

Inducing Stackelberg Equilibrium through Spatio-Temporal Sequential Decision-Making in Multi-Agent Reinforcement Learning

Bin Zhang^{1,2}, Lijuan Li^{1,2}, Zhiwei Xu^{1,2}, Dapeng Li^{1,2}, Guoliang Fan^{1,2}

¹Institute of Automation, Chinese Academy of Sciences

²School of Artificial Intelligence, University of Chinese Academy of Sciences

{zhangbin2020, lijuan.li, xuzhiwei2019, lidapeng2020, guoliang.fan}@ia.ac.cn

Abstract

In multi-agent reinforcement learning (MARL), self-interested agents attempt to establish equilibrium and achieve coordination depending on game structure. However, existing MARL approaches are mostly bound by the simultaneous actions of all agents in the Markov game (MG) framework, and few works consider the formation of equilibrium strategies via asynchronous action coordination. In view of the advantages of Stackelberg equilibrium (SE) over Nash equilibrium, we construct a spatio-temporal sequential decision-making structure derived from the MG and propose an N-level policy model based on a conditional hypernetwork shared by all agents. This approach allows for asymmetric training with symmetric execution, with each agent responding optimally conditioned on the decisions made by superior agents. Agents can learn heterogeneous SE policies while still maintaining parameter sharing, which leads to reduced cost for learning and storage and enhanced scalability as the number of agents increases. Experiments demonstrate that our method effectively converges to the SE policies in repeated matrix game scenarios, and performs admirably in immensely complex settings including cooperative tasks and mixed tasks.

1 Introduction

Reinforcement learning (RL) is a popular method for solving sequential decision-making problems, but it faces several challenges when applied to a multi-agent system (MAS). In a constantly changing environment, agents' rewards are often closely tied to the actions of others. No individual exists in isolation and their actions are often interdependent. To achieve optimal outcomes, they must learn to find a stable equilibrium rather than simply maximizing their own returns. It is, therefore, an eternal topic in multi-agent reinforcement learning (MARL) to learn how to coordinate among agents to achieve the optimal joint policy.

Currently, the majority of prevalent MARL approaches adhere to the centralized training with decentralized execution (CTDE) architecture [Foerster *et al.*, 2016]. In this paradigm,

agent's policy model is trained centrally without communication constraints but executed in a distributed manner. Nevertheless, current CTDE approaches prioritize the acquisition of comprehensive cognition about the environment and the formulation of cooperative policies by directly getting the global state [Rashid *et al.*, 2018] or the local observations and actions of all agents [Lowe *et al.*, 2017]. Some works based on game abstraction [Zhang *et al.*, 2023a] model the relationship of agents based on their observational data, which are state representation learning essentially. However, centralized training is not leveraged to its full potential. Actually, it is feasible to construct an interaction mechanism that effectively guides agents towards mandatory coordination.

Moreover, game theory offers the fundamental framework for the interaction of multiple intelligent agents, and several MARL algorithms, such as Nash Q-Learning [Hu and Wellman, 2003], Mean Field Q-learning [Yang *et al.*, 2018] and HATRPO [Kuba *et al.*, 2022], seek for convergence to Nash Equilibrium (NE). However, it is worth noting that many game problems feature more than one NE, and the diverse NE policies chosen by different individuals frequently lead to undesirable outcomes. In addition, in a two-player zero-sum game, the MaxMin operator [Littman, 1994] can be used to calculate NE strategies for each agent, such that neither player can gain an advantage by deviating from their chosen strategy. But in more complex general-sum game situations, finding NE strategies is typically rather difficult. Consequently, we intend to concentrate on the Stackelberg equilibrium (SE) [Von Stackelberg, 2010], in which agents make decisions in a leader-follower framework, with leaders prioritizing decision-making and enforcing their policies on followers who respond rationally to this enforcement. It has been shown that SE is a superior convergence objective for MARL compared to NE, whether from the certainty of equilibrium solution or Pareto superiority [Zhang *et al.*, 2020]. Actually, when SE encounters MARL, we aim to address the following challenges: (a) How to converge to SE policies that rely on asynchronous action coordination under the Markov game framework where agents act simultaneously? (b) How to extend the method to scenarios with more than two agents?

To address the aforementioned issues, we draw inspiration from techniques used in single-agent RL for high-dimensional continuous control tasks [Metz *et al.*, 2017]. Utilizing significant benefits of centralized training, we model

MARL as a sequential decision-making problem in both temporal and spatial dimensions, and then train the heterogeneous SE policies asymmetrically. Furthermore, to facilitate the execution of SE policies in a communication-free and symmetrical environment, we introduce an N-level policy model based on a conditional hypernetwork shared by all agents, with synthetic targets of generating the weights of target agents' policy networks according to their priority attributes. This approach allows us to bridge the gap between environmental communication restrictions and the requirements that SE policy needs to access superior agents' policies. It mitigates the problem of suboptimal solutions induced by parameter sharing, and overcomes the issue of learning and storage costs caused by heterogeneous policy learning. Our primary contributions are summarized as follows:

- We develop a spatio-temporal sequential Markov game framework based on agent priority that enables agents to establish an efficient interaction mechanism during centralized training.
- We construct an N-level policy model to assist agents in executing SE policies in a fully decentralized setting without imposing a limit on the number of agents.
- We establish the asymmetric training with symmetric execution paradigm, a significant augmentation of CTDE.
- Our method is demonstrated to converge to the SE through repeated matrix game experiments, and the results in more complicated scenarios also illustrate its superiority over powerful benchmarks in terms of sample efficiency and overall performance.

