TRUST REGION POLICY OPTIMISATION IN MULTI-AGENT REINFORCEMENT LEARNING

Anonymous authors
Paper under double-blind review

ABSTRACT

Trust region methods rigorously enabled reinforcement learning (RL) agents to learn monotonically improving policies, leading to superior performance on a variety of tasks. Unfortunately, when it comes to multi-agent reinforcement learning (MARL), the property of monotonic improvement may not simply apply; this is because agents, even in cooperative games, could have conflicting directions of policy updates. As a result, achieving a guaranteed improvement on the joint policy where each agent acts individually remains an open challenge. In this paper, we extend the theory of trust region learning to MARL. Central to our findings are the *multi-agent advantage decomposition lemma* and the *sequential policy update scheme*. Based on these, we develop *Heterogeneous-Agent Trust Region Policy Optimisation* (HATPRO) and *Heterogeneous-Agent Proximal Policy Optimisation* (HAPPO) algorithms. Unlike many existing MARL algorithms, HATRPO/HAPPO do not need agents to share parameters, nor do they need any restrictive assumptions on decomposibility of the joint value function. Most importantly, we justify in theory the monotonic improvement property of HATRPO/HAPPO. We evaluate the proposed methods on a series of Multi-Agent MuJoCo and StarCraftII tasks. Results show that HATRPO and HAPPO significantly outperform strong baselines such as IPPO, MAPPO and MADDPG on all tested tasks, therefore establishing a new state of the art.

1 INTRODUCTION

Policy gradient (PG) methods have played a major role in recent developments of reinforcement learning (RL) algorithms (Silver et al., 2014; Schulman et al., 2015a; Haarnoja et al., 2018). Among the many PG variants, *trust region learning* (Kakade & Langford, 2002), with two typical embodiments of *Trust Region Policy Optimization (TRPO)* (Schulman et al., 2015a) and *Proximal Policy Optimization (PPO)* (Schulman et al., 2017) algorithms, offer supreme empirical performance in both discrete and continuous RL problems (Duan et al., 2016). The effectiveness of trust region methods largely stems from their theoretically-justified policy iteration procedure. By optimising the policy within a *trustable* neighbourhood of the current policy, thus avoiding making aggressive updates towards risky directions, trust region learning enjoys the guarantee of monotonic performance improvement at every iteration.

In multi-agent reinforcement learning (MARL) settings (Yang & Wang, 2020), naively applying PG methods by considering other agents as a part of the environment can lose its effectiveness. This is intuitively clear: once a learning agent updates its policy, so do its opponents; this however changes the loss landscape of the learning agent, thus harming the improvement effect from the PG update. As a result, applying independent PG updates in MARL offers poor convergence property (Claus & Boutilier, 1998). To address this, a learning paradigm named *centralised training with decentralised execution (CTDE)* (Lowe et al., 2017b; Foerster et al., 2018; Zhou et al., 2021) was developed. In CTDE, each agent is equipped with a joint value function which, during training, has access to the global state and opponents' actions. With the help of the centralised value function that accounts for the non-stationarity caused by others, each agent adapts its policy parameters accordingly. As such, the CTDE paradigm allows a straightforward extension of single-agent PG theorems (Sutton et al., 2000; Silver et al., 2014) to multi-agent scenarios (Lowe et al., 2017b; Kuba et al., 2021; Mguni et al., 2021). Consequently, fruitful PG algorithms have been developed for MARL (Foerster et al., 2018; Peng et al., 2017; Zhang et al., 2020; Wen et al., 2018; 2020; Yang et al., 2018).

Unfortunately, existing CTDE methods offer no solution of how to perform trust region learning in MARL. Lack of such an extension impedes agents from learning monotonically improving policies in a stable manner. Recent attempts such as IPPO (de Witt et al., 2020a) and MAPPO (Yu et al., 2021) have been proposed to fill such a gap; however, these methods are designed for agents that are *homogeneous* (i.e., sharing the same action space and policy parameters), which largely limits their applicability and harm the performance. As we show in Proposition 1 later, parameter sharing could suffer from an exponentially-worse suboptimal outcome. On the other hand, although IPPO/MAPPO can be practically applied in a non-parameter sharing way, it still lacks the essential theoretical property of trust region learning, which is the monotonic improvement guarantee.

In this paper, we propose the first theoretically-justified trust region learning framework in MARL. The key to our findings are the *multi-agent advantage decomposition lemma* and the *sequential policy update scheme*. With the advantage decomposition lemma, we introduce a multi-agent policy iteration procedure that enjoys the monotonic improvement guarantee. To implement such a procedure, we propose two practical algorithms: *Heterogeneous-Agent Trust Region Policy Optimisation (HATRPO)* and *Heterogeneous-Agent Proximal Policy Optimisation (HAPPO)*. HATRPO/HAPPO adopts the sequential update scheme, which saves the cost of maintaining a centralised critic for each agent in CTDE. Importantly, HATRPO/HAPPO does not require homogeneity of agents, nor any other restrictive assumptions on the decomposibility of the joint Q-function (Rashid et al., 2018). We evaluate HATRPO and HAPPO on benchmarks of StarCraftII and Multi-Agent MuJoCo against strong baselines such as MADDPG (Lowe et al., 2017a), IPPO (de Witt et al., 2020b) and MAPPO (Yu et al., 2021); results clearly demonstrate its state-of-the-art performance across all tested tasks.

2 Preliminaries

In this section, we first introduce problem formulation and notations for MARL, and then briefly review trust region learning in RL and discuss the difficulty of extending it to MARL. We end by surveying existing MARL work that relates to trust region methods and show their limitations.

2.1 MARL Problem Formulation and Notations

We consider a Markov game (Littman, 1994), which is defined by a tuple $\langle \mathcal{N}, \mathcal{S}, \boldsymbol{\mathcal{A}}, P, r, \gamma \rangle$. Here, $\mathcal{N} = \{1, \dots, n\}$ denotes the set of agents, $\mathcal{S}$ is the finite state space, $\boldsymbol{\mathcal{A}} = \prod_{i=1}^{n} \mathcal{A}^i$ is the product of finite action spaces of all agents, known as the joint action space, $P : \mathcal{S} \times \boldsymbol{\mathcal{A}} \times \mathcal{S} \to [0, 1]$ is the transition probability function, $r : \mathcal{S} \times \boldsymbol{\mathcal{A}} \to \mathbb{R}$ is the reward function, and $\gamma \in [0, 1)$ is the discount factor. The agents interact with the environment according to the following protocol: at time step $t \in \mathbb{N}$, the agents are at state $\mathrm{s}_t \in \mathcal{S}$; every agent i takes an action $\mathrm{a}_t^i \in \mathcal{A}^i$, drawn from its policy $\pi^i(\cdot|\mathrm{s}_t)$, which together with other agents' actions gives a joint action $\mathbf{a}_t = (\mathrm{a}_t^1, \dots, \mathrm{a}_t^n) \in \boldsymbol{\mathcal{A}}$, drawn from the joint policy $\boldsymbol{\pi}(\cdot|\mathrm{s}_t) = \prod_{i=1}^{n} \pi^i(\cdot^i|\mathrm{s}_t)$; the agents receive a joint reward $\mathrm{r}_t = r(\mathrm{s}_t, \mathbf{a}_t) \in \mathbb{R}$, and move to a state s_{t+1} with probability $P(\mathrm{s}_{t+1}|\mathrm{s}_t, \mathbf{a}_t)$. The joint policy $\boldsymbol{\pi}$, the transition probabililty function P, and the initial state distribution ρ^0, induce a marginal state distribution at time t, denoted by $\rho_{\boldsymbol{\pi}}^t$. We define an (improper) marginal state distribution $\rho_{\boldsymbol{\pi}} \triangleq \sum_{t=0}^{\infty} \gamma^t \rho_{\boldsymbol{\pi}}^t$. The state value function and the state-action value function are defined as $V_{\boldsymbol{\pi}}(s) \triangleq \mathbb{E}_{\mathbf{a}_{0:\infty} \sim \boldsymbol{\pi}, \mathrm{s}_{1:\infty} \sim P} \left[\sum_{t=0}^{\infty} \gamma^t \mathrm{r}_t \middle| \mathrm{s}_0 = s \right]$ and $Q_{\boldsymbol{\pi}}(s, \boldsymbol{a}) \triangleq \mathbb{E}_{\mathrm{s}_{1:\infty} \sim P, \mathbf{a}_{1:\infty} \sim \boldsymbol{\pi}} \left[\sum_{t=0}^{\infty} \gamma^t \mathrm{r}_t \middle| \mathrm{s}_0 = s,\ \mathbf{a}_0 = \boldsymbol{a} \right]$. The advantage function is written as $A_{\boldsymbol{\pi}}(s, \boldsymbol{a}) \triangleq Q_{\boldsymbol{\pi}}(s, \boldsymbol{a}) - V_{\boldsymbol{\pi}}(s)$. In this paper, we consider a *fully-cooperative* setting where all agents share the same reward function, aiming to maximise the expected total reward:

