

Bi-Level Actor-Critic for Multi-Agent Coordination*

Haifeng Zhang,¹ Weizhe Chen,² Zeren Huang,² Minne Li,¹ Yaodong Yang,³

Weinan Zhang,² Jun Wang¹

¹University College London

²Shanghai Jiao Tong University

³Huawei R&D UK

{haifeng.zhang, junwang}@ucl.ac.uk, {chenweizhe, sjtu_hzr, wnzhang}@sjtu.edu.cn, minne.li@cs.ucl.ac.uk
yaodong.yang@huawei.com

Abstract

Coordination is one of the essential problems in multi-agent systems. Typically multi-agent reinforcement learning (MARL) methods treat agents equally and the goal is to solve the Markov game to an arbitrary Nash equilibrium (NE) when multiple equilibria exist, thus lacking a solution for NE selection. In this paper, we treat agents *unequally* and consider Stackelberg equilibrium as a potentially better convergence point than Nash equilibrium in terms of Pareto superiority, especially in cooperative environments. Under Markov games, we formally define the bi-level reinforcement learning problem in finding Stackelberg equilibrium. We propose a novel bi-level actor-critic learning method that allows agents to have different knowledge base (thus intelligent), while their actions still can be executed simultaneously and distributedly. The convergence proof is given, while the resulting learning algorithm is tested against the state of the arts. We found that the proposed bi-level actor-critic algorithm successfully converged to the Stackelberg equilibria in matrix games and find a asymmetric solution in a highway merge environment.

Introduction

In a multi-agent system, the effect of any agent’s action on the environment also depends on the actions taken by other agents and coordination is needed to consistently break ties between equally good actions or strategies (Bu et al. 2008). This problem is essential especially in the circumstances where the agents are not able to communicate. In game theory, coordination game is defined as the game with multiple Nash equilibria. Various criteria for Nash equilibrium selection were proposed in the game theory literature such as salience (Vanderschraaf 1995) and fairness (Rabin 1993), where the agents are assumed to know the game model before applying these criteria. For the environments where agents are not able to know the game model but can learn it through interactions with the environments, multi-agent reinforcement learning approaches were proposed to find a Nash equilibrium, including Nash Q-learning (Hu and Wellman 2003), MADDPG (Lowe et al. 2017) and the Mean-

Field Q-learning (Yang et al. 2018). These model-free approaches train the agents centrally to converge to a Nash equilibrium and then execute the agents distributively. However, these approaches can not guarantee a particular converged Nash equilibrium, which leads to uncertainty and sub-optimality.

To tackle this problem, we reconsider the coordination problem from an asymmetric angle. Although the original game model is symmetric that agents make decisions simultaneously, we are still able to define a priority of decision making for the agents in the training phase and keep simultaneous decision making in the execution phase. In this asymmetric game model, the Stackelberg equilibrium (SE) (Von Stackelberg 2010) is naturally set up as the learning objective rather than the Nash equilibrium (NE). The SE optimizes the leader’s policy given that the follower always plays the best-response policy. Despite its discrimination on the follower, we surprisingly find the SE is Pareto superior than the NE in a wide range of environments. For example, in the cooperative games, the SE is guaranteed to be Pareto optimal, whereas only one of the NEs achieves this point, as Table 1a shows. In a non-cooperative case shown in Table 1b, the SE is not included in the set of the NEs and is Pareto superior to any NE. In general, our empirical study shows the SE is likely to be Pareto superior to the average NE in games with high cooperation level.

For solving the SE, a wide variety of bi-level optimization methods were proposed (Dempe 2018). However, our problem setting differs from the traditional bi-level optimization problem in two aspects: 1) we consider a multi-state environment where the objective function is a summation of the sequential discounted rewards; 2) our game model is unknown and can only be learned through interactions. Actually, the traditional bi-level optimization problem can be regarded as a stateless model-based version of our problem. We formally define our problem as the bi-level reinforcement learning problem and propose a novel bi-level actor-critic algorithm to solve it. We train the actor of the follower and the critics of both agents centrally to find an SE and then execute the agents distributively. Our experiments in the small environments and a simulated highway merge environment demonstrate the efficiency of our algorithm, out-

*Weinan Zhang and Jun Wang are the corresponding authors of this paper.

Copyright © 2020, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

	X	Y	Z
A	15, 15	10, 10	0, 0
B	10, 10	10, 10	0, 0
C	0, 0	0, 0	30, 30

(a)

	X	Y	Z
A	20, 15	0, 0	0, 0
B	30, 0	10, 5	0, 0
C	0, 0	0, 0	5, 10

(b)

Table 1: Coordination games. (a) A cooperative game where A-X and C-Z are the NE. C-Z is also the SE and the Perato optimality point. (b) A non-cooperative game where A-X is the SE, B-Y and C-Z are the NE. The SE is Pareto superior to any NE in this game.

performing the state-of-the-art algorithms.