2 Preliminaries

2.1 Markov Game

The multi-agent decision-making problem is typically described as a Markov game (MG), which can be defined by a tuple $\Gamma \triangleq \langle \mathcal{I}, \mathcal{S}, \{\mathcal{A}^i\}_{i \in \mathcal{I}}, P, \{r^i\}_{i \in \mathcal{I}}, \gamma \rangle$. $\mathcal{I} = \{1, 2, \dots, n\}$ denotes the set of agents and $s \in \mathcal{S}$ is the global state set of the environment. \mathcal{A}^i denotes the action space of agent i and the joint action space $\mathcal{A} = \prod_{i=1}^n \mathcal{A}^i$ is the product of action spaces of all agents. $P : \mathcal{S} \times \mathcal{A} \rightarrow \Omega(\mathcal{S})$ denotes state transition function, where $\Omega(X)$ is the set of probability distributions in X space. $r^i : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$ is the reward function of agent i and γ is the discount factor. At time step t , each agent chooses an action $a_t^i \in \mathcal{A}^i$ at state $s_t \in \mathcal{S}$ based on its own policy $\pi^i : \mathcal{S} \rightarrow \Omega(\mathcal{A}^i)$ and receives feedback in the form of $r^i(s_t, \mathbf{a}_t)$, where $\mathbf{a}_t = (a_t^1, \dots, a_t^n) \in \mathcal{A}$. The environment moves to a new state $s_{t+1} \sim P(s_{t+1} | s_t, \mathbf{a}_t)$ as a result of joint action \mathbf{a}_t . The joint policy of all agents is expressed as $\pi(s_t) = \prod_{i=1}^n \pi^i(s_t)$. Under the framework of MG, the state value function of agent i is defined as :

$$V_{\pi}^i(s) = \mathbb{E}_{s \sim P, \mathbf{a} \sim \pi^{-i}} \left[\sum_{t=0}^{\infty} \gamma^t r_t^i(s_t, \mathbf{a}_t) \mid s_0 = s, \mathbf{a}_t \sim \pi^i(s_t) \right], \quad (1)$$

where $-i$ represents all agents except i . According to Bellman equation, the action-state value function is denoted as:

$$Q_{\pi}^i(s, \mathbf{a}) = r^i(s, \mathbf{a}) + \gamma \sum_{s' \in \mathcal{S}} P(s' | s, \mathbf{a}) \cdot V_{\pi}^i(s'). \quad (2)$$




Figure 1: The STMG state transition procedure. It is an extensive game version of MG, which specifies the decision-making sequence of agents simultaneously.

2.2 Multi-Agent Reinforcement Learning

In a multi-agent system, each decision-maker strives to maximize their own expected utility, denoted as $\mathcal{J}^i(\theta) = \mathbb{E}_{s \sim P, \mathbf{a} \sim \pi_{\theta}} \left[\sum_{t=0}^{\infty} \gamma^t r_t^i \right]$, where $\theta = \{\theta^i\}_{i \in \mathcal{I}}$ represents the policy parameters of agents. However, using gradient ascent to directly update the strategy, i.e., $\theta^i \leftarrow \theta^i + \alpha \nabla_{\theta^i} \mathcal{J}^i(\theta)$, usually fails because $\mathcal{J}^i(\theta)$ is affected by all agents and the gradient update directions among agents may conflict. Consequently, it is vital to employ the concept of game equilibrium solutions in order to develop effective coordination strategies. One common solution objective is NE. Nash-Q learning, for example, computes agent i 's value function $V_{Nash}^i(s)$ when all agents follow the NE strategy in each stage of the game, and updates the action-state value function $Q_{Nash}^i(s, \mathbf{a})$ through Eq. (2). The algorithm can converge to NE strategies under the strong assumption that equilibrium exists at every game stage. All agents adopt equilibrium policies as their own convergence objective rather than selfishly maximizing their own individual utilities. With the emergence of deep learning, current popular MARL approaches frequently employ the CTDE paradigm to enable coordination via credit assignment [Rashid *et al.*, 2018] or centralized critics [Lowe *et al.*, 2017].

2.3 Stackelberg Leadership Model

The Stackelberg leadership model dictates the order of actions among agents. Consider a two-player game, the leader can enforce his own strategy π^1 to the follower, who will react after observing the leader's behavior $\pi^2 : \mathcal{S} \times \mathcal{A}^1 \rightarrow \Omega(\mathcal{A}^2)$, where $\pi^i \in \Pi^i$ and Π^i represents the policy space. In this model, leader optimizes its objective on the premise that follower will provide the best response $BR(\pi^1)$, while the follower tends to maximize its expected utility based on leader's preconditions. It can be formalized as:

$$\begin{aligned} \max_{\pi^1 \in \Pi^1} \{ \mathcal{J}^1(\pi^1, \pi^2) \mid \pi^2 = BR(\pi^1) \}, \\ \max_{\pi^2 \in \Pi^2} \{ \mathcal{J}^2(a^1, \pi^2) \mid a^1 \sim \pi^1 \}. \end{aligned} \quad (3)$$

Equilibrium signifies that in a multi-player game, all players have adopted the optimal strategy and none can improve their performance by altering their own strategy. Correspondingly, Stackelberg equilibrium signifies that both agents adhere to the optimal solution strategies (π_{SE}^1, π_{SE}^2) of the the above optimization issue. It satisfies:

$$\begin{aligned} V_{\pi_{SE}^1, BR(\pi_{SE}^1)}^1(s) \geq V_{\pi^1, BR(\pi^1)}^1(s), \\ \pi_{SE}^2 = BR(\pi_{SE}^1). \end{aligned} \quad (4)$$




Figure 2: Exemplification of the execution mechanism under the STMG framework. In the temporal domain, all agents continue to adhere to the MG settings. Simultaneously, agents make decisions according to the Stackelberg leadership model, and subgame states s_t^{i-1} are introduced to maintain each agent’s policy inputs.

As a mandatory equilibrium, SE offers more benefits than NE in terms of the stability, certainty of equilibrium points and Pareto superiority. Consequently, it is a more suitable learning objective.

3 Methodology

In this section, we propose a new method called Spatio-Temporal sequence Equilibrium Policy optimization (STEP). This method leverages the advantages of centralized training to establish an interaction mechanism and employs a more efficient policy network to facilitate the achievement of coordinate policies during decentralized execution. The following parts provide a detailed description of the procedure for implementing STEP.