$$\mathcal{J}(\boldsymbol{\pi}) \triangleq \mathbb{E}_{\mathrm{s}_{0:\infty} \sim \rho_{\boldsymbol{\pi}}^{0:\infty}, \mathbf{a}_{0:\infty} \sim \boldsymbol{\pi}} \left[\sum_{t=0}^{\infty} \gamma^t \mathrm{r}_t \right].$$

Throughout this paper, we pay close attention to the contribution to performance from different subsets of agents. Before proceeding to our methods, we introduce following novel definitions.

Definition 1. *Let $i_{1:m}$ denote an ordered subset $\{i_1, \dots, i_m\}$ of $\mathcal{N}$, and let $-i_{1:m}$ refer to its complement. We write i_k when we refer to the kth agent in the ordered subset. Correspondingly, the multi-agent state-action value function is defined as*

$$Q_{\boldsymbol{\pi}}^{i_{1:m}}\left(s, \boldsymbol{a}^{i_{1:m}}\right) \triangleq \mathbb{E}_{\mathbf{a}^{-i_{1:m}} \sim \boldsymbol{\pi}^{-i_{1:m}}} \left[Q_{\boldsymbol{\pi}}\left(s, \boldsymbol{a}^{i_{1:m}}, \mathbf{a}^{-i_{1:m}}\right) \right],$$

and for disjoint sets $j_{1:k}$ and $i_{1:m}$, the multi-agent advantage function is

$$A_{\boldsymbol{\pi}}^{i_{1:m}}\left(s, \boldsymbol{a}^{j_{1:k}}, \boldsymbol{a}^{i_{1:m}}\right) \triangleq Q_{\boldsymbol{\pi}}^{j_{1:k}, i_{1:m}}\left(s, \boldsymbol{a}^{j_{1:k}}, \boldsymbol{a}^{i_{1:m}}\right) - Q_{\boldsymbol{\pi}}^{j_{1:k}}\left(s, \boldsymbol{a}^{j_{1:k}}\right). \tag{1}$$

Hereafter, the joint policies $\boldsymbol{\pi} = (\pi^1, \ldots, \pi^n)$ and $\bar{\boldsymbol{\pi}} = (\bar{\pi}^1, \ldots, \bar{\pi}^n)$ shall be thought of as the "current", and the "new" joint policy that agents update towards, respectively.

2.2 Trust Region Algorithms in Reinforcement Learning

Trust region methods such as TRPO (Schulman et al., 2015a) were proposed in single-agent RL with an aim of achieving a monotonic improvement of $\mathcal{J}(\pi)$ at each iteration. Formally, it can be described by the following theorem.

Theorem 1. *(Schulman et al., 2015a, Theorem 1) Let π be the current policy and $\bar{\pi}$ be the next candidate policy. We define $L_\pi(\bar{\pi}) = \mathcal{J}(\pi) + \mathbb{E}_{s\sim\rho_\pi, a\sim\bar{\pi}}\left[A_\pi(s,a)\right], D_{KL}^{max}(\pi, \bar{\pi}) = \max_s D_{KL}\left(\pi(\cdot|s), \bar{\pi}(\cdot|s)\right)$. Then the inequality of*

$$\mathcal{J}(\bar{\pi}) \geq L_\pi(\bar{\pi}) - C D_{KL}^{max}(\pi, \bar{\pi}) \tag{2}$$

holds, where $C = \frac{4\gamma \max_{s,a} |A_\pi(s,a)|}{(1-\gamma)^2}$.

The above theorem states that as the distance between the current policy π and a candidate policy $\bar{\pi}$ decreases, the surrogate $L_\pi(\bar{\pi})$, which involves only the current policy's state distribution, becomes an increasingly accurate estimate of the actual performance metric $\mathcal{J}(\bar{\pi})$. Based on this theorem, an iterative trust region algorithm is derived; at iteration $k+1$, the agent updates its policy by

$$\pi_{k+1} = \arg\max_{\pi} \left(L_{\pi_k}(\pi) - C D_{\mathrm{KL}}^{\max}(\pi_k, \pi)\right).$$

Such an update guarantees a monotonic improvement of the policy, i.e., $\mathcal{J}(\pi_{k+1}) \geq \mathcal{J}(\pi_k)$. To implement this procedure in practical settings with parameterised policies π_θ, Schulman et al. (2015a) approximated the KL-penalty with a KL-constraint, which gave rise to the TRPO update of

$$\theta_{k+1} = \arg\max_{\theta} L_{\pi_{\theta_k}}(\pi_\theta), \quad \text{subject to } \mathbb{E}_{s\sim\rho_{\pi_{\theta_k}}}\left[D_{\mathrm{KL}}(\pi_{\theta_k}, \pi_\theta)\right] \leq \delta. \tag{3}$$

At each iteration $k+1$, TRPO constructs a KL-ball $\mathfrak{B}_\delta(\pi_{\theta_k})$ around the policy π_{θ_k}, and optimises $\pi_{\theta_{k+1}} \in \mathfrak{B}_\delta(\pi_{\theta_k})$ to maximise $L_{\pi_{\theta_k}}(\pi_\theta)$. By Theorem 1, we know that the surrogate objective $L_{\pi_{\theta_k}}(\pi_\theta)$ is close to the true reward $\mathcal{J}(\pi_\theta)$ within $\mathfrak{B}_\delta(\pi_{\theta_k})$; therefore, π_{θ_k} leads to improvement.

Furthermore, to save the cost on $\mathbb{E}_{s\sim\rho_{\pi_{\theta_k}}}\left[D_{\mathrm{KL}}(\pi_{\theta_k}, \pi_\theta)\right]$ when computing Equation (3), Schulman et al. (2017) proposed an approximation solution to TRPO that uses only first order derivatives, known as *PPO*. PPO optimises the policy parameter θ_{k+1} by maximising the *PPO-clip* objective of

$$L_{\pi_{\theta_k}}^{\mathrm{PPO}}(\pi_\theta) = \mathbb{E}_{s\sim\rho_{\pi_{\theta_k}}, a\sim\pi_{\theta_k}}\left[\min\left(\frac{\pi_\theta(a|s)}{\pi_{\theta_k}(a|s)} A_{\pi_{\theta_k}}(s,a), \mathrm{clip}\left(\frac{\pi_\theta(a|s)}{\pi_{\theta_k}(a|s)}, 1\pm\epsilon\right) A_{\pi_{\theta_k}}(s,a)\right)\right]. \tag{4}$$

The *clip* operator replaces the ratio $\frac{\pi_\theta(a|s)}{\pi_{\theta_k}(a|s)}$ with $1+\epsilon$ or $1-\epsilon$, depending on whether or not the ratio is beyond the threshold interval. This effectively enables PPO to control the size of policy updates.