Preliminaries

Markov Game

In an n -player Markov game (Littman 1994) (or stochastic game) $\langle S, A_i, P, R_i, \gamma \rangle$, S denotes the state space, A_i denotes agent i 's action space and A denotes the joint action space, $P : S \times A \rightarrow \text{PD}(S)$ ¹ denotes the transition function, $R_i : S \times A_i \rightarrow \mathcal{R}$ denotes the reward function for agent i , and γ denotes the discount factor. Agents take actions simultaneously in each state following their policies $\pi_i : S \rightarrow \text{PD}(A_i)$. The objective of agent i is to maximize its discounted cumulative reward $\sum_t \gamma^t r_i^t$, where r_i^t is the reward agent i receives in time-step t . We also call Markov games as multi-agent reinforcement learning (MARL) problems.

Related MARL Solutions

For a Markov game, we have Bellman equations that characterize the optimal state-values $V_i^*(s)$ and action-values $Q_i^*(s, a)$:

$$Q_i^*(s, \vec{a}) = R(s, \vec{a}) + \gamma \sum_{s'} P(s, \vec{a}, s') V_i^*(s'). \quad (1)$$

The Minimax-Q method (Littman 1994) attempts to find the highest worst case values in zero-sum games whose state-values are computed as:

$$V_1^*(s) = \max_{\pi_1 \in \Pi_1} \min_{a_2 \in A_2} Q_1^*(s, \pi_1, a_2) = -V_2^*(s), \quad (2)$$

where $Q_1^*(s, \pi_1, a_2) = \sum_{a_1 \in A_1} \pi_1(s, a_1) Q_1^*(s, a_1, a_2)$ and Π_1 denotes the policy space of agent 1. Our bi-level method generalizes the minimax method from zero-sum games to general-sum games.

The Nash-Q method (Hu and Wellman 2003) attempts to find the Nash equilibrium whose state-values are computed as:

$$V_i^*(s) = \text{NASH}_i(Q_1^*(s), Q_2^*(s), \dots, Q_n^*(s)) \quad (3)$$

where $\text{NASH}_i(\vec{x}_1, \vec{x}_2, \dots, \vec{x}_n)$ denotes the i -th agent's payoff in a Nash equilibrium of the matrix game formed by

¹In this paper, $\text{PD}(X)$ denotes the probability distribution space over discrete set X .

$\vec{x}_1, \vec{x}_2, \dots, \vec{x}_n$. The Nash-Q method also generalizes the minimax method to general-sum games, but in a different direction compared to our method. Our bi-level method attempts to find Stackelberg Equilibrium rather than Nash equilibrium.

Bi-level Optimization

In this paper, we assume the agents in a two-player Markov game are asymmetric that the following agent observes the actions of the leading agent, which results in solving a bi-level optimization problem for a Markov game. The original bi-level optimization problem is formulated as below:

$$\begin{aligned} \min_{x_1} \quad & f_1(x_1, x_2) \\ \text{s.t.} \quad & g_1(x_1, x_2) \leq 0 \\ \min_{x_2} \quad & f_2(x_1, x_2) \\ \text{s.t.} \quad & g_2(x_1, x_2) \leq 0 \end{aligned} \quad (4)$$

where $f_i, i = 1, 2$ are the objective functions and $g_i, i = 1, 2$ are the constraint functions in each level.

The bi-level optimization problem can be equivalently described as a Stackelberg game where the upper-level optimizer is the leader and the lower-level optimizer is the follower and the solution of the bi-level optimization problem is the Stackelberg equilibrium.

Bi-level Reinforcement Learning

Problem Formulation

Connecting bi-level optimization with Markov game, x_i in Eq. (4) corresponds to agent i 's policy π_i , f_i corresponds to agent i 's cumulative reward and g_i corresponds to the constraint of action space. Assuming Agent 1 as the leader and Agent 2 as the follower, our problem is formulated as:

$$\begin{aligned} \max_{\pi_1} \quad & \mathbb{E}_{r_1^1, r_1^2 \dots \sim \pi_1, \pi_2} \sum_{t=1}^{\infty} \gamma^t r_1^t \\ \text{s.t.} \quad & \pi_1 \in \Pi_1 \\ \max_{\pi_2} \quad & \mathbb{E}_{r_2^1, r_2^2 \dots \sim \pi_1, \pi_2} \sum_{t=1}^{\infty} \gamma^t r_2^t \\ \text{s.t.} \quad & \pi_2 \in \Pi_2. \end{aligned} \quad (5)$$

We call this problem *bi-level reinforcement learning* (BiRL). BiRL can be viewed as a multi-state version of the Stackelberg game (Von Stackelberg 2010) and extends the standard bi-level optimization problem in two dimensions: 1) the objective is a summation of the discounted rewards in sequential states; 2) the form of the objective function is unknown and can only be learned through interactions with the environment in a model-free way. The standard bi-level optimization problem can be viewed as a stateless model-based version of our problem.

Stackelberg Equilibrium vs. Nash equilibrium

We formulate BiRL to tackle the coordination problem in MARL. In game theory, coordination game is defined as

a game with multiple Nash equilibria and the coordination problem can be regarded as a Nash equilibrium selection problem. In this paper, we consider Stackelberg equilibria as a potentially better solution for coordination games. Fig. 1 is an example demonstrating the difference between NE and SE in a Markov game. We figure out two advantages of the SE over the NE.

The first advantage of SE is the certainty or uniqueness. Multiple NEs may exist in a game while multiple SEs only exist under very strict conditions.² Existing MARL methods mainly converge to an arbitrary NE, which leads to uncertainty. Since the SE is unique in most games, it is more clear and stable to be a learning objective. By setting the SE as the objective, we actually attempt to avoid the coordination problem (or the NE selection problem) rather than solving it.