3.1 Spatio-Temporal Sequential Markov Game

Current CTDE approaches, especially policy-based methods, tend to center around enhancing the richness of input information rather than boosting the effectiveness of collaboration. Actually, centralized training liberates us from restrictive circumstances. With centralized training, we can definitely create an interaction mechanism among agents for completing their tasks. Furthermore, the concept of the SE is more in line with the goals of RL as it requires merely additional constraints, which can be easily achieved by setting a specific interaction mechanism. So naturally, we utilize the form of multi-player Stackelberg leadership model and construct the spatio-temporal sequential Markov game (STMG) in the training phase to direct agents coordination and facilitate the implementation of SE policies.

Definition 1. *STMG can be formalized as a tuple $\langle \mathcal{I}, \mathcal{S}, \{\mathcal{A}^i\}_{i \in \mathcal{I}}, P, \{\nabla^i\}_{i \in \mathcal{I}}, \gamma, \{o^i\}_{i \in \mathcal{I}} \rangle$. In addition to the MG defined in Section 2.1, STMG add the term o^i , which denotes the action order of agent i . $\mathcal{O} = (o^1, \dots, o^n)$ represents the action order of all agents, indicating the priority/importance of agents at the decision-making stage.*

Figure 1 shows the state transition process in STMG. For the sake of simplicity, we assume that the agent ID i is assigned by priority o^i . Compared with MG, STMG assumes the form of a sequence decision in both temporal and spatial domains. Agents with higher priorities have greater initiative, whereas agents with lower priorities must respond to the

actions of those with higher priorities. Correspondingly, the policy of each agent changes to $\pi^i : \mathcal{S} \times \mathcal{A}^1 \dots \mathcal{A}^{i-1} \rightarrow \Omega(\mathcal{A}^i)$. The action-state value function of agent i can be written as:

$$Q_{\pi}^i(s, a^{1:i-1}, a^i) = \mathbb{E}_{s \sim P, a^{i+1:n} \sim \pi^{i+1:n}} \left[\sum_{t=0}^{\infty} \gamma^t r_t^i(s_t, \mathbf{a}_t) \mid s_0 = s, \mathbf{a}_0 = \mathbf{a}, a_t^i \sim \pi^i(s_t, a_t^1, \dots, a_t^{i-1}) \right]. \quad (5)$$

The state value function is denoted as:

$$V_{\pi}^i(s, a^{1:i-1}) = \sum_{a^i \in \mathcal{A}^i} \pi^i(a^i | s, a^{1:i-1}) Q_{\pi}^i(s, a^{1:i-1}, a^i). \quad (6)$$

And then we have the advantage function:

$$A_{\pi}^i(s, a^{1:i-1}, a^i) = Q_{\pi}^i(s, a^{1:i-1}, a^i) - V_{\pi}^i(s, a^{1:i-1}). \quad (7)$$

Within the STMG framework, agent with priority i can consider the actions of the preceding $i - 1$ superior agents as prerequisites, maximize its own expected return $\mathcal{J}^i(\theta) = \mathbb{E}_{s \sim P, a^{i+1:n} \sim \pi_{\theta}^{i+1:n}} [\sum_{t=0}^{\infty} \gamma^t r_t^i]$ according to the current subgame state $s^{i-1} = (s, a^1, \dots, a^{i-1})$, which is equivalent to learning the best response to the strategies of superior agents. In summary, during the training phase, agents’ execution mechanism can be divided into the Markov game process in the time domain and the Leader-Follower model in the space domain. The specific execution process is depicted in Figure 2.

3.2 Parameterized N-level Policy Model Using A Conditional Hypernetwork

Within the STMG framework, agents are able to acquire the actions of superior agents directly. This asymmetric approach arises naturally in the centralized training phase without communication restrictions, but it is obviously illegal in the decentralized execution phase. To this end, several approaches can be taken. We can train a policy network that is shared by all agents, allowing them to calculate the actions of other agents using the shared parameters. Designing a communication module that enables superior agents to broadcast their decisions to inferior agents is also an alternative. Moreover, agents can keep copies of other players’

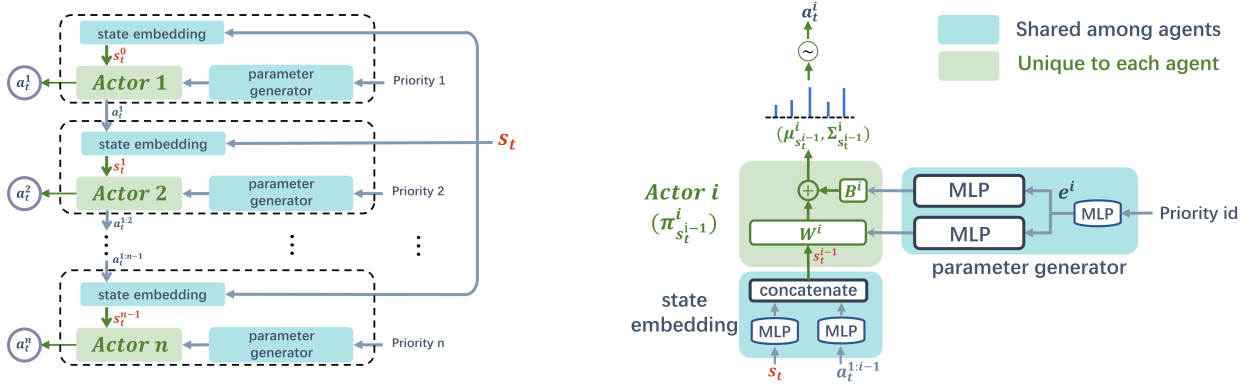


Figure 3: The overall architecture of STEP. *Left*: The workflow of STEP for a comprehensive decision in a time step. Agents make their decisions based on the current situation s_t , their self-positioning *Priority* ID, and the prerequisite actions $a_t^{1:i-1}$ of superior agents. *Right*: The structure of N-level policy model. It allows for the implementation of heterogeneous policies under parameter sharing and the Stackelberg equilibrium policies under symmetric conditions.

policy or value networks to independently calculate their own policies. Nonetheless, parameter sharing can result in suboptimal solutions, and communication may not be feasible in fully decentralized execution settings or with communication bandwidth constraints. Additionally, the scalability of the algorithm may be limited as the number of agents increases and the volume of saved model copies grows.

