pi_{\text{old}}^{k+1}(a_i^{k+1} | o_i^{k+1})$  represents the ratio for the action  $a_i^{k+1}$  at the next local state  $o_i^{k+1}$ . Additionally,  $A^{\pi_{\text{old}}^k}$  is the advantage function calculated from the old policy by GAE, which helps estimate the value of each state-action pair in a multi-agent context. The hyperparameter  $\alpha$  controls the clipping bounds and balances the trade-off between different components of the objective function.

The clipping mechanism ensures that policy updates remain within a reasonable range, thereby avoiding large updates that could destabilize the training process. It is defined as:

$$c(\rho_i^k, \rho_i^{k+1}) = \text{clip}(\rho_i^k(a_i^k | o_i^k), x_1, x_2) \cdot \text{clip}(\rho_i^{k+1}(a_i^{k+1} | o_i^{k+1}), x'_1, x'_2), \quad (8)$$

where  $x_1, x_2$  and  $x'_1, x'_2$  are clipping bounds that regulate the magnitude of policy updates, by incorporating future state-action pairs, agents can account for the long-term impact of their actions, thereby enhancing global coordination and overall policy quality. This integration leverages future state-action information to support more informed and stable policy updates. Our approach enhances coordination, accelerates training, and stabilizes learning, effectively addressing key challenges commonly faced by traditional MARL methods.

### 4.2 Dynamic Adaptive Asymmetric Clipping Mechanism

In the previous section, we introduced a multi-agent reflective mechanism in which the policy optimization objective considers not only the reward at the current time step but also the potential influence of future actions. However, under the reflection mechanism, the two importance sampling ratios  $\rho^k$  and  $\rho^{k+1}$  co-occur in the loss, making the objective more sensitive to policy drift. To tackle this issue, we introduce a Dynamic Asymmetric Clipping Mechanism, which leverages approximate estimates of the KL divergence to dynamically adjust the bounds  $x_1$  and  $x_2$ , thereby enhancing the stability and constraint of policy updates.

**Theoretical Analysis and Enhancement of Clipping Mechanisms Using KL Divergence.** The standard clipping mechanism in PPO can be viewed as a heuristic approximation to the TV distance, designed to constrain the magnitude of policy updates. Specifically, the TV distance between the old and new policies can be expressed as:

$$D_{\text{TV}}(\pi_{\text{old}}||\pi_{\text{new}}) = \frac{1}{2} \sum_a |\pi_{\text{old}} - \pi_{\text{new}}| = \frac{1}{2} \mathbb{E}_{\pi_{\text{old}}} \left| 1 - \frac{\pi_{\text{new}}}{\pi_{\text{old}}} \right|. \quad (9)$$

To ensure stable policy updates, PPO constrains the importance sampling ratio  $\left| 1 - \frac{\pi_{\text{new}}}{\pi_{\text{old}}} \right|$  within a fixed range, effectively bounding policy deviation through a surrogate loss that approximates total variation (TV) distance. While TV distance controls the magnitude of policy shifts, it only reflects the scalar distance between distributions and overlooks the geometric structure of the policy space.

In contrast, KL divergence provides an asymmetric and differentiable measure that more accurately captures directional changes in policies, making it better suited for gradient-based optimization, particularly in high-dimensional or complex multi-agent reinforcement learning. Formally, the KL divergence between the old and new policies is defined as:

$$D_{\text{KL}}(\pi_{\text{old}}||\pi_{\text{new}}) = \sum_a \pi_{\text{old}} \cdot \log \frac{\pi_{\text{old}}}{\pi_{\text{new}}} = \mathbb{E}_{\pi_{\text{old}}} \left( \log \frac{\pi_{\text{old}}}{\pi_{\text{new}}} \right). \quad (10)$$