2.3 Limitations of Existing Trust Region Methods in MARL

Extending trust region methods to MARL is highly non-trivial. One naive approach is to equip all agents with one shared set of parameters and use agents' aggregated trajectories to conduct policy optimisation at every iteration. This approach was adopted by MAPPO (Yu et al., 2021) in which the policy parameter θ is optimised by maximising the objective of

$$L_{\pi_{\theta_k}}^{\mathrm{MAPPO}}(\pi_\theta) = \sum_{i=1}^{n} \mathbb{E}_{s\sim\rho_{\boldsymbol{\pi}_{\theta_k}}, \mathbf{a}\sim\boldsymbol{\pi}_{\theta_k}}\left[\min\left(\frac{\pi_\theta(a^i|s)}{\pi_{\theta_k}(a^i|s)} A_{\boldsymbol{\pi}_{\theta_k}}(s,\mathbf{a}), \mathrm{clip}\left(\frac{\pi_\theta(a^i|s)}{\pi_{\theta_k}(a^i|s)}, 1\pm\epsilon\right) A_{\boldsymbol{\pi}_{\theta_k}}(s,\mathbf{a})\right)\right]. \tag{5}$$

Unfortunately, this simple approach has significant drawbacks. An obvious demerit is that parameter sharing requires that all agents have identical action spaces, *i.e.*, $\mathcal{A}^i = \mathcal{A}^j, \forall i, j \in \mathcal{N}$, which limits the class of MARL problems to solve. Importantly, enforcing parameter sharing is equivalent to putting a constraint $\theta^i = \theta^j, \forall i, j \in \mathcal{N}$ on the joint policy space. In principle, this can lead to a suboptimal solution. To elaborate, we demonstrate through an example in the following proposition.

Proposition 1. *Let's consider a fully-cooperative game with an even number of agents n, one state, and the joint action space $\{0,1\}^n$, where the reward is given by $r(\mathbf{0}^{n/2}, \mathbf{1}^{n/2}) = r(\mathbf{1}^{n/2}, \mathbf{0}^{n/2}) = 1$, and $r(\boldsymbol{a}^{1:n}) = 0$ for all other joint actions. Let $\mathcal{J}^*$ be the optimal joint reward, and $\mathcal{J}^*_{share}$ be the optimal joint reward under the shared policy constraint. Then*

$$\frac{\mathcal{J}^*_{share}}{\mathcal{J}^*} = \frac{2}{2^n}.$$

For proof see Appendix B. In the above example, we show that parameter sharing can lead to a suboptimal outcome that is exponentially-worse with the increasing number of agents.

Apart from parameter sharing, a more general approach to extend trust region methods in MARL is to endow all agents with their own parameters, and at each iteration $k+1$, agents construct trust regions of $\{\mathfrak{B}_\delta(\pi^i_{\theta^i_k})\}_{i\in\mathcal{N}}$, and optimise their objectives $\{L_{\boldsymbol{\pi}_{\boldsymbol{\theta}_k}}(\pi^i_{\theta^i}\boldsymbol{\pi}^{-i}_{\boldsymbol{\theta}^{-i}_k})\}_{i\in\mathcal{N}}$. Admittedly, this approach can still be supported by the current MAPPO implementation (Yu et al., 2021) if one turns off parameter sharing, thus distributing the summation in Equation (5) to all agents. However, such an approach cannot offer a rigorous guarantee of monotonic improvement during training. In fact, agents' local improvements in performance can jointly lead to a worse outcome. For example, in Figure 1, we design a single-state differential game where two agents draw their actions from Gaussian distributions with learnable means μ^1, μ^2 and unit variance, and the reward function is $r(a^1, a^2) = a^1 a^2$. The failure of MAPPO-style approach comes from the fact that, although the reward function increases in each of the agents' (one-dimensional) update directions, it decreases in the joint (two-dimensional) update direction.

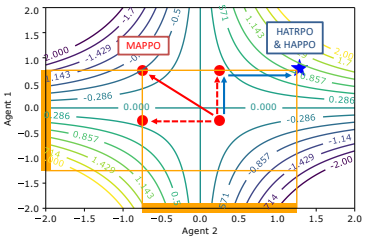

Figure 1: Example of differentiable game with $r(a^1, a^2) = a^1 a^2$. We initialise agents' 1-D Gaussian policies with $\mu^1 = -0.25$, $\mu^2 = 0.25$. Orange intervals represent the KL-ball of $\delta = 0.5$. Individual trust region update (red) jointly decreases the performance, whereas our sequential updates (blue) lead to improvement.

Having seen the limitations of existing trust region methods in MARL, in the following sections, we first introduce a multi-agent policy iteration procedure that enjoys theoretically-justified monotonic improvement guarantee. To implement this procedure, we propose *HATRPO* and *HAPPO* algorithms, which offer practical solutions to apply trust region learning in MARL without the necessity of assuming homogeneous agents while still maintaining the monotonic improvement property.

3 Multi-Agent Trust Region Learning

The purpose of this section is to develop a theoretically-justified trust region learning procedure in the context of multi-agent learning. In Subsection 3.1 we present the policy iteration procedure with monotonic improvement guarantee, and in Subsection 3.2 we analyse its properties during training and at convergence. Throughout the work, we make the following regularity assumptions.

Assumption 1. *There exists $\eta \in \mathbb{R}$, such that $0 < \eta \ll 1$, and for every agent $i \in \mathcal{N}$, the policy space Π^i is η-soft; that means that for every $\pi^i \in \Pi^i$, $s \in \mathcal{S}$, and $a^i \in \mathcal{A}^i$, we have $\pi^i(a^i|s) \geq \eta$.*

3.1 Trust Region Learning in MARL with Monotonic Improvement Guarantee

We start by introducing a pivotal lemma which shows that the joint advantage function can be decomposed into a summation of each agent's local advantages. Importantly, this lemma also offers a critical intuition behind the sequential policy-update scheme that our algorithms later apply.

Lemma 1 (Multi-Agent Advantage Decomposition)**.** *In any cooperative Markov games, given a joint policy $\boldsymbol{\pi}$, for any state s, and any agent subset $i_{1:m}$, the below equations holds.*

$$A^{i_{1:m}}_{\boldsymbol{\pi}}\left(s, \boldsymbol{a}^{i_{1:m}}\right) = \sum_{j=1}^{m} A^{i_j}_{\boldsymbol{\pi}}\left(s, \boldsymbol{a}^{i_{1:j-1}}, a^{i_j}\right).$$

For proof see Appendix C.2. Notably, Lemma 1 holds in general for cooperative Markov games, with no need for any assumptions on the decomposibility of the joint value function such as those in VDN (Sunehag et al., 2018), QMIX (Rashid et al., 2018) or Q-DPP (Yang et al., 2020).