In contrast to the aforementioned methods, we develop an N-level policy model using a conditional hypernetwork. It consists of a parameter generator, a state embedding, and a target policy module. Figure 3 illustrates its structure in detail. MARL is essentially a multi-task learning (multi-objective regression) problem [Iqbal and Sha, 2019]. Instead of explicitly training and storing individual policy modules for each agent, we maintain a set of policy parameters $\{\theta_{tar}^i\}_{i \in \mathcal{I}}$ by training a meta-model (conditional hypernetwork) $\mathcal{H}(e^i, \theta_h) = \theta_{tar}^i$ with the weights θ_h shared by all agents. It maps the embedding of each agent’s priority ID e^i to the parameter configuration of its policy network. It has been demonstrated that hypernetworks can handle continuous learning tasks [von Oswald *et al.*, 2020], where a model is trained on a series of tasks in sequence. This is highly compatible with our expectations. In addition, the input s^{i-1} to the target policy module is encoded by the state embedding network $\mathcal{E}(s, a^1, \dots, a^{i-1}; \theta_s)$, which receives as input the current state of the environment and the actions of superior agents. Specifically, each component of the model is described as follows:

$$\text{State Embedding: } s^{i-1} = \mathcal{E}(s, a^1, \dots, a^{i-1}; \theta_s), \quad (8)$$

$$\text{Parameter Generator: } \theta_{tar}^i = \mathcal{H}(e^i; \theta_h), \quad (9)$$

$$\text{Target Policy: } a^i \sim \pi^i(s^{i-1}; \theta_{tar}^i). \quad (10)$$

It is important to note that the parameters of the policy model consist of $\theta^i = (\theta_s, \theta_h, \theta_{tar}^i)$, where θ_s, θ_h are learnable parameters shared by all agents and θ_{tar}^i can be accessed by all agents through the parameter generator. By using the N-level policy model, we are able to share policy parameters while avoiding the suboptimal solutions and storage issues

previously mentioned, resulting in efficient decentralized coordination.

3.3 Implementation

The N-level policy model allows for the implementation of the asymmetric training with symmetric/decentralized execution (ATSE) paradigm in STEP. Meanwhile, we select Proximal Policy Optimization (PPO) [Schulman *et al.*, 2017] as the underlying algorithm due to its superior performance. To ensure the balance between exploration and exploitation during the training phase, agents sample actions from the categorized distribution or multivariate Gaussian distribution generated by the policy model. However, due to the unpredictability of sampling, agents are unable to calculate the actions of superior agents. As a result, we directly transmit the behaviors of superior agents to inferior agents during the training process. During the execution process, each agent selects the action with the highest probability, and inferior agents can easily compute actions of superior agents through the shared policy model to establish corresponding coordination. Additionally, throughout each epoch of the training process, each agent trains the identical policy module using its own data, which can result in catastrophic forgetting. To address this issue, we incorporate a regularization item to ensure that the policy network can train the current agent’s policy parameters while maintaining the capacity to fit the previously updated policies of other agents. The objective function of actor i is expressed as:

$$\begin{aligned} \mathcal{L}^i(\theta) &= \mathbb{E}_{s \sim P, \mathbf{a} \sim \pi} [\mathcal{L}_{clip} - \mathcal{L}_h(\theta_h, \theta_{h_{old}}, e^{1:i-1}) \\ &\quad + \eta S(\pi_\theta^i(s, a^{1:i-1}))], \\ \mathcal{L}_{clip} &= \min(r_\theta^i A_\pi^i, clip(r_\theta^i, 1 \pm \epsilon) A_\pi^i), \\ \mathcal{L}_h(\cdot) &= \frac{\beta}{i-1} \sum_{j=1}^{i-1} \|\mathcal{H}(e^j; \theta_h) - \mathcal{H}(e^j; \theta_{h_{old}})\|_2^2, \end{aligned} \quad (11)$$

where $S(\cdot)$ is Shannon entropy used for strengthening exploration, $r_\theta^i = \frac{\pi_\theta^i(a^i | s, a^{1:i-1})}{\pi_{\theta_{old}}^i(a^i | s, a^{1:i-1})}$ is the likelihood ratio

between the current and previous policies, ϵ is the clipping ratio, η and β are coefficients of entropy and regularization term respectively. Critic network is used to fit the value function, and its loss function is expressed as:

$$\mathcal{L}(\phi^i) = \max \left[(V_\phi(s, a^{1:i-1}) - R^i)^2, \right. \\ \left. (\text{clip}(V_\phi(s, a^{1:i-1}), V_{\phi_{old}}(s, a^{1:i-1}) \pm \epsilon) - R^i)^2 \right], \quad (12)$$

where R^i is the cumulative return and ϵ is the clipping ratio.

The pseudo-code of STEP can be found in Appendix. Notably, our method provides comprehensive advantages over currently popular MARL techniques. Unlike prior work based on SE [Könönen, 2004; Zhang *et al.*, 2020], STEP is easily extensible to include additional players; Unlike prior work based on parameter sharing [Rashid *et al.*, 2018], STEP is able to learn heterogeneous policies; In contrast to previous work on learning heterogeneous policies [Kuba *et al.*, 2022], STEP does not increase the learning cost as the number of agents increases; In contrast to previous CTDE work [Lowe *et al.*, 2017; Yu *et al.*, 2021], we no longer focus on merely expanding the information space, but on specific forms of coordination.

4 Related Work

Our study, inspired by the Stackelberg leadership model, builds an asymmetric cooperative connection within the MG framework and applies RL to discover the Stackelberg equilibrium solution. In certain publications, SE strategies also serve as the convergence objective. Similar to Nash-Q learning, Asymmetric Q-learning (AQL) [Könönen, 2004] updates the action-state value function in an asymmetric setting by calculating the SE of the stage game at each iteration. Leaders are able to access followers' reward information and save copies of each follower's value function. Bi-AC [Zhang *et al.*, 2020] presents a bi-level actor-critic method based on CTDE paradigm that employs a Q-learning-based leader [Mnih *et al.*, 2013] and DDPG-based follower [Lillicrap *et al.*, 2016]. Both the leader and follower must save the leader's critic network and the follower's actor network during the execution process in order to compute and execute their policies. Additionally, Bully [Littman and Stone, 2001] and DeDOL [Tian *et al.*, 2019] are also dedicated to the solution of SE. However, all of the previous works only divide the agents into two levels, and extending them to more agents is not as straightforward as they anticipated. Furthermore, HATRPO [Kuba *et al.*, 2022] uses random sequence updating to ensure monotonic policy improvement in multi-agent settings, but it is only suitable for cooperative environments. The algorithm requires to define a joint advantage function for all agents, limiting its application to mixed tasks.