Although Eqn. (10) offers a rigorous definition, directly constraining  $\log \frac{\pi_{\text{old}}}{\pi_{\text{new}}}$  can lead to incorrect behavior, as it does not provide a proper unbiased estimator and may even yield negative values, violating the non-negativity of KL divergence. To address this, we first establish the following identity, which forms the foundation for an unbiased KL estimator:

$$\begin{aligned} & E_{\pi_{\text{old}}} \left( -\log \frac{\pi_{\text{new}}}{\pi_{\text{old}}} + \frac{\pi_{\text{new}}}{\pi_{\text{old}}} - 1 \right) \\ &= E_{\pi_{\text{old}}} \left( \log \frac{\pi_{\text{old}}}{\pi_{\text{new}}} + \frac{\pi_{\text{new}}}{\pi_{\text{old}}} - 1 \right) \\ &= D_{\text{KL}}(\pi_{\text{old}}||\pi_{\text{new}}) + \sum_a \pi_{\text{new}} - \sum_a \pi_{\text{old}} \\ &= D_{\text{KL}}(\pi_{\text{old}}||\pi_{\text{new}}). \end{aligned} \quad (11)$$

where the last equality follows from the fact that both policies are normalized distributions, i.e.  $\sum_a \pi_{\text{new}} = \sum_a \pi_{\text{old}} = 1$ .

This motivates the definition of  $f(x) = x - 1 - \log x$ , where  $x = \frac{\pi_{\text{new}}}{\pi_{\text{old}}}$ . Crucially,  $f(x)$  satisfies  $E_{\pi_{\text{old}}}[f(x)] = D_{\text{KL}}(\pi_{\text{old}}||\pi_{\text{new}})$ , thereby providing an unbiased and theoretically sound surrogate for KL divergence that is inherently non-negative. This property allows  $f(x)$  to serve as a robust alternative to the heuristic TV-based clipping mechanism in PPO and forms the basis of our enhanced clipping strategy.

**Theorem 1.** *We define the unbiased estimation function  $f(x) = x - 1 - \log(x)$ , derived from Eqn. (11), where  $x = \frac{\pi_{\text{new}}}{\pi_{\text{old}}}$ . This function satisfies:*

1.  $f(x) \geq 0$  for all  $x > 0$  with equality if and only if  $x = 1$ .




Figure 1: Root solving for dynamic clipping bounds based on KL targets.

2.  $f(x)$  is convex in  $x$ , ensuring stable optimization when used for clipping.

*Proof.* (1) **Non-negativity:** We first compute the derivative of  $f(x)$ :

$$f'(x) = 1 - \frac{1}{x}.$$

For  $x > 1$ ,  $f'(x) > 0$ , and for  $0 < x < 1$ ,  $f'(x) < 0$ . Therefore,  $f(x)$  decreases on  $(0, 1]$  and increases on  $[1, \infty)$ , attaining its global minimum at  $x = 1$ . Evaluating  $f(1)$ :  $f(1) = 1 - 1 - \log 1 = 0$ . Thus,  $f(x) \geq 0$  for all  $x > 0$ , with equality only at  $x = 1$ .

- (2) **Convexity:** We compute the second derivative:

$$f''(x) = \frac{1}{x^2} > 0, \quad \forall x > 0,$$

which implies that  $f(x)$  is convex. Convexity ensures that local clipping constraints defined by  $f(x)$  form a valid and stable bound in optimization.  $\square$

**Advantages of the unbiased estimator.** The proposed function  $f(x)$  exhibits several advantages over the standard TV-based clipping mechanism in PPO:

- **Non-negativity and convexity:**  $f(x)$  avoids invalid negative divergence values and provides a convex structure for stable optimization.
- **Compatibility with KL-based policy optimization:** By aligning clipping directly with KL divergence, the method captures both the magnitude and directional information of policy updates, improving optimization in high-dimensional reinforcement learning settings.

**Construction of Clipping Bounds.** To enhance the stability and adaptability of reflective policy optimization, we propose a novel clipping mechanism that dynamically determines asymmetric clipping bounds based on the inverse of a function  $f(x)$  related to the true KL divergence. Specifically, the true KL divergence between the current and the old policy is computed directly, and the function  $f(x)$  is used to find the inverse, yielding two roots that define the clipping interval. This target KL value is subsequently smoothed using an exponential moving average (EMA), forming the dynamic KL constraint.

$$D_{\text{KL}_t}^{\text{target}} = \beta \cdot D_{\text{KL}_{t-1}}^{\text{target}} + (1 - \beta) \cdot D_{\text{KL}_t}. \quad (12)$$