The second advantage of SE is the performance. The SE may achieve better payoff than the average NE in coordination environments in terms of Pareto superiority. An extreme example is the cooperative games. The SE always achieves the Pareto optimality point in a cooperative game while only the best NE achieves so, as we showed in Table 1b and Fig. 1. In other words, both the leader and the follower achieve higher payoffs in the SE than in the average NE. We intuitively believe that this result would still hold in games with less (but still high) cooperation levels.

For demonstrating our belief, we formally define the cooperation level of two-player Markov games as the correlation between the cumulative rewards of the agents:

$$CL = \frac{\sum_{\bar{\pi}} (V_1^{\bar{\pi}} - \bar{V}_1)(V_2^{\bar{\pi}} - \bar{V}_2)}{\sqrt{\sum_{\bar{\pi}} (V_1^{\bar{\pi}} - \bar{V}_1)^2 \sum_{\bar{\pi}} (V_2^{\bar{\pi}} - \bar{V}_2)^2}} \quad (6)$$

where $V_i^{\bar{\pi}}$ is short for $V_i^{\bar{\pi}}(s_0)$ denoting the average discounted cumulative reward for agent i from the start state s_0 following the joint policy $\bar{\pi}$ and $\bar{V}_i = \frac{1}{|\bar{\pi}|} \sum_{\bar{\pi}} V_i^{\bar{\pi}}$. Under this definition, the cooperation levels of a cooperative game and a zero-sum game are 1 and -1 respectively.

We empirically study the relationship between the cooperation level of a game and the average payoff achieved by the agents in the average NE and the SE. The results in Fig. 2 demonstrate that both the leader and the follower achieve higher payoff in the SE not only in fully cooperative games but also in the games with high cooperation level. We also find that the number of Nash equilibria in a game is positively correlated with the cooperation level, which suggests that the coordination problem is more likely to occur in games with high cooperation level. Hence, we argue that the SE may in general be Pareto superior to the average NE in coordination problems, especially in highly cooperative games.

Bi-level Tabular Q-learning

Similar to the minimax-Q and Nash-Q, we can define the bi-level Bellman equation by specifying the calculation method

²Multiple SEs only exist when given the policy of the leader, multiple policies of the follower achieve the maximal payoff, or given the best response of the follower, multiple policies of the leader achieve the maximal payoff.

A		10
B		20

Figure 1: A cooperative game example for BiRL. Agent A and B move simultaneously in the grid and receive the common reward only when they are both in the 10 or 20 square. Joint policies lead both agents to 10 or 20 are Nash equilibria, but only joint policies lead both agents to 20 are Stackelberg equilibria and are the solutions for BiRL.




Figure 2: SE vs. NE. We sample the payoffs for a 10×10 matrix game using the multivariate normal distribution with 0 for mean, 1 for variance and various parameters for covariance, which represents the cooperation level of the generated matrix game. When the covariance equals to 1, the game is fully cooperative. The result for each covariance parameter is averaged from 2000 independent trials. Similar results are found in matrix games with different sizes ranging from 5×5 to 100×100 , which is omitted due to the space limit.

for the optimal state-values in Eq. (1):

$$V_i^*(s) = \text{Stackelberg}_i(Q_1^*(s), Q_2^*(s)), \quad (7)$$

where $\text{Stackelberg}_i(\vec{x}_1, \vec{x}_2)$ denotes the i -th agent's payoff in the Stackelberg Equilibrium of the matrix game formed by \vec{x}_1, \vec{x}_2 .

Based on the bi-level Bellman equation, we are able to update the Q-values iteratively by Eq. (1) and (7). Formally, We have the update rules for Q_1 and Q_2 tables given a transaction $\langle s, a_1, a_2, s', r_1, r_2 \rangle$ with learning rate α_i :

$$a'_1 \leftarrow \underset{a_1}{\operatorname{argmax}} Q_1(s', a_1, \underset{a_2}{\operatorname{argmax}} Q_2(s', a_1, a_2)), \quad (8)$$

$$a'_2 \leftarrow \underset{a_2}{\operatorname{argmax}} Q_2(s', a'_1, a_2), \quad (9)$$

$$Q_1(s, a_1, a_2) \leftarrow (1 - \alpha_1)Q_1(s, a_1, a_2) + \alpha_1(r_1 + \gamma Q_1(s', a'_1, a'_2)), \quad (10)$$

$$Q_2(s, a_1, a_2) \leftarrow (1 - \alpha_2)Q_2(s, a_1, a_2) + \alpha_2(r_2 + \gamma Q_2(s', a'_1, a'_2)). \quad (11)$$

This tabular method was also studied in (Littman and Stone 2001) and (Könönen 2004). However, these works mainly focused on solving asymmetric problems while our motivation is to solve symmetric coordination problems using an asymmetric method.