When dealing with structured combination action spaces, single-agent RL also uses similar spatial-temporal sequence learning approaches. SQL [Metz *et al.*, 2017] imitates the sequence-to-sequence model of the structured prediction problem, predicts the policy of each dimension simultaneously, and combines them into a comprehensive high-dimensional policy. Our technique can be considered as an extension of SQL for multi-agent systems. The actions

| $a^1 \backslash a^2$ | a_1^2 | a_2^2 | a_3^2 |
|----------------------|----------|----------|-----------|
| a_1^1 | k | 0 | 10 |
| a_2^1 | 0 | 2 | 0 |
| a_3^1 | 8 | 0 | k |

| $a^1 \backslash a^2$ | a_1^2 | a_2^2 | a_3^2 |
|----------------------|-------------|--------------|---------|
| a_1^1 | 0, 5 | -10,-5 | -8, 4 |
| a_2^1 | -5,-10 | -5, 0 | -15,-5 |
| a_3^1 | 5, 0 | -10,-5 | -10, 5 |

Figure 4: Matrix games. *Left*: The common-payoff matrix of the Penalty game, where $k = [0, -25, -50, -75, -100]$. (a_1^1, a_3^2) , (a_2^1, a_2^2) and (a_3^1, a_2^2) are NE points, and (a_1^1, a_3^2) is the only SE point. *Right*: The payoff matrix of the Mixing game. It has only one NE point (a_2^1, a_2^2) and only one SE point (a_1^1, a_1^2) . The SE is Pareto superior to the NE.

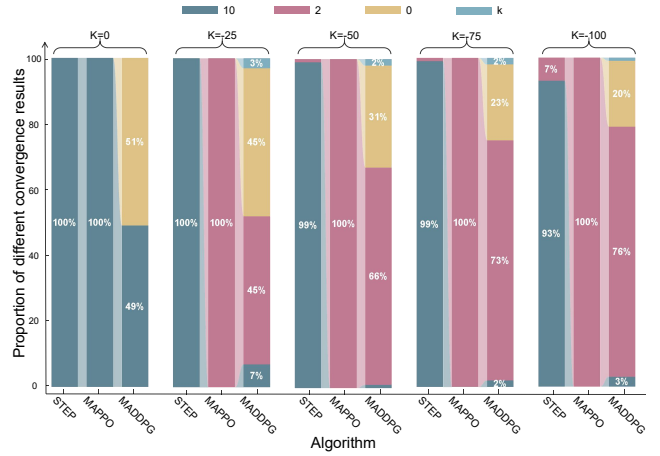


Figure 5: Proportion of different convergence results in Penalty.

of each dimension in SQL corresponds to the actions of each agent in STEP, and the complete high-dimensional policy relates to the joint policy of all agents. In addition, BQD's [Tavakoli *et al.*, 2018] use of the action branch structure is also an effective way to decompose and separately control high-dimensional actions. This form of single-agent algorithm lacks communication restrictions and can freely access each dimension's policy, hence reducing the complexity of problem-solving. In MARL, however, it is required to build extra procedures to address communication restriction-related issues.

5 Experiments

We evaluate the performance of our proposed algorithm, STEP, in three benchmark environments: the Repeated Matrix Game, the Multi-Agent MuJoCo [Peng *et al.*, 2021], and the Highway On-Ramp Merging [Chen *et al.*, 2021]. Appendix contains thorough descriptions of the three environments. Based on the reward settings (shared or individual rewards) and control types (discrete or continuous action space) in various environments, we compare STEP with various state-of-the-art MARL algorithms. The main objectives of these experiments are to: (a) Validate STEP's ability to identify the Stackelberg equilibrium policies; (b) Evaluate its performance in more challenging cooperative and mixed tasks; and (c) Assess the effect of the hyperparameter on its performance.

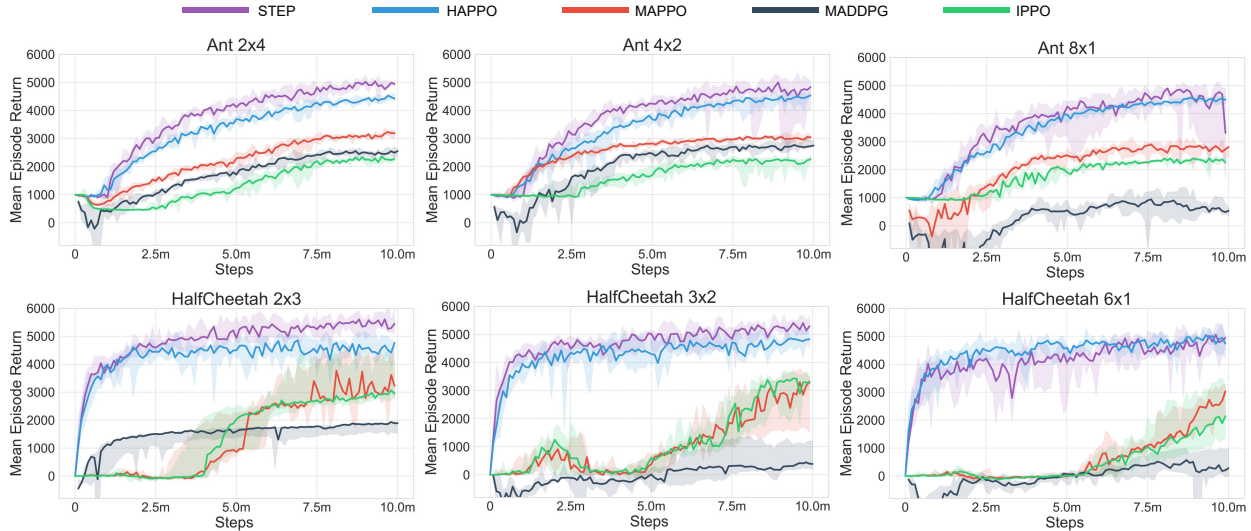


Figure 6: Performance comparison with baselines on Multi-Agent MuJoCo tasks. Error bars are a 95% confidence interval across 5 runs.