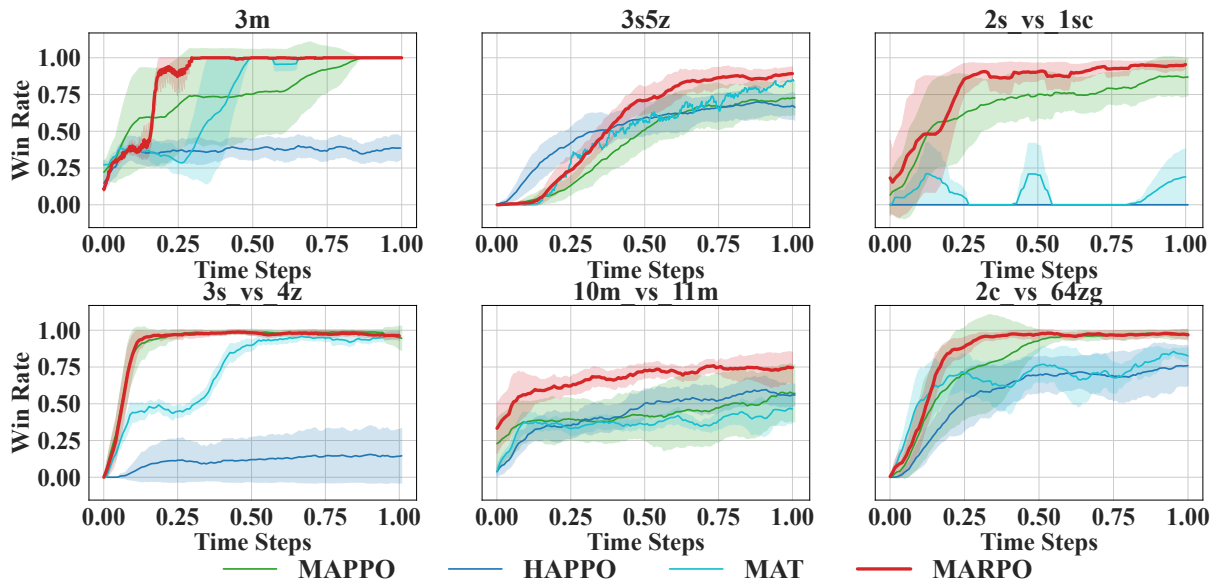


Figure 2: Win rate curves for the main experiments on SMAC-Hard environments. Shaded regions represent the standard deviation across five random seeds. The X-axis denotes timesteps, and the Y-axis denotes win rate. This experiment primarily compares policy-based methods with baselines.




Figure 3: Win rate curves from the ablation study on SMAC-Hard environments, validating the impact of different modules. Shaded regions represent the standard deviation across five random seeds.

In contrast to fixed or heuristically scheduled KL bounds, our formulation induces a self-adjusting trust region that evolves automatically with the policy. The clipping thresholds are obtained by numerically solving for the roots of a KL-based approximation function, as introduced in Eqn. (12), yielding an asymmetric range  $[x_1, x_2]$  that adapts continuously to the current policy and observed KL statistics. In our approach, these dynamic thresholds are derived analytically from the target KL divergence and updated online, which is further smoothed via an EMA mechanism to ensure temporal consistency and reduce short-term variance. As a result, the method eliminates the need for manually chosen KL bounds or hand-crafted annealing schedules. It offers a principled, stable, and adaptive clipping mechanism—distinct from PPO’s symmetric and fixed-range clipping (see Fig. 1)—while maintaining robust performance across diverse environments and training regimes.

## 5 Experiments

### 5.1 Experiment Setup and Evaluations

**StarCraft II Multi-Agent Challenge.** SMAC is a widely used benchmark for cooperative MARL, featuring diverse combat scenarios based on the StarCraft II engine. Agents must collaborate to defeat built-in AI opponents. However, the original SMAC’s limited opponent diversity often leads to overfitting and poor generalization. To address this, SMAC-Hard introduces mixed scripted opponents, randomized strategy switching, and a self-play interface, enhancing adversarial variability and robustness evaluation. With its open-source implementation publicly available, SMAC-Hard establishes a new benchmark for evaluating the robustness, adaptability, and strategic coverage of MARL algorithms in dynamic, partially observable environments. We conduct experiments on both SMAC and SMAC-Hard, cov-