Lemma 1 suggests a feasible way to find the direction of performance improvement (i.e., joint actions with positive advantage values) in multi-agent learning. Specifically, let agents take actions sequentially by following an arbitrary order $i_{1:n}$, assuming agent i_1 takes an action $\bar{a}^{i_1}$ such that $A^{i_1}(s, \bar{a}^{i_1}) > 0$, and for the rest $m = 2, \ldots, n$, the agent i_m takes an action $\bar{a}^{i_m}$ such that $A^{i_m}(s, \bar{\boldsymbol{a}}^{i_{1:m-1}}, \bar{a}^{i_m}) > 0$. For the joint action $\bar{\boldsymbol{a}}$, Lemma 1 assures that the advantage of $A_{\boldsymbol{\pi_\theta}}(s, \bar{\boldsymbol{a}})$ is positive, thus the performance will be improved. To extend the above process into a policy iteration procedure with monotonic improvement guarantee, we need the following definitions.

Definition 2. *Let $\boldsymbol{\pi}$ be a joint policy, $\bar{\boldsymbol{\pi}}^{i_{1:m-1}}$ be some other joint policy of agents $i_{1:m-1}$, and $\hat{\pi}^{i_m}$ be some other policy of agent i_m. Then*

$$L_{\boldsymbol{\pi}}^{i_{1:m}}\left(\bar{\boldsymbol{\pi}}^{i_{1:m-1}}, \hat{\pi}^{i_m}\right) \triangleq \mathbb{E}_{\mathrm{s}\sim\rho_{\boldsymbol{\pi}}, \mathbf{a}^{i_{1:m-1}}\sim\bar{\boldsymbol{\pi}}^{i_{1:m-1}}, \mathrm{a}^{i_m}\sim\hat{\pi}^{i_m}}\left[A_{\boldsymbol{\pi}}^{i_m}\left(\mathrm{s}, \mathbf{a}^{i_{1:m-1}}, \mathrm{a}^{i_m}\right)\right].$$

Note that, for any $\bar{\boldsymbol{\pi}}^{i_{1:m-1}}$, we have

$$\begin{aligned} L_{\boldsymbol{\pi}}^{i_{1:m}}\left(\bar{\boldsymbol{\pi}}^{i_{1:m-1}}, \pi^{i_m}\right) &= \mathbb{E}_{\mathrm{s}\sim\rho_{\boldsymbol{\pi}}, \mathbf{a}^{i_{1:m-1}}\sim\bar{\boldsymbol{\pi}}^{i_{1:m-1}}, \mathrm{a}^{i_m}\sim\pi^{i_m}}\left[A_{\boldsymbol{\pi}}^{i_m}\left(\mathrm{s}, \mathbf{a}^{i_{1:m-1}}, \mathrm{a}^{i_m}\right)\right] \\ &= \mathbb{E}_{\mathrm{s}\sim\rho_{\boldsymbol{\pi}}, \mathbf{a}^{i_{1:m-1}}\sim\bar{\boldsymbol{\pi}}^{i_{1:m-1}}}\left[\mathbb{E}_{\mathrm{a}^{i_m}\sim\pi^{i_m}}\left[A_{\boldsymbol{\pi}}^{i_m}\left(\mathrm{s}, \mathbf{a}^{i_{1:m-1}}, \mathrm{a}^{i_m}\right)\right]\right] = 0. \end{aligned} \tag{6}$$

Building on Lemma 1 and Definition 2, we can finally generalise Theorem 1 of TRPO to MARL.

Lemma 2. *Let $\boldsymbol{\pi}$ be a joint policy. Then, for any joint policy $\bar{\boldsymbol{\pi}}$, we have*

$$\mathcal{J}(\bar{\boldsymbol{\pi}}) \geq \mathcal{J}(\boldsymbol{\pi}) + \sum_{m=1}^{n}\left[L_{\boldsymbol{\pi}}^{i_{1:m}}\left(\bar{\boldsymbol{\pi}}^{i_{1:m-1}}, \bar{\pi}^{i_m}\right) - CD_{KL}^{max}(\pi^{i_m}, \bar{\pi}^{i_m})\right].$$

For proof see Appendix C.2. This lemma provides an idea about how a joint policy can be improved. Namely, by Equation (6), we know that if any agents were to set the values of the above summands $L_{\boldsymbol{\pi}}^{i_{1:m}}(\bar{\boldsymbol{\pi}}^{i_{1:m-1}}, \bar{\pi}^{i_m}) - CD_{\mathrm{KL}}^{\max}(\pi^{i_m}, \bar{\pi}^{i_m})$ by sequentially updating their policies, each of them can always make its summand be zero by making no policy update (i.e., $\bar{\pi}^{i_m} = \pi^{i_m}$). This implies that any positive update will lead to an increment in summation. Moreover, as there are n agents making policy updates, the compound increment can be large, leading to a substantial improvement. Lastly, note that this property holds regardless of the order in which the agents make the updates; this allows scheduling the order at every iteration. To summarise, we propose Algorithm 1.

Algorithm 1 Multi-Agent Policy Iteration with Monotonic Improvement Guarantee

1: Initialise the joint policy $\boldsymbol{\pi}_0 = (\pi_0^1, \ldots, \pi_0^n)$.
2: **for** $k = 0, 1, \ldots$ **do**
3: Compute the advantage function $A_{\boldsymbol{\pi}_k}(s, \boldsymbol{a})$ for all state-(joint)action pairs $(s, \boldsymbol{a})$.
4: Compute $\epsilon = \max_{s,\boldsymbol{a}} |A_{\boldsymbol{\pi}_k}(s, \boldsymbol{a})|$ and $C = \frac{4\gamma\epsilon}{(1-\gamma)^2}$.
5: Draw a permutaion $i_{1:n}$ of agents at random.
6: **for** $m = 1 : n$ **do**
7: Make an update $\pi_{k+1}^{i_m} = \arg\max_{\pi^{i_m}}\left[L_{\boldsymbol{\pi}_k}^{i_{1:m}}\left(\boldsymbol{\pi}_{k+1}^{i_{1:m-1}}, \pi^{i_m}\right) - CD_{\mathrm{KL}}^{\max}(\pi_k^{i_m}, \pi^{i_m})\right]$.
8: **end for**
9: **end for**

We want to highlight that Algorithm 1 is markedly different from naively applying the TRPO update in Equation (3) on the joint policy of all agents. Firstly, our Algorithm 1 does not update the entire joint policy at once, but rather update each agent's individual policy sequentially. Secondly, during the sequential update, each agent has a unique optimisation objective that takes into account all previous agents' updates, which is also the key for the monotonic improvement property to hold.

3.2 Theoretical Analysis

Now we justify by the following theorm that Algorithm 1 enjoys monotonic improvement property.

Theorem 2. *A sequence $(\boldsymbol{\pi}_k)_{k=0}^{\infty}$ of joint policies updated by Algorithm 1 has the monotonic improvement property, i.e., $\mathcal{J}(\boldsymbol{\pi}_{k+1}) \geq \mathcal{J}(\boldsymbol{\pi}_k)$ for all $k \in \mathbb{N}$.*

For proof see Appendix C.2. With the above theorem, we finally claim a successful introduction of trust region learning to MARL, as this generalises the monotonic improvement property of TRPO. Moreover, we take a step further to study the convergence property of Algorithm 1. Before stating the result, we introduce the following solution concept.

Definition 3. *In a fully-cooperative game, a joint policy $\boldsymbol{\pi}_* = (\pi^1_*, \dots, \pi^n_*)$ is a Nash equilibrium (NE) if for every $i \in \mathcal{N}$, $\pi^i \in \Pi^i$ implies $\mathcal{J}(\boldsymbol{\pi}_*) \geq \mathcal{J}(\pi^i, \boldsymbol{\pi}^{-i}_*)$.*

NE (Nash, 1951) is a well-established game-theoretic solution concept. Definition 3 characterises the equilibrium point at convergence for cooperative MARL tasks. Based on this, we have the following result that describes Algorithm 1's asymptotic convergent behaviour towards NE.