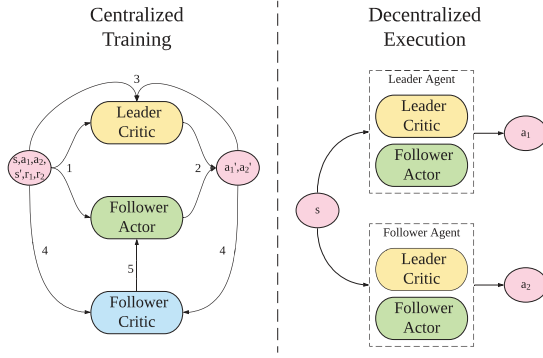


Figure 3: Structure of bi-level actor-critic. In the training phase, the joint action in the next state is computed firstly, then the three models are updated accordingly. In the execution phase, both agents have the leader critic and the follower actor models and perform Stackelberg equilibria independently.

Bi-level Actor-Critic

In Eq. (8), we need to enumerate the actions in both levels to select action a'_1 , which leads to $|A_1| \cdot |A_2|$ visits to the Q_2 table. When Q_2 is modeled by an approximation function, i.e. a neural network, the calculations of Eq. (8) could be time-consuming. Furthermore, if we extend the bi-level Q-learning methods to multi-level, the computation complexity of a'_1 would increase in exponential w.r.t. the number of level.

For solving this problem, we propose the bi-level actor-critic (Bi-AC) method which introduces an actor for the follower while keeping the leader as a Q-learner. Formally, let $\pi_2(s, a_1; \phi_2) \in \text{PD}(A_2)$ denote the policy model (or actor) of agent 2, which takes agent 1's action as its input in addition to the current state. We also model the two critics using approximation functions for both agents. We have the following update rules given a transaction $\langle s, a_1, a_2, s', r_1, r_2 \rangle$ with learning rate α_i, β :

$$a'_1 \leftarrow \underset{a_1}{\operatorname{argmax}} Q_1(s', a_1, \pi_2(s', a_1; \phi_2); \theta_1), \quad (12)$$

$$a'_2 \leftarrow \pi_2(s', a'_1; \phi_2), \quad (13)$$

$$\delta_i \leftarrow r_i + \gamma Q_i(s', \vec{a}'; \theta_i) - Q_i(s, \vec{a}; \theta_i), i = 1, 2, \quad (14)$$

$$\theta_i \leftarrow \theta_i + \alpha_i \delta_i \nabla_{\theta_i} Q_i(s, \vec{a}; \theta_i), i = 1, 2, \quad (15)$$

$$\phi_2 \leftarrow \phi_2 + \beta \nabla_{\phi_2} \log \pi_2(s, \vec{a}; \phi_2) Q_2(s, \vec{a}; \theta_2). \quad (16)$$

where $\pi'_2(s, a_1; \phi_2)$ is modeled by a Gumbel-Softmax estimator (Jang, Gu, and Poole 2016) which computes a'_2 directly.

For the environments with continuous action space, we model agent 2's policy using a deterministic model $\mu_2(s, a_1; \phi_2) \in A_2$ which is updated by the deterministic policy gradient method (Silver et al. 2014). The Q-network of agent 1 can be updated by the soft Q-learning (Haarnoja et al. 2017) method.

Bi-AC is a centralized-training-decentralized-execution method as Fig. 3 shows. The three models are trained together given off-policy episodes. In execution, the trained

leader critic model and follower actor model are both allocated to and executed by the leader and the follower. In such way, the two agents are able to achieve the Stackelberg equilibrium distributively.

In the partially observable environments, we train two additional actors $\pi'_i(o_i)$ as the approximators for each agent, where the corresponding critics are the trained leader and follower critics. The approximators allow the two agents play joint actions forming Stackelberg equilibria based on their own observations.

Bi-AC can be naturally extended to n -level actor-critic. In the case of continuous action space, we define deterministic policy models:

$$\mu_i(s, a_1, a_2, \dots, a_{i-1}; \theta_i), i = 1..n. \quad (17)$$

We also model the Q-functions for each agent as:

$$Q_i(s, a_1, a_2, \dots, a_n; \phi_i), i = 1..n. \quad (18)$$

In each training step, the actions in the next step are determined one by one from the upper-level agent to the lower-level agent and the models are updated accordingly:

$$a'_i \leftarrow \mu_i(s', a'_{1..i-1}; \theta_i), \quad (19)$$

$$\delta_i \leftarrow r_i + \gamma Q_i(s', \vec{a}'; \phi_i) - Q_i(s, \vec{a}; \phi_i), \quad (20)$$

$$\theta_i \leftarrow \theta_i + \alpha \delta_i \nabla_{\theta_i} Q_i(s, \vec{a}; \theta_i), \quad (21)$$

$$\phi_i \leftarrow \phi_i + \alpha \nabla_{\phi_i} \mu_i(s, a_{1..i-1}; \phi_i) \cdot \nabla_{a_i} Q_i(s, \vec{a}; \theta_i). \quad (22)$$

In practical problems, the models can be modified slightly to contain multiple agents in each level, where agents in the same level take actions simultaneously and the lower-level agents observe the actions of the upper-level agents.

Convergence and Limitation

Bi-AC will converge to the Stackelberg equilibrium under the following assumptions:

1. Every stage game $(Q_1^t(s), Q_2^t(s))$ for all t and s , has a global optimal point, and agents' payoffs in this point are selected by the actor function to update the critic functions with probability 1.
2. The critic learning rates α_t for the t -th transaction satisfies $\sum_{t=0}^{\infty} \alpha^t(s, \vec{a}) = \infty$, $\sum_{t=0}^{\infty} [\alpha^t(s, \vec{a})]^2 < \infty$ hold uniformly and with probability 1.
3. Every stage game $(Q_1^t(s), Q_2^t(s))$ for all t and s has a global optimal point, and the agents' payoffs in this point are selected by Bi-AC to update the critic functions with probability 1.