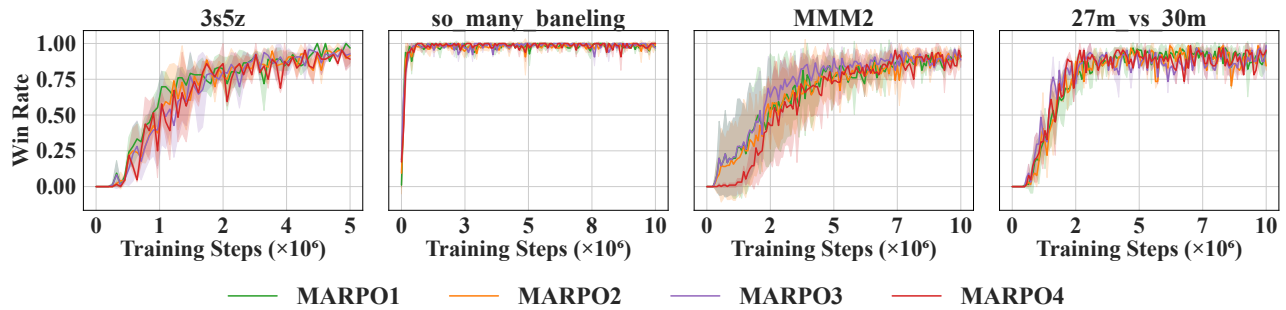


Figure 4: Hyperparameter Sensitivity Study: Impact of KL Bias and Sliding Average Update Rate ( $\beta$ ) on Win Rate in SMAC Environments. Shaded regions represent the standard deviation across five random seeds.




Figure 5: Win rate curves for the main experiments on GRF environments. Shaded regions represent the standard deviation across three random seeds. The X-axis denotes timesteps, and the Y-axis denotes win rate.

ering scenarios of varying difficulty and asymmetry to assess the adaptability and generalization of our method.

**Main Experiment Setup.** Experiments are conducted on a subset of scenarios from the SMAC-Hard benchmark, selected to reflect a broad range of task difficulty and coordination complexity. We compare our method against several strong baselines on SMAC-Hard, including value-based methods (QMIX (Rashid et al. 2020), LDSA (Yang et al. 2022), QPLEX (Wang et al. 2021)) and policy-based methods (MAPPO (Yu et al. 2022), HAPPO (Kuba et al. 2021), and MAT (Wen et al. 2022)), providing a comprehensive evaluation across diverse algorithmic paradigms. All agents are trained for 10 million environment steps. For evaluation, we report the average win rate over the final 2 million training steps. To ensure fair comparisons, all methods adopt identical neural network architectures and optimization settings. Additionally, we conduct supplementary experiments on the original SMAC benchmark to further evaluate the generalization ability of our approach.

## 5.2 Main Results on SMAC-Hard

Our method consistently outperforms existing baselines across diverse SMAC-Hard scenarios, achieving higher win rates and improved robustness (Table 1). This gain stems from better utilization of low-probability, high-impact samples during updates, mitigating issues in clipped objectives. As shown in Figure 2, our approach exhibits stable convergence and enhanced sample efficiency in challenging multi-agent tasks.

## 5.3 Analysis of Algorithm Performance

**Ablation Experiments.** We conduct ablation experiments to evaluate the contribution of each module within our framework. Additionally, we analyze the computational overhead to assess the efficiency and scalability of the proposed algorithm. As shown in Figure 3, we conducted ablation experiments in most environments by removing each module individually. The results demonstrate that removing any module results in performance degradation, indicating that each component contributes to the overall effectiveness of the algorithm. This highlights that all modules play essential roles across different scenarios, and their combined use is crucial for achieving improved performance.

**Sensitivity Analysis of Hyperparameters.** In our method, the KL divergence is computed as the true KL between the new and old policy distributions and averaged over trajectories. Instead of imposing an upper bound, we maintain a lower bound to prevent the KL from collapsing to zero. This value is tracked using an Exponential Moving Average (EMA) with two hyperparameters: the KL bias and the update rate  $\beta$  defined in Eqn. (12). To examine the algorithm’s sensitivity to these hyperparameters, we test several representative settings and analyze their effects on stability, convergence, and sample efficiency.

Figure 4 summarizes the results across four configurations—MARPO1 (0.05, 0.05), MARPO2 (0.08, 0.08), MARPO3 (0.10, 0.08), and MARPO4 (0.10, 0.01)—where each pair denotes the KL bias and EMA update rate used in the *second*, KL-guided clipping term, while the




Figure 6: Win rate curves for the main experiments on SMACv2 environments. Shaded regions represent the standard deviation across three random seeds. The X-axis denotes timesteps, and the Y-axis denotes win rate.