Proposition 2. *Supposing in Algorithm 1 any permutation of agents has a fixed non-zero probability, a sequence $(\boldsymbol{\pi}_k)_{k=0}^{\infty}$ of joint policies generated by Algorithm 1, in a one-shot Markov game, has a non-empty set of limit points, each of which is a Nash equilibrium.*

For proof see Appendix C.3. Note that here we only study the convergent behaviour of Algorithm 1 in one-shot games. Understanding the learning dynamics of trust region methods in multi-state Markov games is highly non-trivial; we leave it for future work.

4 Practical Algorithms

When implementing Algorithm 1 in practice, large state and action spaces could prevent agents from designating policies $\pi^i(\cdot|s)$ for each state s separately. To handle this, we parameterise each agent's policy $\pi^i_{\theta^i}$ by θ^i, which, together with other agents' policies, forms a joint policy $\boldsymbol{\pi}_{\boldsymbol{\theta}}$ parametrised by $\boldsymbol{\theta} = (\theta^1, \dots, \theta^n)$. In this section, we develop two deep MARL algorithms to optimise $\boldsymbol{\theta}$.

4.1 HATRPO

Computing $D^{\max}_{\mathrm{KL}}\left(\pi^{i_m}_{\theta^{i_m}_k}, \pi^{i_m}_{\theta^{i_m}}\right)$ in Algorithm 1 is challenging; it requires evaluating the KL-divergence for all states at each iteration. Similar to TRPO, one can ease this maximal KL-divergence penalty $D^{\max}_{\mathrm{KL}}\left(\pi^{i_m}_{\theta^{i_m}_k}, \pi^{i_m}_{\theta^{i_m}}\right)$ by replacing it with the expected KL-divergence constraint $\mathbb{E}_{\mathrm{s}\sim\rho_{\boldsymbol{\pi}_{\boldsymbol{\theta}_k}}}\left[D_{\mathrm{KL}}\left(\pi^{i_m}_{\theta^{i_m}_k}(\cdot|\mathrm{s}), \pi^{i_m}_{\theta^{i_m}}(\cdot|\mathrm{s})\right)\right] \leq \delta$ where δ is a threshold hyperparameter, and the expectation can be easily approximated by stochastic sampling. With the above amendment, we propose practical HATRPO algorithm in which, at every iteration $k+1$, given a permutation of agents $i_{1:n}$, agent $i_{m\in\{1,\dots,n\}}$ sequentially optimises its policy parameter $\theta^{i_m}_{k+1}$ by maximising a constrained objective:

$$\begin{aligned}\theta^{i_m}_{k+1} = \arg\max_{\theta^{i_m}} \mathbb{E}_{\mathrm{s}\sim\rho_{\boldsymbol{\pi}_{\boldsymbol{\theta}_k}}, \mathbf{a}^{i_{1:m-1}}\sim\boldsymbol{\pi}^{i_{1:m-1}}_{\boldsymbol{\theta}_{k+1}}, \mathrm{a}^{i_m}\sim\pi^{i_m}_{\theta}}\left[A^{i_m}_{\boldsymbol{\pi}_{\boldsymbol{\theta}_k}}(\mathrm{s}, \mathbf{a}^{i_{1:m-1}}, \mathrm{a}^{i_m})\right],\\ \text{subject to } \mathbb{E}_{\mathrm{s}\sim\rho_{\boldsymbol{\pi}_{\boldsymbol{\theta}_k}}}\left[D_{\mathrm{KL}}\left(\pi^{i_m}_{\theta^{i_m}_k}(\cdot|\mathrm{s}), \pi^{i_m}_{\theta}(\cdot|\mathrm{s})\right)\right] \leq \delta.\end{aligned} \tag{7}$$

To compute the above equation, one can apply a linear approximation to the objective function and a quadratic approximation to the KL constraint; the optimisation problem in Equation (7) can be solved by a closed-form update rule as

$$\theta^{i_m}_{k+1} = \theta^{i_m}_k + \alpha^j \sqrt{\frac{2\delta}{\boldsymbol{g}^{i_m}_k (\boldsymbol{H}^{i_m}_k)^{-1} \boldsymbol{g}^{i_m}_k}} (\boldsymbol{H}^{i_m}_k)^{-1} \boldsymbol{g}^{i_m}_k\,, \tag{8}$$

where $\boldsymbol{H}^{i_m}_k = \nabla^2_{\theta^{i_m}} \mathbb{E}_{\mathrm{s}\sim\rho_{\boldsymbol{\pi}_{\boldsymbol{\theta}_k}}}\left[D_{\mathrm{KL}}(\pi^{i_m}_{\theta^{i_m}_k}(\cdot|\mathrm{s}), \pi^{i_m}_{\theta^{i_m}}(\cdot|\mathrm{s}))\right]\big|_{\theta^{i_m}=\theta^{i_m}_k}$ is the Hessian of the expected KL-divergence, $\boldsymbol{g}^{i_m}_k$ is the gradient of the objective in Equation (7), $\alpha^j < 1$ is a positive coefficient that is found via backtracking line search, and the product of $(\boldsymbol{H}^{i_m}_k)^{-1}\boldsymbol{g}^{i_m}_k$ can be efficiently computed with conjugate gradient algorithm. The last missing piece for HATRPO is how to estimate $\mathbb{E}_{\mathbf{a}^{i_{1:m-1}}\sim\boldsymbol{\pi}^{i_{1:m-1}}_{\boldsymbol{\theta}_{k+1}}, \mathrm{a}^{i_m}\sim\pi^{i_m}_{\theta}}\left[A^{i_m}_{\boldsymbol{\pi}_{\boldsymbol{\theta}_k}}(\mathrm{s}, \mathbf{a}^{i_{1:m-1}}, \mathrm{a}^{i_m})\right]$, which poses new challenges because each agent's objective has to take into account all previous agents' updates, and the size of input vaires.

Fortunately, with the following proposition, we can efficiently estimate this objective by employing a joint advantage estimator.

Proposition 3. *Let $\boldsymbol{\pi}$ be a joint policy, and $A_{\boldsymbol{\pi}}(\mathrm{s}, \mathbf{a})$ be its joint advantage function. Let $\bar{\boldsymbol{\pi}}^{i_{1:m-1}}$ be a joint policy of agents $i_{1:m-1}$, and $\hat{\pi}^{i_m}$ be a policy of agent i_m. Then, for every state s,*

$$\mathbb{E}_{\mathbf{a}^{i_{1:m-1}} \sim \bar{\boldsymbol{\pi}}^{i_{1:m-1}}, \mathrm{a}^{i_m} \sim \hat{\pi}^{i_m}} \left[A_{\boldsymbol{\pi}}^{i_m}\left(s, \mathbf{a}^{i_{1:m-1}}, \mathrm{a}^{i_m}\right) \right] = \mathbb{E}_{\mathbf{a} \sim \boldsymbol{\pi}} \left[\left(\frac{\hat{\pi}^{i_m}(\mathrm{a}^{i_m}|s)}{\pi^{i_m}(\mathrm{a}^{i_m}|s)} - 1 \right) \frac{\bar{\boldsymbol{\pi}}^{i_{1:m-1}}(\mathbf{a}^{i_{1:m-1}}|s)}{\boldsymbol{\pi}^{i_{1:m-1}}(\mathbf{a}^{i_{1:m-1}}|s)} A_{\boldsymbol{\pi}}(s, \mathbf{a}) \right]. \tag{9}$$