Given these assumptions, we can use the method of Cauchy sequence convergence to prove the convergence of our algorithm. The detailed proof is omitted due to the space limit.

The assumption 3 above is a strong assumption which may not be met even in the very start of the training process. However, as shown later in the experiment part, our algorithm does converge in many cases.

Our algorithm is based on the bi-level Bellman equation, which is a necessary but not sufficient condition of the solution of BiRL. Therefore, there may exist convergent points other than the solution of BiRL.

Related Work

In MARL, various approaches were proposed to tackle the coordination problem (Bu et al. 2008), especially for the cooperative environments. A general approach is applying the social convention which breaks ties by ordering of agents and actions (Boutilier 1996). Our method is compatible with social convention in the sense that we find the SE as the common knowledge of the agents about the game, based on which they can form social conventions. For cooperative games, the optimistic exploration were proposed (Claus and Boutilier 1998) for reaching optimal equilibrium. Lauer and Riedmiller (2000) used maximal estimation to update Q-value which ensures convergence to the optimal equilibrium given the reward function is deterministic. For the case of stochastic reward function, FMQ (Kapetanakis and Kudenko 2002), SOoN (Matignon, Laurent, and Le Fort-Piat 2009) and LMRL2 (Wei and Luke 2016) were proposed. These works share the idea of optimistic expectation on the cooperative opponent, which could not be extended to general-sum games.

Communication is an essential method to facilitate coordination. CommNet (Sukhbaatar, Fergus, and others 2016) used a centralized network architecture to exchange information between agents. BicNet (Peng et al. 2017) proposed the bidirectional RNNs to exchange information between agents in an actor-critic setting. MADDPG (Lowe et al. 2017) proposed the centralized-training-decentralized-execution scheme which is also adopted by our method. Other works in this area include DIAL (Foerster et al. 2016), COMA (Foerster et al. 2018) and MD-MADDPG (Pesce and Montana 2019).

Understanding other agents in multi-agent environments is of vital importance. Opponent modeling methods are helpful for coordination in many circumstances. ROMMEO (Tian et al. 2019) applied maximum entropy to model the opponent. Wen et al. (2019) introduced the idea of recursive reasoning between two agents. Opponent modeling methods adopt the decentralized training scheme while our training is centralized.

Our work tackle the coordination problem from a asymmetric angle. BiRL is an extension of the bi-level optimization (Dempe 2018) or the Stackelberg game (Von Stackelberg 2010). For solving the original bi-level problem, stateless model-based evolutionary algorithms were proposed, such as BLEAQ (Sinha, Malo, and Deb 2014). Extensively, a stateless model-free leader-follower problem was studied (Zhang and Lin 2012) where the objective functions are not visible. In the other dimension, BiMPC (Mintz et al. 2018) studied the multi-state model-based Stackelberg game under the linear-quadratic assumption. In our paper, we formulate the multi-state model-free problem of BiRL, which extends the original bi-level problem in both dimensions.

Our Bi-AC method contains critics which are iteratively updated by the Bellman equation. There are a series of MARL methods adopting the similar update scheme. Minimax-Q (Littman 1994) solved the two-player zero-sum Markov games. Afterwards, fiend-and-foe learning (Littman 2001), Nash-Q (Hu and Wellman 2003), CE-Q (Greenwald, Hall, and Serrano 2003), Coco-Q (Sodomka et al. 2013) and

AQL (Könönen 2004) were proposed successively. Among them, AQL updates the Q-value by solving the Stackelberg Equilibrium in each iteration, which can be regarded as the value-based version of Bi-AC. Compared to AQL, Bi-AC is able to work in multi-level or continuous action space environments. Another difference between AQL and our work is the motivation that we propose the Stackelberg Equilibrium as a potentially better solution for solving symmetric coordination problems while AQL focused on asymmetric problems. Other works applying Stackelberg Equilibrium to solve asymmetric problems include Bully (Littman and Stone 2001) and DeDOL (Wang et al. 2019).

Experiment

We performed the experiments in three coordination environments comparing Bi-AC with the state-of-the-art MARL algorithms.³

Algorithms

We compared Bi-AC with I-DQN, MADDPG, TD3, Hysteretic-DQN and Lenient-DQN, where I-DQN, Hysteretic-DQN and Lenient-DQN are decentralized training algorithms and MADDPG and TD3 are centralized training algorithms. In each experiment, we fully explored the actions for 1000 steps in the beginning and then applied the decaying ϵ -greedy method for explorations. A three layer fully connected neural network with ReLU (Nair and Hinton 2010) activation function was applied for the models in each algorithm to approximate the actor and critic functions. We trained the critic models with the replay buffer and the target network introduced in DQN (Mnih et al. 2015). We used the Gumbel-Softmax estimator (Jang, Gu, and Poole 2016) in the actor function when applicable.

Bi-AC We realized the Bi-AC algorithm where the two critics were modeled by DQN.