Env Name	MAPPO	HAPPO	LDSA	QMIX	QPLEX	MAT	MARPO (Ours)
3m	99.2 ± 4.7	37.3 ± 8.7	99.6 ± 1.5	99.8 ± 0.8	5.7 ± 19.3	99.9 ± 0.6	<b>100.0 ± 0.0</b> (✓)
3s5z	71.0 ± 1.7	68.1 ± 1.1(✓)	12.3 ± 0.7	34.9 ± 0.8	26.4 ± 3.0	79.7 ± 0.5	<b>87.2 ± 0.4</b>
2s_vs_1sc	85.2 ± 13.5	0.0 ± 0.0	93.9 ± 9.5	61.4 ± 36.7	81.5 ± 15.5	10.3 ± 15.0	<b>94.8 ± 4.2</b> (✓)
3s_vs_4z	<b>97.8 ± 7.5</b>	14.4 ± 18.6	95.0 ± 6.5	75.8 ± 20.2	78.8 ± 22.8	94.4 ± 6.3	97.0 ± 3.2 (✓)
10m_vs_11m	53.0 ± 4.7	57.6 ± 1.0	45.7 ± 2.6	65.2 ± 1.1	0.4 ± 0.0	43.2 ± 1.7	<b>74.3 ± 1.2</b> (✓)
2c_vs_64zg	97.3 ± 3.3	73.3 ± 16.9	87.0 ± 9.5	69.7 ± 32.3	35.1 ± 32.5	79.7 ± 12.3	<b>97.4 ± 3.3</b> (✓)

Table 1: Performance comparison across 6 SMAC-Hard environments. The values represent the average performance over the last 2 million steps. A check mark (✓) in parentheses indicates the highest value for the first 2.5 million steps. The values are reported as mean ± standard deviation, with the variance represented as  $10^{-2}$ .

standard PPO-style ratio clipping threshold is kept fixed across all runs. The curves are almost indistinguishable across these configurations, indicating that once the KL-based clipping is enabled, MARPO is fairly insensitive to the precise choice of these two hyperparameters in our tested range. This suggests that MARPO does not rely on careful tuning of the KL-related clipping term for stable performance, and a more exhaustive joint exploration of both clipping mechanisms is left for future work.

#### 5.4 Experiments on Google Football and SMACv2

We evaluate MARPO on two cooperative multi-agent benchmarks: Google Research Football (GRF) and SMACv2. GRF features fast-paced, partially observable gameplay that stresses real-time coordination among agents, while SMACv2 introduces stochastic unit behaviors, delayed rewards, and non-stationary opponents, making it substantially harder than the original SMAC and a good test of robustness.

On GRF, we focus on generalization and sample efficiency. As shown in Figure 5, MARPO attains a clear early performance lead over MAPPO across multiple scenarios, indicating that the reflective mechanism helps extract useful trajectory information more quickly and accelerates policy adaptation. Throughout training, MARPO maintains stable performance that remains competitive with MAPPO, suggesting that the early gains do not come at the cost of instability. On SMACv2, we assess robustness under stochastic and non-stationary dynamics. Figure 6 shows that MARPO yields consistently higher win rates and more stable learn-

ing than MAPPO on a range of maps, and often converges faster despite the increased task difficulty and noisy credit assignment.

Taken together, the GRF and SMACv2 results show that MARPO improves early-phase sample efficiency while preserving strong robustness in challenging multi-agent environments, supporting its effectiveness as a general-purpose on-policy method for complex cooperative control in realistic, large-scale, and highly non-stationary settings with diverse tasks, agents, difficulty levels, and evaluation scenarios, where its reflective mechanism and adaptive clipping strategy enable stable learning dynamics and reliable coordination across heterogeneous benchmarks.

## 6 Conclusion

We propose Multi-Agent Reflective Policy Optimization (MARPO), a framework that enhances sample efficiency through a reflection mechanism and stabilizes training via a KL-guided asymmetric clipping strategy. Experiments on StarCraft II and Google Research Football demonstrate that MARPO achieves competitive performance across diverse settings and yields clear gains in highly stochastic, non-stationary scenarios where optimization is more challenging. On relatively simple tasks, where MAPPO is already strong, the improvements from reflection are understandably less pronounced, but MARPO preserves comparable performance while offering stable learning dynamics. Future work includes exploring a broader range of hyperparameters, further refining the reflection mechanism, and extending MARPO to complex, large-scale tasks to fully exploit its advantages in challenging multi-agent environments.