| | Agent 1 | | | Agent 2 | | | Reward | | |
|--------|---------------------|-------------|--------------|---------------------|-------------|--------------|------------|-------------|--------------------|
| | a_1^1 | a_2^1 | a_3^1 | a_1^2 | a_2^2 | a_3^2 | agent 1 | agent 2 | Avg. |
| STEP | 99.99 (0.01) | 0.00(0.00) | 0.01(0.01) | 99.99 (0.01) | 0.00(0.00) | 0.01(0.01) | 0.00(0.00) | 5.00(0.00) | 2.50 (0.00) |
| MAPPO | 32.93(46.20) | 0.03(0.20) | 67.04(46.21) | 61.98(47.31) | 0.04(0.22) | 37.97(47.31) | 2.72(2.38) | -1.29(6.43) | 0.72(2.33) |
| MADDPG | 21.39(37.14) | 5.62(20.52) | 72.99(39.82) | 14.46(31.81) | 7.86(22.51) | 77.68(37.30) | 3.20(3.26) | -7.9(5.12) | -2.35(2.84) |

Table 1: Results of different methods in Mixing game. The first two columns show the average probability of two agents choosing different actions. The third column displays the individual rewards earned by each agent and the average rewards of the two agents. The values in parentheses correspond to a single standard deviation over trials.

5.1 Repeated Matrix Game

We test STEP using the cooperative matrix game Penalty proposed by Claus and Boutilier [1998] and the non-cooperative matrix game Mixing shown in Figure 4. Agents receive constant observations at each time step in both games, and the length of repeated matrix games is set to 25. To cultivate formidable critics, the popular CTDE paradigm utilises combined observation and action to enlarge the information space. However, there is no noticeable performance difference between independent learning (IL) and CTDE for empty state matrix games. Therefore, only MAPPO [Yu *et al.*, 2021] and MADDPG [Lowe *et al.*, 2017] are considered as baselines for our research. We present the experimental results over 100 trials of 10000 time steps in Figure 5 and Table 1, which show that STEP outperforms the baselines in both types of matrix games. More detailed results can be found in Appendix.

In Penalty, there are three NE points, and the only SE point corresponds to the optimal NE point. The challenge lies in the fact that during the exploration phase, any deviation from the optimal NE policy by any of the players will result in severe punishment for all players. The likelihood that agents choose the best NE action decreases as the penalty term increases. Consequently, MAPPO always falls into the trap of penalty and converges to the sub-optimal NE point (a_2^1, a_2^2) at a rate of 100%, unless $k = 0$. Similarly, when k declines,

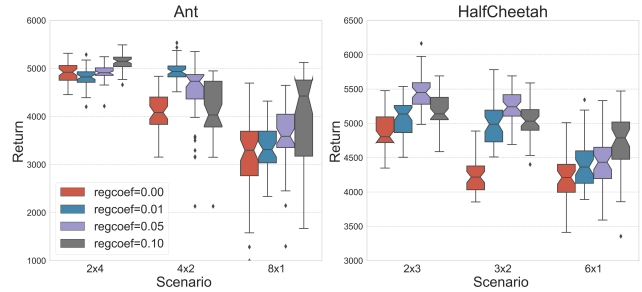


Figure 7: Results of STEP with different regularization coefficients. Note that a larger coefficient produces better results in the case of more agents.

MADDPG has a greater probability of converging to the sub-optimal NE point. Compared to the baselines, SE is the convergence objective for STEP. By computing the leader’s action, the follower is always able to select the best response action, allowing it to efficiently arrive at the optimal solution regardless of the configuration. Even under a high penalty of $k = -100$, there is a 93% probability of convergence to the optimal value. Due to the use of neural network approximation, modest probability errors are permitted.

Table 1 displays the results of different methods in Mixing




Figure 8: Learning curves for the Highway On-Ramp Merging tasks. Error bars are a 95% confidence interval across 5 runs.

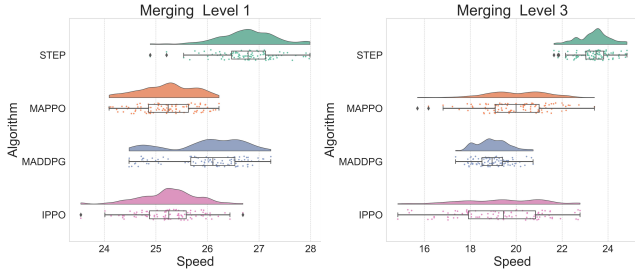


Figure 9: Comparing the testing performance of the average speed among different methods.

game. We evaluate the agents’ coordination based on the average return proposed by Zhang [2018], as each agent has its own particular reward. It is clear that there are two optimal solutions in this scenario, with (a_1^1, a_1^2) corresponding to the SE. In addition, the SE is Pareto superior to the only NE. As expected, STEP is able to rapidly converge to the SE with a probability of nearly 100%. In contrast, MAPPO and MADDPG are utterly ineffective and cannot even converge. It also demonstrates that STEP can effectively solve the problem of selecting multiple optimal solutions to some extent due to the certainty of the SE.

5.2 Multi-Agent MuJoCo

The Multi-Agent MuJoCo environment, which divides each part of the robot into separate agents that operate each part independently to achieve optimal motion, is used to assess the performance of the algorithm in a fully cooperative setting, as well as the impact of the regularization term coefficient on STEP’s performance. Specifically, to measure the collaboration quality of the agents, we choose representative cooperation scenarios such as Ant and HalfCheetah. And the results in Many-Agent Ant scenario can be found in Appendix due to the space limit.