For proof see Appendix D.1. One benefite of applying Equation (9) is that agents only need to maitain a joint advantage estiator $A_{\boldsymbol{\pi}}(\mathrm{s}, \mathbf{a})$ rather than one centralised critic for each individual agent (e.g., CTDE methods). Another practical benefit one can draw is that, given an estimator $\hat{A}(\mathrm{s}, \mathrm{a})$ of the advantage function $A_{\boldsymbol{\pi}_{\boldsymbol{\theta}_k}}(\mathrm{s}, \mathbf{a})$, for example GAE (Schulman et al., 2015b), we can estimate $\mathbb{E}_{\mathbf{a}^{i_{1:m-1}} \sim \boldsymbol{\pi}_{\boldsymbol{\theta}_{k+1}}^{i_{1:m-1}}, \mathrm{a}^{i_m} \sim \pi_{\theta}^{i_m}} \left[A_{\boldsymbol{\pi}_{\boldsymbol{\theta}_k}}^{i_m}\left(s, \mathbf{a}^{i_{1:m-1}}, \mathrm{a}^{i_m}\right) \right]$ with an estimator of

$$\left(\frac{\pi_{\theta}^{i_m}(\mathrm{a}^{i_m}|s)}{\pi_{\boldsymbol{\theta}_k}^{i_m}(\mathrm{a}^{i_m}|s)} - 1 \right) M^{i_{1:m}}(s, \mathbf{a}), \quad \text{where } M^{i_{1:m}} = \frac{\bar{\boldsymbol{\pi}}^{i_{1:m-1}}(\mathbf{a}^{i_{1:m-1}}|s)}{\boldsymbol{\pi}^{i_{1:m-1}}(\mathbf{a}^{i_{1:m-1}}|s)} \hat{A}(s, \mathbf{a}). \tag{10}$$

Notably, Equation (10) aligns nicely with the sequential update scheme in HATRPO. For agent i_m, since previous agents $i_{1:m-1}$ have already made their updates, the compound policy ratio for $M^{i_{1:m}}$ in Equation (10) is easy to compute. Given a batch $\mathcal{B}$ of trajectories with length T, we can estimate the gradient with respect to policy parameters as follows,

$$\hat{\boldsymbol{g}}_k^{i_m} = \frac{1}{|\mathcal{B}|} \sum_{\tau \in \mathcal{B}} \sum_{t=0}^{T} M^{i_{1:m}}(\mathrm{s}_t, \mathbf{a}_t) \nabla_{\theta^{i_m}} \log \pi_{\theta^{i_m}}^{i_m}(\mathrm{a}_t^i|\mathrm{s}_t) \Big|_{\theta^{i_m} = \theta_k^{i_m}}.$$

Next, we compute the Hessian of the expected KL-divergence, *i.e.*, $\boldsymbol{H}_k^{i_m}$. Finally, we can update $\theta_{k+1}^{i_m}$ by following Equation (8). The detailed pseudocode of HATRPO is listed in Appendix D.2.

4.2 HAPPO

To further alleviate the computation burden from $\boldsymbol{H}_k^{i_m}$ in HATRPO, one can follow the idea of PPO in Equation (4) by considering only using first order derivatives. This is achieved by making agent i_m choose a policy parameter $\theta_{k+1}^{i_m}$ which maximises the clipping objective of

$$\mathbb{E}_{\mathrm{s} \sim \rho_{\boldsymbol{\pi}_{\boldsymbol{\theta}_k}}, \mathbf{a} \sim \boldsymbol{\pi}_{\boldsymbol{\theta}_k}} \left[\min \left(\frac{\pi_{\theta^{i_m}}^{i_m}(\mathrm{a}^i|\mathrm{s})}{\pi_{\theta_k^{i_m}}^{i_m}(\mathrm{a}^i|\mathrm{s})} M^{i_{1:m}}(\mathrm{s}, \mathbf{a}), \mathrm{clip}\left(\frac{\pi_{\theta^{i_m}}^{i_m}(\mathrm{a}^i|\mathrm{s})}{\pi_{\theta_k^{i_m}}^{i_m}(\mathrm{a}^i|\mathrm{s})}, 1 \pm \epsilon \right) M^{i_{1:m}}(\mathrm{s}, \mathbf{a}) \right) \right]. \tag{11}$$

The optimisation process can be performed by stochastic gradient methods such as Adam (Kingma & Ba, 2014). We refer to the above procedure as HAPPO and Appendix D.3 for its full pseudocode.

4.3 Related Work

We are fully aware of previous attempts that tried to extend TRPO/PPO into MARL. Despite empirical successes, none of them managed to propose a theoretically-justified trust region protocol in multi-agent learning, or maintain the monotonic improvement property. Instead, they tend to impose certain assumptions to enable direct implementations of TRPO/PPO in MARL problems. For example, IPPO (de Witt et al., 2020a) assume homogeneity of action spaces for all agents and enforce parameter sharing. Yu et al. (2021) proposed MAPPO which enhances IPPO by considering a joint critic function and finer implementation techniques for on-policy methods. Yet, it still suffers similar drawbacks of IPPO due to the lack of monotonic improvement guarantee especially when the parameter-sharing condition is switched off. Wen et al. (2021) adjusted PPO for MARL by considering a game-theoretical approach at the meta-game level among agents. Unfortunately, it can only deal with two-agent cases due to the intractability of Nash equilibrium. Recently, Li & He (2020) tried to implement TRPO for MARL through distributed consensus optimisation; however, they enforced the same ratio $\bar{\pi}^i(a^i|s)/\pi^i(a^i|s)$ for all agents (see their Equation (7)) which, similar to parameter sharing, largely limits the policy space for optimisation. Moreover, their method comes with a δ/n KL-constraint threshold that fails to consider scenarios with large agent number.

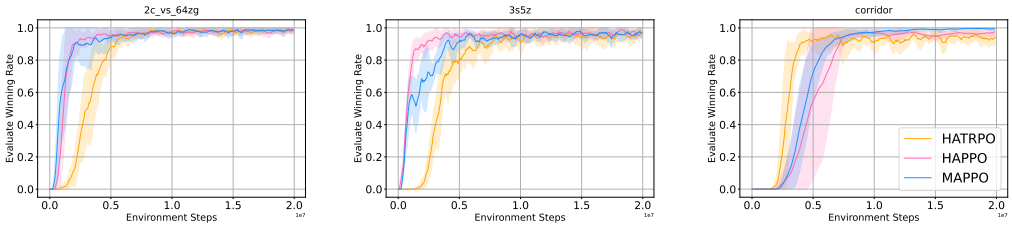

(a) 2c-vs-64zg (hard) (b) 3s5z (hard) (c) corridor (super hard)

Figure 2: Performance comparisons between HATRPO/HAPPO and MAPPO on three SMAC tasks. Since all methods achieve 100% win rate, we believe SMAC is not sufficiently difficult to discriminate the capabilities of these algorithms, especially when non-parameter sharing is not required.

One of the key ideas behind our HATRPO/HAPPO is the sequential update scheme. A similar idea of multi-agent sequential update was also discussed in the context of dynamic programming (Bertsekas, 2019) where artificial "in-between" states have to be considered. On the contrary, our sequential update sceheme is developed based on Lemma 1, which does not require any artificial assumptions and hold for any cooperative games. Furthermore, Bertsekas (2019) requires to maintain a fixed order of updates that is pre-defined for the task, whereas the order in HATRPO/MAPPO is randomised at each iteration, which also offers desirable convergence property (see Proposition 2).

Lastly, we would like to highlight the importance of the decomposition result in Lemma 1. This result could serve as an effective solution to value-based methods in MARL where tremendous efforts have been made to decompose the joint Q-function into individual Q-functions when the joint Q-function are decomposable (Rashid et al., 2018). Lemma 1, in contrast, is a general result that holds for any cooperative MARL problems regardless of decomposibility. As such, we think of it as an appealing contribution to future developments on value-based MARL methods.