## Acknowledgments

The work is supported in part by National Key R&D Program of China under Grant No. 2024YFB4505500 & 2024YFB4505503 and in part by the National Natural Science Foundation of China (62506191, 62222606, 62331006, and 62171038), and the Fundamental Research Funds for the Central Universities.

## References

- Canese, L.; Cardarilli, G. C.; Di Nunzio, L.; Fazzolari, R.; Giardino, D.; Re, M.; and Spanò, S. 2021. Multi-agent reinforcement learning: A review of challenges and applications. *Applied Sciences*, 11(11): 4948.
- De Witt, C. S.; Gupta, T.; Makoviichuk, D.; Makoviychuk, V.; Torr, P. H.; Sun, M.; and Whiteson, S. 2020. Is independent learning all you need in the starcraft multi-agent challenge? *arXiv preprint arXiv:2011.09533*.
- Deng, Y.; Wang, Z.; and Zhang, Y. 2024. Improving multi-agent reinforcement learning with stable prefix policy. In *Proceedings of the International Joint Conference on Artificial Intelligence*, 49–57.
- Deng, Y.; Yu, Y.; Ma, W.; Wang, Z.; Zhu, W.; Zhao, J.; and Zhang, Y. 2024. SMAC-Hard: Enabling Mixed Opponent Strategy Script and Self-play on SMAC. *arXiv preprint arXiv:2412.17707*.
- Ellis, B.; Cook, J.; Moalla, S.; Samvelyan, M.; Sun, M.; Mahajan, A.; Foerster, J.; and Whiteson, S. 2023. Smacv2: An improved benchmark for cooperative multi-agent reinforcement learning. In *Advances in Neural Information Processing Systems*, 37567–37593.
- Gan, Y.; Yan, R.; Wu, Z.; and Xing, J. 2024. Reflective Policy Optimization. In *Proceedings of the International Conference on Machine Learning*, 14471–14490.
- Haarnoja, T.; Zhou, A.; Abbeel, P.; and Levine, S. 2018. Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. In *Proceedings of the International Conference on Machine Learning*, 1861–1870.
- He, C.; Zou, B.; Xing, J.; Chen, J.; Shi, Y.; and Ma, H. 2025. DeCoDe: Defer-and-Complement Decision-Making via Decoupled Concept Bottleneck Models. *arXiv preprint arXiv:2505.19220*.
- Hernandez-Leal, P.; Kaisers, M.; Baarslag, T.; and De Cote, E. M. 2017. A survey of learning in multiagent environments: Dealing with non-stationarity. *arXiv preprint arXiv:1707.09183*.
- Jin, W.; Du, H.; Zhao, B.; Tian, X.; Shi, B.; and Yang, G. 2025. A comprehensive survey on multi-agent cooperative decision-making: Scenarios, approaches, challenges and perspectives. *arXiv preprint arXiv:2503.13415*.
- Kapturowski, S.; Ostrovski, G.; Dabney, W.; Quan, J.; and Munos, R. 2019. Recurrent Experience Replay in Distributed Reinforcement Learning. In *International Conference on Learning Representations*.
- Kuba, J. G.; Chen, R.; Wen, M.; Wen, Y.; Sun, F.; Wang, J.; and Yang, Y. 2021. Trust region policy optimisation in multi-agent reinforcement learning. In *International Conference on Learning Representations*.
- Kurach, K.; Raichuk, A.; Bachem, O.; Espeholt, L.; Riquelme, C.; Vincent, D.; Michalski, M.; Bousquet, O.; et al. 2020. Google research football: A novel reinforcement learning environment. In *Proceedings of the AAAI conference on artificial intelligence*, 4501–4510.
- Liu, C.; and Liu, G. 2024. JointPPO: diving deeper into the effectiveness of PPO in multi-agent reinforcement learning. *arXiv preprint arXiv:2404.11831*.
- Long, T.; Chen, P.; Xia, Y.; Ma, Y.; Sun, X.; Zhao, J.; and Lyu, Y. 2024. A Deep Deterministic Policy Gradient-Based Method for Enforcing Service Fault-Tolerance in MEC. *Chinese Journal of Electronics*, 33(4): 899–909.
- Massaoudi, M.; and Davis, K. R. 2025. Adaptive PPO With Multi-Armed Bandit Clipping and Meta-Control for Robust Power Grid Operation Under Adversarial Attacks. *IEEE Access*, 13: 73586–73602.
- Na, H.; Seo, Y.; and Moon, I.-c. 2024. Efficient episodic memory utilization of cooperative multi-agent reinforcement learning. *arXiv preprint arXiv:2403.01112*.
- Ng, A. Y.; Russell, S.; et al. 2000. Algorithms for inverse reinforcement learning. In *Proceedings of the International Conference on Machine Learning*, 663–670.
- Oliehoek, F. A.; and Amato, C. 2016. *A concise introduction to decentralized POMDPs*, volume 1. Springer.
- Prokopenko, M.; Lizier, J. T.; Obst, O.; and Wang, X. R. 2011. Relating Fisher information to order parameters. *Physical Review E—Statistical, Nonlinear, and Soft Matter Physics*, 84(4): 041116.
- Qin, G.; Meng, X.; Wen, T.; and Cai, B. 2024. Virtual Coupling Trains Based on Multi-Agent System Under Communication Delay. *Chinese Journal of Electronics*, 33(6): 1545–1554.
- 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. J.; Hung, C.-M.; Torr, P. H. S.; Foerster, J.; and Whiteson, S. 2019. The StarCraft Multi-Agent Challenge. *arXiv preprint arXiv:1902.04043*.
- Schaul, T.; Quan, J.; Antonoglou, I.; and Silver, D. 2015. Prioritized experience replay. *arXiv preprint arXiv:1511.05952*.
- Schulman, J.; Levine, S.; Moritz, P.; and Jordan, M. 2015. Trust Region Policy Optimization. In *Proceedings of the International Conference on Machine Learning*, 1889–1897.
- Schulman, J.; Moritz, P.; Levine, S.; Jordan, M.; and Abbeel, P. 2016. High-dimensional continuous control using generalized advantage estimation. In *International Conference on Learning Representations*.