I-DQN We tested independent DQN which regards other agents as a part of the environments. I-DQN is not guaranteed to converge and if it converges it will converge to a Nash equilibrium.

MADDPG We tested MADDPG (Lowe et al. 2017) as the baseline of the centralized-training-decentralized-execution algorithm.

TD3 Twin Delayed Deep Deterministic policy gradient (TD3) (Fujimoto, van Hoof, and Meger 2018) enhances MADDPG by addressing function approximation error, which is the state-of-the-art MARL algorithm.

Hysteretic-DQN Hysteretic-DQN (Omidshafiei et al. 2017) is based on hysteretic Q-learning (Matignon, Laurent, and Fort-Piat 2007) where a smaller learning rate is used when an update would reduce a Q-value, leading to an optimistic update function which puts more weights on positive experiences.

³Our experiments are repeatable and the source code is provided in <https://github.com/laonahongchen/Bilevel-Optimization-in-Coordination-Game>.






Figure 4: Convergence of Bi-AC in Escape game



(a) Bi-AC



(b) MADDPG

Figure 5: Distribution of entry selected by different algorithms while training. Each color represent the possibility in the recent 100 steps to choose the corresponding entry.

Lenient-DQN Lenient-DQN (Palmer et al. 2018) combines DQNs with leniency, which increases the likelihood of convergence towards the global optimal solution within fully cooperative environments that require implicit coordination strategies.

Matrix Game

We tested two matrix games named Escape and Maintain whose payoff tables are shown in Table 1a and Table 1b respectively.

We designed the Escape game in Table 1a to show that our algorithm has the capability to converge to the SE which is better than the average NE in a cooperative game. We expected Bi-AC converged to the C-Z point while other algorithms converged to either the A-X or the C-Z point. Note that the higher-left 2×2 part of the matrix will lead to the sub-optimal A-X point for an NE learner, unless the joint action C-Z is explored enough for an escape.

We ran each algorithm for 100 times and the results are provided in Table 2. We see that Bi-AC achieved higher re-

	Leader	Follower	Optimality
Bi-AC	28.5	28.5	90%
I-DQN	20.55	20.55	37%
MADDPG	23.4	23.4	56%
TD3	15.3	15.3	2%
Hysteretic-DQN	13.25	13.25	0%
Lenient-DQN	13.4	13.4	0%

Table 2: Result of Escape game. The first two column means the average reward while the third column means the percentage of converging to C-Z, the global optimal point.

	Leader	Follower	Optimality
Bi-AC	20	15	100%
I-DQN	9.5	5.5	0%
MADDPG	12	4.5	0%
TD3	11.5	5	0%
Hysteretic-DQN	9.1	5.9	0%
Lenient-DQN	8.8	6.2	0%

Table 3: Result of Maintain game. The first two column means the average reward while the third column means the percentage of converging to A-X.

wards for both agents than all the baseline algorithms. Also, Bi-AC converged to the global optimal point C-Z in 90% trials leading the baseline algorithms by a large margin. Note that Bi-AC did not converge to the optimal point with 100% probability because of the usage of neural network function approximations.⁴ We also found MADDPG outperformed Hysteretic-DQN, Lenient-DQN and TD3 in this game. The reason may be that these three algorithms converges with lower bias which could also make it more difficult to converge to the isolated optimal point C-Z. The convergence curve of Bi-AC is provided in Fig. 4. From the very start of training, the Q-value was trained to the correct estimated value 30. And once the Q-value of the follower was trained to make C-Z higher than C-X and C-Y, the follower started to learn the best response toward action C of the leader. Then after the follower chose Z as the response with high probability, the probability for leader to choose action C started to increase, from about the 100th episode. The reason that the Q-value of the leader for C-Z converged before the follower’s Q-value is that the learning rate for Q-function of the leader is higher than the follower. This asymmetric parameter setting is because the asymmetric architecture in our algorithm. We also find that the probability of action C of leader agent did not converge to 1.0. It was because there was still exploration.

We designed the Maintain game in Table 1b to show that Bi-AC algorithm is able to achieve the SE which is Pareto superior to all the NEs. As discussed, the A-X point is the SE and Pareto optimality point but not an NE. We expected our algorithm to converge to the A-X point while other NE learners converge to B-Y or C-Z. Particularly, NE learners are not able to maintain in A-X because B-X is a better point for

⁴We performed the tabular version of Bi-AC to confirm this.

	Leader	Follower
Bi-AC	72.2%	25.2%
I-DQN	58.2%	41.3%
MADDPG	37.5%	60.2%
TD3	38.6%	55.1%

Table 4: Result of Traffic Merge. The Leader column shows the rate that the car goes first is from the main lane. The sum of the two columns is not 100% due to a small rate of crash.



Figure 6: An illustration of the Traffic Merge environment.

the row player. The result is provided in table 3. We can see that Bi-AC achieved higher rewards for both agents and converged to the A-X point in all the trials while the baselines never converged to the optimality. Fig. 5 shows the learning processes of the joint policies for Bi-AC and MADDPG. Bi-AC converged to the A-X point smoothly while MADDPG converged to the B-Y point as expected.