Figure 6 demonstrates that STEP outperforms heterogeneous and homogeneous methods in all cases, which is certainly a result of the hierarchical structure among agents. Notably, the performance gap between STEP and HAPPO narrows as the number of agents increases. We think it’s because STEP achieves the effect of heterogeneous policies through parameter sharing. The number of STEP parameters remains constant regardless of the number of agents, whereas HAPPO with heterogeneous policies builds an individual policy network for each agent. In this sense, STEP shows a higher capacity to capture and represent diverse and intricate policies.

Figure 7 depicts the influence of various regularization coefficients on the experimental results. When the number of agents is small, the performance of various coefficients is similar. When there are more agents, however, greater coefficients lead to better results. The absence of regularization coefficients, on the other hand, has a negative impact on the experimental outcomes, demonstrating the efficiency of the regularization term. This outcome is expected and intuitive.

5.3 Highway On-Ramp Merging

Lane merging in high-density, high-speed traffic presents a significant challenge for both autonomous vehicles and human drivers. To guarantee safe and efficient merging, vehicles must apply coordination mechanisms that allow them to merge at precise speeds and times. In this scenario, each agent has its own goal, and all of them seek to pass through the main road quickly. Due to the fact that HAPPO is built for a fully cooperative environment, we only employ MAPPO, IPPO, and MADDPG to compare with STEP in this mixed task scenario. The road conditions vary at different levels.

Results in Figure 8 show that STEP outperforms the other approaches in the mixed task scenario, with the performance gap between STEP and the other algorithms growing as traffic density and complexity rise. Additionally, the average speed of STEP, as illustrated in Figure 9, is also faster than that of the other algorithms after convergence occurs.

6 Conclusion

In this study, we propose a multi-agent reinforcement learning method called STEP, which is inspired by the Stackelberg leadership model, multi-task learning, as well as high-dimensional continuous action decomposition for single agents. The core insight behind STEP is to make sequential decisions in both temporal and spatial domains that drives MARL to converge to the Stackelberg equilibrium. During the centralized training phase, we construct an efficient interaction mechanism, namely the spatio-temporal sequence Markov game, and utilize a shared N-level policy model to learn policies that aid the agents in achieving efficient coordination during the decentralized execution phase. The experimental results demonstrate that our STEP approach achieves state-of-the-art performance. We believe that asynchronous action coordination and the spatio-temporal sequential decision-making model have further development potential, and how to leverage the advantages of the game structure is a topic worthy of further research.

References

- [Chen *et al.*, 2021] Dong Chen, Zhaojian Li, Yongqiang Wang, Longsheng Jiang, and Yue Wang. Deep multi-agent reinforcement learning for highway on-ramp merging in mixed traffic. *arXiv preprint arXiv:2105.05701*, 2021.
- [Chen *et al.*, 2022] Yiqun Chen, Hangyu Mao, Tianle Zhang, Shiguang Wu, Bin Zhang, Jianye Hao, Dong Li, Bin Wang, and Hongxing Chang. Ptdc: Personalized training with distilled execution for multi-agent reinforcement learning. *arXiv preprint arXiv:2210.08872*, 2022.
- [Claus and Boutilier, 1998] Caroline Claus and Craig Boutilier. The dynamics of reinforcement learning in cooperative multiagent systems. In Jack Mostow and Chuck Rich, editors, *Proceedings of the Fifteenth National Conference on Artificial Intelligence and Tenth Innovative Applications of Artificial Intelligence Conference, AAAI 98, IAAI 98, July 26-30, 1998, Madison, Wisconsin, USA*, pages 746–752. AAAI Press / The MIT Press, 1998.
- [Foerster *et al.*, 2016] Jakob N. Foerster, Yannis M. Assael, Nando de Freitas, and Shimon Whiteson. Learning to communicate with deep multi-agent reinforcement learning. In Daniel D. Lee, Masashi Sugiyama, Ulrike von Luxburg, Isabelle Guyon, and Roman Garnett, editors, *Advances in Neural Information Processing Systems 29: Annual Conference on Neural Information Processing Systems 2016, December 5-10, 2016, Barcelona, Spain*, pages 2137–2145, 2016.
- [Hu and Wellman, 2003] Junling Hu and Michael P. Wellman. Nash q-learning for general-sum stochastic games. *J. Mach. Learn. Res.*, 4:1039–1069, 2003.
- [Iqbal and Sha, 2019] Shariq Iqbal and Fei Sha. Actor-attention-critic for multi-agent reinforcement learning. In Kamalika Chaudhuri and Ruslan Salakhutdinov, editors, *Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA*, volume 97 of *Proceedings of Machine Learning Research*, pages 2961–2970. PMLR, 2019.
- [Könönen, 2004] Ville Könönen. Asymmetric multiagent reinforcement learning. *Web Intell. Agent Syst.*, 2(2):105–121, 2004.
- [Kuba *et al.*, 2022] Jakub Grudzien Kuba, Ruiqing Chen, Muning Wen, Ying Wen, Fanglei Sun, Jun Wang, and Yaodong Yang. Trust region policy optimisation in multi-agent reinforcement learning. In *The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022*. OpenReview.net, 2022.
- [Lillicrap *et al.*, 2016] Timothy P. Lillicrap, Jonathan J. Hunt, Alexander Pritzel, Nicolas Heess, Tom Erez, Yuval Tassa, David Silver, and Daan Wierstra. Continuous control with deep reinforcement learning. In Yoshua Bengio and Yann LeCun, editors, *4th International Conference on Learning Representations, ICLR 2016, San Juan, Puerto Rico, May 2-4, 2016, Conference Track Proceedings*, 2016.
- [Littman and Stone, 2001] Michael L. Littman and Peter Stone. Leading best-response strategies in repeated games. In *In Seventeenth Annual International Joint Conference on Artificial Intelligence Workshop on Economic Agents, Models, and Mechanisms*. Citeseer, 2001.
- [Littman, 1994] Michael L. Littman. Markov games as a framework for multi-agent reinforcement learning. In William W. Cohen and Haym Hirsh, editors, *Machine Learning, Proceedings of the Eleventh International Conference, Rutgers University, New Brunswick, NJ, USA, July 10-13, 1994*, pages 157–163. Morgan Kaufmann, 1994.
- [Lowe *et al.*, 2017] Ryan Lowe, Yi Wu, Aviv Tamar, Jean Harb, Pieter Abbeel, and Igor Mordatch. Multi-agent actor-critic for mixed cooperative-competitive environments. In Isabelle Guyon, Ulrike von Luxburg, Samy Bengio, Hanna M. Wallach, Rob Fergus, S. V. N. Vishwanathan, and Roman Garnett, editors, *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pages 6379–6390, 2017.
- [Metz *et al.*, 2017] Luke Metz, Julian Ibarz, Navdeep Jaitly, and James Davidson. Discrete sequential prediction of continuous actions for deep RL. *arXiv preprint arXiv:1705.05035*, 2017.
- [Mnih *et al.*, 2013] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wierstra, and Martin Riedmiller. Playing atari with deep reinforcement learning. *arXiv preprint arXiv:1312.5602*, 2013.
- [Peng *et al.*, 2021] Bei Peng, Tabish Rashid, Christian Schröder de Witt, Pierre-Alexandre Kamienny, Philip H. S. Torr, Wendelin Boehmer, and Shimon Whiteson. FACMAC: factored multi-agent centralised policy gradients. In Marc’Aurelio Ranzato, Alina Beygelzimer, Yann N. Dauphin, Percy Liang, and Jennifer Wortman Vaughan, editors, *Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, December 6-14, 2021, virtual*, pages 12208–12221, 2021.
- [Rashid *et al.*, 2018] Tabish Rashid, Mikayel Samvelyan, Christian Schröder de Witt, Gregory Farquhar, Jakob N. Foerster, and Shimon Whiteson. QMIX: monotonic value function factorisation for deep multi-agent reinforcement learning. In Jennifer G. Dy and Andreas Krause, editors, *Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholmsmässan, Stockholm, Sweden, July 10-15, 2018*, volume 80 of *Proceedings of Machine Learning Research*, pages 4292–4301. PMLR, 2018.
- [Schulman *et al.*, 2017] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- [Tavakoli *et al.*, 2018] Arash Tavakoli, Fabio Pardo, and Petar Kormushev. Action branching architectures for deep reinforcement learning. In Sheila A. McIlraith and Kilian Q. Weinberger, editors, *Proceedings of the Thirty-*

