

# Learning from Multiple Independent Advisors in Multi-agent Reinforcement Learning

Sriram Ganapathi Subramanian  
Vector Institute, Toronto, Canada  
University of Waterloo, Waterloo, Canada  
sriram.subramanian@vectorinstitute.ai

Kate Larson  
University of Waterloo  
Waterloo, Canada  
kate.larson@uwaterloo.ca

Matthew E. Taylor  
University of Alberta, Edmonton, Canada  
Alberta Machine Intelligence Institute, Edmonton, Canada  
matthew.e.taylor@ualberta.ca

Mark Crowley  
University of Waterloo  
Waterloo, Canada  
mcrowley@uwaterloo.ca

## ABSTRACT

Multi-agent reinforcement learning typically suffers from the problem of sample inefficiency, where learning suitable policies involves the use of many data samples. Learning from external demonstrators is a possible solution that mitigates this problem. However, most prior approaches in this area assume the presence of a single demonstrator. Leveraging multiple knowledge sources (i.e., *advisors*) with expertise in distinct aspects of the environment could substantially speed up learning in complex environments. This paper considers the problem of simultaneously learning from multiple independent advisors in multi-agent reinforcement learning. The approach leverages a two-level  $Q$ -learning architecture, and extends this framework from single-agent to multi-agent settings. We provide principled algorithms that incorporate a set of advisors by both evaluating the advisors at each state and subsequently using the advisors to guide action selection. We also provide theoretical convergence and sample complexity guarantees. Experimentally, we validate our approach in three different test-beds and show that our algorithms give better performances than baselines, can effectively integrate the combined expertise of different advisors, and learn to ignore bad advice.

## KEYWORDS

Multi-agent systems, Multi-agent reinforcement learning, Learning from action advising, Reinforcement learning, Sample efficiency

### ACM Reference Format:

Sriram Ganapathi Subramanian, Matthew E. Taylor, Kate Larson, and Mark Crowley. 2023. Learning from Multiple Independent Advisors in Multi-agent Reinforcement Learning. In *Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023)*, London, United Kingdom, May 29 – June 2, 2023, IFAAMAS, 27 pages.

## 1 INTRODUCTION

Reinforcement learning (RL) has been successful in obtaining super-human performances in a wide range of challenges such as Atari games [26], Go [34], and simple robotic tasks like opening doors and learning visuomotor policies [20]. However, it has not been straightforward to replicate these successes in complex real-world

problems. One reason is that these problems often have a multi-agent structure, where more than one learning agent participates at the same time, resulting in complicated dynamics. Despite research advances in multi-agent reinforcement learning (MARL) [11], poor sample efficiency in existing algorithms is one issue that still causes significant hurdles in applying MARL to complex problems [32].

Using external sources of knowledge that help in accelerating MARL training is one solution [2], which has extensive support in literature [32]. However, most prior work include two limiting assumptions. First, all demonstrations need to come from a single demonstrator [4]. In complex MARL environments, since agents learn policies that meet the twin goals of responding to changing opponent(s) and environments [22], a learner can likely benefit from multiple knowledge sources that have expertise in different parts of the environment or different aspects of the task. Second, all demonstrations are near-optimal (i.e., from an “expert”) [28]. In practice, these knowledge sources are typically sub-optimal, and we broadly refer to them as *advisors* (to differentiate from experts).

In this paper, we provide an approach that simultaneously leverages multiple different (sub-optimal) advisors for MARL training. Since the advisors may provide conflicting advice in different states, an algorithm needs to resolve such conflicts to take advantage of all the advisors effectively. We propose a two-level learning architecture and formulate a  $Q$ -learning algorithm, improving upon the work of Li et al. [21] in single-agent RL. This architecture uses one level to evaluate advisors and the other learns values for actions. Further, we extend our approach to an actor-critic variant that applies to the centralized training and decentralized execution (CTDE) setting [24]. Since RL is a fixed point iterative method [39], we provide convergence and finite-time analysis of our algorithm. Finally, we experimentally study our approach in three different multi-agent test-beds, in relation to standard baselines.

Since we relax the two limiting assumptions regarding learning from demonstrators in MARL, our hope is that this approach will spur successes in real-world applications, such as autonomous driving [10] and fighting wildfires [15], where MARL methods could use existing (sub-optimal) solutions as advisors to accelerate training.