Highway Merge

We also designed a driving environment, where one car drives on the main lane and another car from the auxiliary lane wants to merge into the main lane. We used a slightly modified version of the Highway environment (Leurent 2018), in which an agent can observe the kinematics of the nearby agent including its position and velocity, and has a discrete action space of 5 actions including: LANE LEFT, LANE RIGHT, FASTER, SLOWER and IDLE, where IDLE means neither changing the current speed nor changing the LANE it is running on. The agents are rewarded for 50 if it passes first, rewarded for 10 if it passes second, and rewarded for -10 if two cars bump together as a penalty. An example overview of the environment is given in Fig. 6. Note that if the cars from auxiliary lane do not choose LANE LEFT before his own lane disappears, the environment will automatically drive the cars to perform LANE LEFT right before they step away from the auxiliary lane. For the deployment of our algorithm Bi-AC, we set the car in main lane to be the leader and the car in auxiliary lane to be the follower.

We ran our algorithm Bi-AC together with the baselines for 10 times using 10 random seeds to ensure all the algorithms face the same difficulty in training. The result is shown in Table 4 and Fig. 7. We found that in our setting, Bi-AC converged to a situation of going first with high probability of 70%. Note that the rest 30% probability came to the situation that the auxiliary lane car started with a much higher speed which made it impossible to wait for the leader. We also found that the other baselines failed to choose which car to go first because they do not have a preference so their estimation of the main lane car going first is about 50%, as shown in Table 4. In Fig. 7 we see that from the very beginning Bi-AC has almost the same possibility of three results: car crash, main lane car going first and auxiliary lane car go-

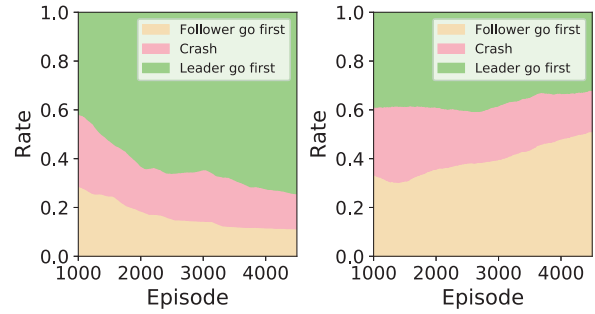


Figure 7: The training curve of highway merge problem. The left and right figures are the curve of Bi-AC and MADDPG respectively. The figures of other baselines are similar with MADDPG.

ing first. With the training, Bi-AC makes the possibility of main lane car going first improve steadily until convergence, which shows that Bi-AC can solve real world problems like traffic merge on highway.

Conclusion

In this paper, we consider Stackelberg equilibrium as a potentially better learning objective than Nash equilibrium in coordination environments due to its certainty and optimality. We formally define the bi-level reinforcement learning problem as the multi-state model-free Stackelberg equilibrium learning problem and empirically study the relationship between the cooperation level and the superiority of Stackelberg equilibrium to Nash equilibrium. We then propose a novel bi-level actor-critic algorithm which is trained centrally and executed decentrally. Our experiments on matrix games and a highway merge environment demonstrate the effectiveness of our algorithm to find the Stackelberg solutions.

Acknowledgments

This work is supported by "New Generation of AI 2030" Major Project 2018AAA0100900 and NSFC 61702327, 61632017.

References

- Boutilier, C. 1996. Planning, learning and coordination in multiagent decision processes. In *6th TARK*, 195–210. Morgan Kaufmann Publishers Inc.
- Bu, L.; Babu, R.; De Schutter, B.; et al. 2008. A comprehensive survey of multiagent reinforcement learning. *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)* 38(2):156–172.
- Claus, C., and Boutilier, C. 1998. The dynamics of reinforcement learning in cooperative multiagent systems. *AAAI/IAAI* 1998:746–752.
- Dempe, S. 2018. *Bilevel optimization: theory, algorithms and applications*. TU Bergakademie Freiberg, Fakultät für Mathematik und Informatik.