Schulman, J.; Wolski, F.; Dhariwal, P.; Radford, A.; and Klimov, O. 2017. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*.

Silver, D.; Hubert, T.; Schrittwieser, J.; Antonoglou, I.; Lai, M.; Guez, A.; Lanctot, M.; Sifre, L.; Kumaran, D.; Graepel, T.; et al. 2017. Mastering chess and shogi by self-play with a general reinforcement learning algorithm. *arXiv preprint arXiv:1712.01815*.

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

Wang, Y.; He, H.; Tan, X.; and Gan, Y. 2019. Trust region-guided proximal policy optimization. In *Proceedings of the Advances in Neural Information Processing Systems*.

Wen, M.; Kuba, J. G.; Lin, R.; Zhang, W.; Wen, Y.; Wang, J.; and Yang, Y. 2022. Multi-Agent Reinforcement Learning is a Sequence Modeling Problem. In *Proceedings of the Advances in Neural Information Processing Systems*, 16509–16521.

Xiong, N.; Liu, Z.; Wang, Z.; and Yang, Z. 2023. Sample-efficient multi-agent rl: An optimization perspective. *arXiv preprint arXiv:2310.06243*.

Yang, M.; Zhao, J.; Hu, X.; Zhou, W.; Zhu, J.; and Li, H. 2022. LDSA: learning dynamic subtask assignment in cooperative multi-agent reinforcement learning. In *Proceedings of the Advances in Neural Information Processing Systems*, 1698–1710.

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 *Proceedings of the Advances in Neural Information Processing Systems*, 24611–24624.

Yue, S.; Deng, Y.; Wang, G.; Ren, J.; and Zhang, Y. 2024. Federated Offline Reinforcement Learning with Proximal Policy Evaluation. *Chinese Journal of Electronics*, 33(6): 1360–1372.

Zhang, T.; Fu, Y.; and Zhang, J. 2024. Deep Guided Attention Network for Joint Denoising and Demosaicing in Real Image. *Chinese Journal of Electronics*, 33(1): 303–312.

Zhang, Y.; Lai, Z.; Zhang, T.; Fu, Y.; and Zhou, C. 2025. Unaligned RGB Guided Hyperspectral Image Super-Resolution with Spatial-Spectral Concordance. *International Journal of Computer Vision*, 133(9): 6590–6610.

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 *Proceedings of Advances in Neural Information Processing Systems*, 3757–3769.