- Second AAI Conference on Artificial Intelligence, (AAAI-18), the 30th innovative Applications of Artificial Intelligence (IAAI-18), and the 8th AAI Symposium on Educational Advances in Artificial Intelligence (EAAI-18), New Orleans, Louisiana, USA, February 2-7, 2018*, pages 4131–4138. AAAI Press, 2018.
- [Tian *et al.*, 2019] Zheng Tian, Ying Wen, Zhichen Gong, Faiz Punakkath, Shihao Zou, and Jun Wang. A regularized opponent model with maximum entropy objective. In Sarit Kraus, editor, *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI 2019, Macao, China, August 10-16, 2019*, pages 602–608. ijcai.org, 2019.
- [von Oswald *et al.*, 2020] Johannes von Oswald, Christian Henning, João Sacramento, and Benjamin F. Grewe. Continual learning with hypernetworks. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net, 2020.
- [Von Stackelberg, 2010] Heinrich Von Stackelberg. *Market structure and equilibrium*. Springer Science & Business Media, 2010.
- [Xu *et al.*, 2021] Zhiwei Xu, Bin Zhang, Yunpeng Bai, Dapeng Li, and Guoliang Fan. Learning to coordinate via multiple graph neural networks. In *International Conference on Neural Information Processing*, pages 52–63. Springer, 2021.
- [Yang *et al.*, 2018] Yaodong Yang, Rui Luo, Minne Li, Ming Zhou, Weinan Zhang, and Jun Wang. Mean field multi-agent reinforcement learning. In Jennifer G. Dy and Andreas Krause, editors, *Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholmsmässan, Stockholm, Sweden, July 10-15, 2018*, volume 80 of *Proceedings of Machine Learning Research*, pages 5567–5576. PMLR, 2018.
- [Yu *et al.*, 2021] Chao Yu, Akash Velu, Eugene Vinitzky, Yu Wang, Alexandre Bayen, and Yi Wu. The surprising effectiveness of ppo in cooperative, multi-agent games. *arXiv preprint arXiv:2103.01955*, 2021.
- [Zhang *et al.*, 2018] Kaiqing Zhang, Zhuoran Yang, Han Liu, Tong Zhang, and Tamer Basar. Fully decentralized multi-agent reinforcement learning with networked agents. In Jennifer G. Dy and Andreas Krause, editors, *Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholmsmässan, Stockholm, Sweden, July 10-15, 2018*, volume 80 of *Proceedings of Machine Learning Research*, pages 5867–5876. PMLR, 2018.
- [Zhang *et al.*, 2020] Haifeng Zhang, Weizhe Chen, Zeren Huang, Minne Li, Yaodong Yang, Weinan Zhang, and Jun Wang. Bi-level actor-critic for multi-agent coordination. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020, The Thirty-Second Innovative Applications of Artificial Intelligence Conference, IAAI 2020, The Tenth AAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2020, New York, NY, USA, February 7-12, 2020*, pages 7325–7332. AAAI Press, 2020.
- [Zhang *et al.*, 2023a] Bin Zhang, Yunpeng Bai, Zhiwei Xu, Dapeng Li, and Guoliang Fan. Efficient policy generation in multi-agent systems via hypergraph neural network. In *Neural Information Processing: 29th International Conference, ICONIP 2022, Virtual Event, November 22–26, 2022, Proceedings, Part II*, pages 219–230. Springer, 2023.
- [Zhang *et al.*, 2023b] Bin Zhang, Zhiwei Xu, Yiqun Chen, Dapeng Li, Yunpeng Bai, Guoliang Fan, and Lijuan Li. Multi-agent hyper-attention policy optimization. In *Neural Information Processing: 29th International Conference, ICONIP 2022, Virtual Event, November 22–26, 2022, Proceedings, Part I*, pages 76–87. Springer, 2023.