5 Experiments and Results

We consider two most common benchmarks—StarCraftII Multi-Agent Challenge (SMAC) (Samvelyan et al., 2019) and Multi-Agent Mujoco (de Witt et al., 2020b)—for evaluating MARL algorithms. All hyperparameter settings and implementations details can be found in Appendix E.

StarCraftII Multi-Agent Challenge (SMAC). SMAC contains a set of StarCraft maps in which a team of ally units aims to defeat the opponent team. IPPO (de Witt et al., 2020a) and MAPPO (Yu et al., 2021) are known to achieve supreme results on this benchmark. By adopting parameter sharing, these methods achieve a winning rate of 100% on most maps, even including the maps that have heterogeneous agents. Therefore, we hypothesise that non-parameter sharing is not necessarily required and sharing policies is sufficient to solve SMAC tasks. We test our methods on two hard maps and one super-hard; results on Figure 2 confirm that SMAC is not sufficiently difficult to show off the capabilities of HATRPO/HAPPO when compared against parameter-sharing methods.

Multi-Agent MuJoCo. In comparison to SMAC, we believe Mujoco enviornment provides a more suitable testing case for our methods. MuJoCo tasks challenge a robot to learn an optimal way of motion; Multi-Agent MuJoCo models each part of a robot as an independent agent, for example, a leg for a spider or an arm for a swimmer. With the increasing variety of the body parts, modelling heterogeneous policies becomes necessary. Figure 3 demonstrate that, in all scenarios, HATRPO and HAPPO enjoy superior performance over those of parameter-sharing methods: IPPO and MAPPO, and also outperform non-parameter sharing MADDPG with (Lowe et al., 2017b) both in terms of reward values and variance. It is also worth noting that the performance gap between HATRPO and its rivals enlarges with the increasing number of agents. Meanwhile, we can observe that HATRPO outperforms HAPPO in almost all tasks; we believe it is because the hard KL constraint in HATRPO, compared to the clipping version in HAPPO, relates more closely to Algorithm 1 that attains monotonic improvement guarantee.

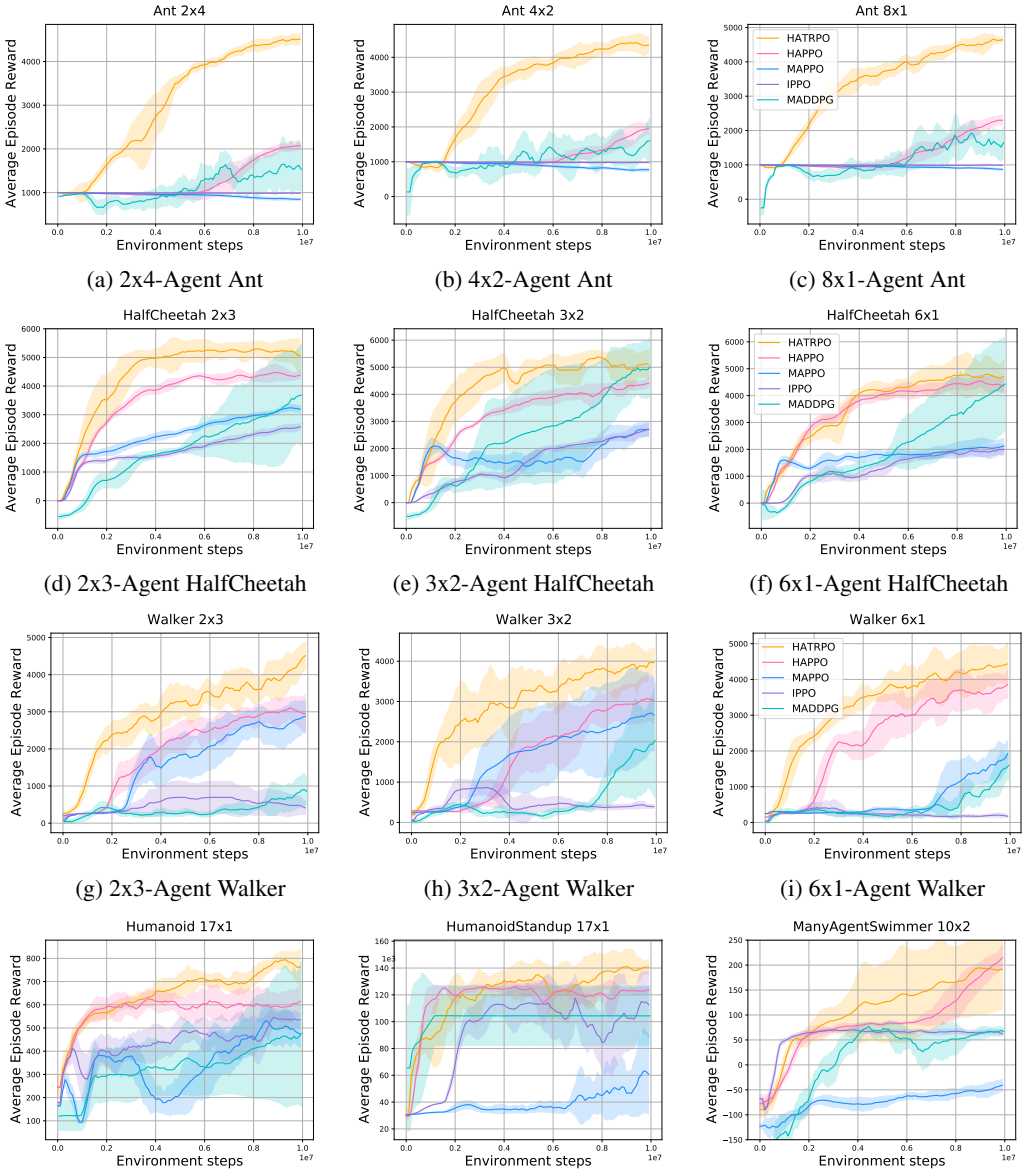

(a) 2x4-Agent Ant (b) 4x2-Agent Ant (c) 8x1-Agent Ant

(d) 2x3-Agent HalfCheetah (e) 3x2-Agent HalfCheetah (f) 6x1-Agent HalfCheetah

(g) 2x3-Agent Walker (h) 3x2-Agent Walker (i) 6x1-Agent Walker

(j) 17x1-Agent Humanoid (k) 17x1-Agent HumanoidStandup (l) 10x2-Agent Swimmer

Figure 3: Performance comparison on multiple Multi-Agent MuJoCo tasks. HAPPO and HATRPO consistently outperform their rivals, thus establishing a new state-of-the-art algorithm for MARL. The performance gap enlarges with increasing number of agents.

6 Conclusion

In this paper, we successfully apply trust region learning to multi-agent settings by proposing the first MARL algorithm that attains theoretically-justified monotonical improvement property. The key to our development is the multi-agent advantage decomposition lemma that holds in general with no need for any assumptions on agents sharing parameters or the joint value function being decomposable. Based on this, we introduced two practical deep MARL algorithms: HATRPO and HAPPO. Experimental results on both discrete and continuous control tasks (*i.e.*, SMAC and Multi-Agent Mujoco) confirm their state-of-the-art performance. For future work, we will consider incorporating the safety constraint into HATRPO/HAPPO so as to propose safety-aware MARL solutions.

ACKNOWLEDGEMENT

We would like to thank Shangding Gu and Derrick Goh Xin Deik for offering computational resources to run experiments, and Wang Jiaruijue for an early implementation of HATRPO.

REFERENCES

Dimitri Bertsekas. Multiagent rollout algorithms and reinforcement learning. *arXiv preprint arXiv:1910.00120*, 2019.