- Foerster, J.; Assael, I. A.; de Freitas, N.; and Whiteson, S. 2016. Learning to communicate with deep multi-agent reinforcement learning. In *NIPS*, 2137–2145.
- Foerster, J. N.; Farquhar, G.; Afouras, T.; Nardelli, N.; and Whiteson, S. 2018. Counterfactual multi-agent policy gradients. In *AAAI*.
- Fujimoto, S.; van Hoof, H.; and Meger, D. 2018. Addressing function approximation error in actor-critic methods. *arXiv preprint arXiv:1802.09477*.
- Greenwald, A.; Hall, K.; and Serrano, R. 2003. Correlated q-learning. In *ICML*, volume 3, 242–249.
- Haarnoja, T.; Tang, H.; Abbeel, P.; and Levine, S. 2017. Reinforcement learning with deep energy-based policies. In *34th ICML-Volume 70*, 1352–1361. JMLR. org.
- Hu, J., and Wellman, M. P. 2003. Nash q-learning for general-sum stochastic games. *JMLR* 4(Nov):1039–1069.
- Jang, E.; Gu, S.; and Poole, B. 2016. Categorical reparameterization with gumbel-softmax. *arXiv:1611.01144*.
- Kapetanakis, S., and Kudenko, D. 2002. Improving on the reinforcement learning of coordination in cooperative multi-agent systems. In *2nd AAMAS*.
- Könönen, V. 2004. Asymmetric multiagent reinforcement learning. *Web Intelligence and Agent Systems: An international journal* 2(2):105–121.
- Lauer, M., and Riedmiller, M. 2000. An algorithm for distributed reinforcement learning in cooperative multi-agent systems. In *ICML*. Citeseer.
- Leurent, E. 2018. An environment for autonomous driving decision-making. <https://github.com/eleurent/highway-env>.
- Littman, M. L., and Stone, P. 2001. Leading best-response strategies in repeated games. In *7th IJCAI*. Citeseer.
- Littman, M. L. 1994. Markov games as a framework for multi-agent reinforcement learning. In *Machine learning proceedings 1994*. Elsevier. 157–163.
- Littman, M. L. 2001. Friend-or-foe q-learning in general-sum games. In *ICML*, volume 1, 322–328.
- Lowe, R.; Wu, Y.; Tamar, A.; Harb, J.; Abbeel, O. P.; and Mordatch, I. 2017. Multi-agent actor-critic for mixed cooperative-competitive environments. In *NIPS*, 6379–6390.
- Matignon, L.; Laurent, G. J.; and Fort-Piat, N. L. 2007. Hysteretic q-learning : an algorithm for decentralized reinforcement learning in cooperative multi-agent teams. In *2007 IEEE/RSJ ICIRS*, 64–69.
- Matignon, L.; Laurent, G. J.; and Le Fort-Piat, N. 2009. Coordination of independent learners in cooperative markov games.
- Mintz, Y.; Cabrera, J. A.; Pedrasa, J. R.; and Aswani, A. 2018. Control synthesis for bilevel linear model predictive control. In *ACC 2018*, 2338–2343. IEEE.
- Mnih, V.; Kavukcuoglu, K.; Silver, D.; Rusu, A. A.; Veness, J.; Bellemare, M. G.; Graves, A.; Riedmiller, M.; Fidjeland, A. K.; Ostrovski, G.; et al. 2015. Human-level control through deep reinforcement learning. *Nature* 518(7540):529.
- Nair, V., and Hinton, G. E. 2010. Rectified linear units improve restricted boltzmann machines. In *ICML-10*, 807–814.
- Omidshafiei, S.; Pazis, J.; Amato, C.; How, J. P.; and Vian, J. 2017. Deep decentralized multi-task multi-agent reinforcement learning under partial observability. In *ICML*, 2681–2690.
- Palmer, G.; Tuyls, K.; Bloembergen, D.; and Savani, R. 2018. Lenient multi-agent deep reinforcement learning. In *AAMAS*, 443–451.
- Peng, P.; Wen, Y.; Yang, Y.; Yuan, Q.; Tang, Z.; Long, H.; and Wang, J. 2017. Multiagent bidirectionally-coordinated nets: Emergence of human-level coordination in learning to play starcraft combat games. *arXiv preprint arXiv:1703.10069*.
- Pesce, E., and Montana, G. 2019. Improving coordination in multi-agent deep reinforcement learning through memory-driven communication. *arXiv preprint arXiv:1901.03887*.
- Rabin, M. 1993. Incorporating fairness into game theory and economics. *The American economic review* 1281–1302.
- Silver, D.; Lever, G.; Heess, N.; Degris, T.; Wierstra, D.; and Riedmiller, M. 2014. Deterministic policy gradient algorithms.
- Sinha, A.; Malo, P.; and Deb, K. 2014. An improved bilevel evolutionary algorithm based on quadratic approximations. In *2014 CEC*, 1870–1877. IEEE.
- Sodomka, E.; Hilliard, E.; Littman, M.; and Greenwald, A. 2013. Coco-q: Learning in stochastic games with side payments. In *ICML*, 1471–1479.
- Sukhbaatar, S.; Fergus, R.; et al. 2016. Learning multiagent communication with backpropagation. In *NIPS*, 2244–2252.
- Tian, Z.; Wen, Y.; Gong, Z.; Punakkath, F.; Zou, S.; and Wang, J. 2019. A regularized opponent model with maximum entropy objective. *arXiv preprint arXiv:1905.08087*.
- Vanderschraaf, P. 1995. Convention as correlated equilibrium. *Erkenntnis* 42(1):65–87.
- Von Stackelberg, H. 2010. *Market structure and equilibrium*. Springer Science & Business Media.
- Wang, Y.; Shi, Z. R.; Yu, L.; Wu, Y.; Singh, R.; Joppa, L.; and Fang, F. 2019. Deep reinforcement learning for green security games with real-time information. In *AAAI*, volume 33, 1401–1408.
- Wei, E., and Luke, S. 2016. Lenient learning in independent-learner stochastic cooperative games. *JMLR* 17(1):2914–2955.
- Wen, Y.; Yang, Y.; Luo, R.; Wang, J.; and Pan, W. 2019. Probabilistic recursive reasoning for multi-agent reinforcement learning. *arXiv preprint arXiv:1901.09207*.
- Yang, Y.; Luo, R.; Li, M.; Zhou, M.; Zhang, W.; and Wang, J. 2018. Mean field multi-agent reinforcement learning. *arXiv preprint arXiv:1802.05438*.
- Zhang, D., and Lin, G.-H. 2012. Bilevel direct search method for leader-follower equilibrium problems and applications.