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

Christian Schroeder de Witt, Tarun Gupta, Denys Makoviichuk, Viktor Makoviychuk, Philip HS Torr, Mingfei Sun, and Shimon Whiteson. Is independent learning all you need in the starcraft multi-agent challenge? *arXiv preprint arXiv:2011.09533*, 2020a.

Christian Schröder de Witt, Bei Peng, Pierre-Alexandre Kamienny, Philip H. S. Torr, Wendelin Böhmer, and Shimon Whiteson. Deep multi-agent reinforcement learning for decentralized continuous cooperative control. *CoRR*, abs/2003.06709, 2020b.

Yan Duan, Xi Chen, Rein Houthooft, John Schulman, and Pieter Abbeel. Benchmarking deep reinforcement learning for continuous control. In *International conference on machine learning*, pp. 1329–1338. PMLR, 2016.

Jakob Foerster, Gregory Farquhar, Triantafyllos Afouras, Nantas Nardelli, and Shimon Whiteson. Counterfactual multi-agent policy gradients. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32, 2018.

Tuomas Haarnoja, Aurick Zhou, Pieter Abbeel, and Sergey Levine. Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. In *International Conference on Machine Learning*, pp. 1861–1870. PMLR, 2018.

Sham Kakade and John Langford. Approximately optimal approximate reinforcement learning. In *In Proc. 19th International Conference on Machine Learning*. Citeseer, 2002.

Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.

Jakub Grudzien Kuba, Muning Wen, Yaodong Yang, Linghui Meng, Shangding Gu, Haifeng Zhang, David Henry Mguni, and Jun Wang. Settling the variance of multi-agent policy gradients. *arXiv preprint arXiv:2108.08612*, 2021.

Hepeng Li and Haibo He. Multi-agent trust region policy optimization. *arXiv e-prints*, pp. arXiv–2010, 2020.

Michael L Littman. Markov games as a framework for multi-agent reinforcement learning. In *Machine learning proceedings 1994*, pp. 157–163. Elsevier, 1994.

Ryan Lowe, Yi Wu, Aviv Tamar, Jean Harb, Pieter Abbeel, and Igor Mordatch. Multi-agent actor-critic for mixed cooperative-competitive environments. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, pp. 6382–6393, 2017a.

Ryan Lowe, Yi Wu, Aviv Tamar, Jean Harb, Pieter Abbeel, and Igor Mordatch. Multi-agent actor-critic for mixed cooperative-competitive environments. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, pp. 6382–6393, 2017b.

David H Mguni, Yutong Wu, Yali Du, Yaodong Yang, Ziyi Wang, Minne Li, Ying Wen, Joel Jennings, and Jun Wang. Learning in nonzero-sum stochastic games with potentials. In Marina Meila and Tong Zhang (eds.), *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pp. 7688–7699. PMLR, 18–24 Jul 2021.

John Nash. Non-cooperative games. *Annals of mathematics*, pp. 286–295, 1951.

P Peng, Q Yuan, Y Wen, Y Yang, Z Tang, H Long, and J Wang. Multiagent bidirectionally-coordinated nets for learning to play starcraft combat games. arxiv 2017. *arXiv preprint arXiv:1703.10069*, 2017.

Tabish Rashid, Mikayel Samvelyan, Christian Schroeder, Gregory Farquhar, Jakob Foerster, and Shimon Whiteson. Qmix: Monotonic value function factorisation for deep multi-agent reinforcement learning. In *International Conference on Machine Learning*, pp. 4295–4304. PMLR, 2018.

Mikayel Samvelyan, Tabish Rashid, Christian Schroeder De Witt, Gregory Farquhar, Nantas Nardelli, Tim GJ Rudner, Chia-Man Hung, Philip HS Torr, Jakob Foerster, and Shimon Whiteson. The starcraft multi-agent challenge. *arXiv preprint arXiv:1902.04043*, 2019.

John Schulman, Sergey Levine, Pieter Abbeel, Michael Jordan, and Philipp Moritz. Trust region policy optimization. In *International conference on machine learning*, pp. 1889–1897. PMLR, 2015a.

John Schulman, Philipp Moritz, Sergey Levine, Michael Jordan, and Pieter Abbeel. High-dimensional continuous control using generalized advantage estimation. *arXiv preprint arXiv:1506.02438*, 2015b.

John Schulman, F. Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *ArXiv*, abs/1707.06347, 2017.

David Silver, Guy Lever, Nicolas Heess, Thomas Degris, Daan Wierstra, and Martin Riedmiller. Deterministic policy gradient algorithms. In *International conference on machine learning*, pp. 387–395. PMLR, 2014.

Peter Sunehag, Guy Lever, Audrunas Gruslys, Wojciech Marian Czarnecki, Vinicius Zambaldi, Max Jaderberg, Marc Lanctot, Nicolas Sonnerat, Joel Z Leibo, Karl Tuyls, et al. 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*, pp. 2085–2087, 2018.

R. S. Sutton, D. Mcallester, S. Singh, and Y. Mansour. Policy gradient methods for reinforcement learning with function approximation. In *Advances in Neural Information Processing Systems 12*, volume 12, pp. 1057–1063. MIT Press, 2000.

Ying Wen, Yaodong Yang, Rui Luo, Jun Wang, and Wei Pan. Probabilistic recursive reasoning for multi-agent reinforcement learning. In *International Conference on Learning Representations*, 2018.

Ying Wen, Yaodong Yang, and Jun Wang. Modelling bounded rationality in multi-agent interactions by generalized recursive reasoning. In Christian Bessiere (ed.), *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI-20*, pp. 414–421. International Joint Conferences on Artificial Intelligence Organization, 7 2020. Main track.

Ying Wen, Hui Chen, Yaodong Yang, Zheng Tian, Minne Li, Xu Chen, and Jun Wang. A game-theoretic approach to multi-agent trust region optimization. *arXiv preprint arXiv:2106.06828*, 2021.

Jiayi Weng, Huayu Chen, Dong Yan, Kaichao You, Alexis Duburcq, Minghao Zhang, Hang Su, and Jun Zhu. Tianshou: a highly modularized deep reinforcement learning library. *arXiv preprint arXiv:2107.14171*, 2021.

Yaodong Yang and Jun Wang. An overview of multi-agent reinforcement learning from game theoretical perspective. *arXiv preprint arXiv:2011.00583*, 2020.

Yaodong Yang, Rui Luo, Minne Li, Ming Zhou, Weinan Zhang, and Jun Wang. Mean field multi-agent reinforcement learning. In *International Conference on Machine Learning*, pp. 5571–5580. PMLR, 2018.

Yaodong Yang, Ying Wen, Jun Wang, Liheng Chen, Kun Shao, David Mguni, and Weinan Zhang. Multi-agent determinantal q-learning. In *International Conference on Machine Learning*, pp. 10757–10766. PMLR, 2020.

Chao Yu, A. Velu, Eugene Vinitsky, Yu Wang, A. Bayen, and Yi Wu. The surprising effectiveness of mappo in cooperative, multi-agent games. *ArXiv*, abs/2103.01955, 2021.

Haifeng Zhang, Weizhe Chen, Zeren Huang, Minne Li, Yaodong Yang, Weinan Zhang, and Jun Wang. Bi-level actor-critic for multi-agent coordination. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pp. 7325–7332, 2020.

Ming Zhou, Ziyu Wan, Hanjing Wang, Muning Wen, Runzhe Wu, Ying Wen, Yaodong Yang, Weinan Zhang, and Jun Wang. Malib: A parallel framework for population-based multi-agent reinforcement learning. *arXiv preprint arXiv:2106.07551*, 2021.