## 2 RELATED WORK

This work is most related to the approach of *reinforcement learning from expert demonstrations* (RLED) [28]. A well-known RLED technique is *deep  $Q$ -learning from demonstrations* (DQfD) [12], which

combines a temporal difference (TD) loss, an L2 regularization loss, and a classification loss that encourages actions to be close to that of the demonstrator. Another method, *normalized actor-critic* (NAC) [16], drops the classification loss and is more robust under imperfect demonstrations. However, NAC is prone to weaker performances than DQfD under good demonstrations due to the absence of classification loss. A different approach, *human agent transfer* (HAT) [41], extracts information from limited demonstrations using a classifier, while *confidence-based human-agent transfer* (CHAT) [46] improves HAT by using a confidence measurement to safeguard against sub-optimal demonstrations. A related approach is the teacher-student framework [43], where a pretrained policy (teacher) can be used to provide limited advice to a learning agent (student). Subsequent works expand this framework towards interactive learning [1], however, almost all works in this area assume a moderate level expertise for the teacher. Moreover, these are all independent methods primarily suited for single-agent environments, and may not be directly applicable in MARL context.

Furthermore, external knowledge sources have also been used in MARL [32], where prior works often assume near optimal experts [29, 48] or are only applicable to restrictive settings, such as fully cooperative or zero-sum competitive games [27, 33, 34, 45, 51]. Leno et al. [33] introduced a framework where an agent can learn from its peers in a shared learning environment, in addition to learning from the environmental rewards. Here the peers can be sub-optimal, however this work only applies to cooperative environments. Other works have provided a cooperative teaching framework for hierarchical learning [18, 49]. For multi-agent general-sum environments, *advising multiple intelligent reinforcement agents - decision making* (ADMIRAL-DM) [37] is a  $Q$ -learning approach that incorporates real-time information from a single online sub-optimal advisor.

One limitation of many prior works is the assumption of a single source of demonstration. In MARL, it may be possible to obtain advisors from different sources of knowledge that provide conflicting advice. For single-agent settings, Li et al. [21] provides the two-level  $Q$ -learning (TLQL) algorithm that incorporates multiple advisors in RL. The TLQL maintains two  $Q$ -networks, where the first  $Q$ -network (high-level) keeps track of each advisor's performance and the second  $Q$ -network (low-level) learns the quality of each action. We improve upon TLQL and make it applicable to MARL settings.

### 3 BACKGROUND

**Stochastic Games:** A  $N$ -player stochastic game is represented by a tuple  $\langle S, A^1, \dots, A^N, r^1, \dots, r^N, P, \gamma \rangle$ , where  $S$  is the state space,  $A^j$  is the action space of the agent  $j \in \{1, \dots, N\}$ , and  $r^j : S \times A^1 \times \dots \times A^N \rightarrow \mathcal{R}$  is the reward function of  $j$ . Also,  $P : S \times A^1 \times \dots \times A^N \rightarrow \Omega(S)$  is the transition probability that determines the next state given the current state and the joint action of all agents, where  $\Omega$  is a probability distribution. Finally,  $\gamma \in [0, 1)$  is the discount factor. At each time  $t$ , all agents observe the global state  $s$  and take a local action  $a^j$  [31]. The joint action  $\mathbf{a} = \{a^1, \dots, a^N\}$  determines the immediate reward  $r^j$  for  $j$  and the next state of the system  $s'$ . Each agent learns a suitable policy that gives the best responses to its opponent(s). The policy is denoted by  $\pi^j : S \rightarrow \Omega(A^j)$ . Let  $\pi \triangleq (\pi^1, \dots, \pi^N)$  be the joint policy of all agents. At a state  $s$ , the value function of  $j$  under the joint policy  $\pi$  is

$v_\pi^j(s) = \sum_{t=0}^{\infty} \gamma^t \mathbb{E}_{\pi, P}[r_t^j | s_0 = s, \pi]$ . This represents the expected discounted future reward of  $j$ , when all agents follow the policy  $\pi$  from the state  $s$ . Related to the value function, is the action-value function or the  $Q$ -function. The  $Q$ -function of agent  $j$ , under the policy  $\pi$ , is given by,  $Q_\pi^j(s, \mathbf{a}) = r^j(s, \mathbf{a}) + \gamma \mathbb{E}_{s' \sim P}[v_\pi^j(s')]$ .

The setting we consider is general-sum stochastic games, where the reward functions of the different agents can be related in any arbitrary fashion. In this setting, the *Nash equilibrium* is typically considered as the solution concept [13], where the joint policy  $\pi_* = [\pi_*^1, \dots, \pi_*^N]$  for all  $s \in S$  and all  $j$  satisfies  $v^j(s; \pi_*^j, \pi_*^{-j}) \geq v^j(s; \pi^j, \pi_*^{-j})$ . Here,  $\pi_*^{-j} \triangleq [\pi_*^1, \dots, \pi_*^{j-1}, \pi_*^{j+1}, \dots, \pi_*^N]$  represents the joint policy of all agents except  $j$ . In a Nash equilibrium, each agent plays the best response to the other agents and any deviation from this response is guaranteed to be worse off. Further, Hu and Wellman [13] proved that the  $Q$ -updates of an agent  $j$ , using the Nash payoff at each stage eventually converges to its Nash  $Q$  value ( $Q_*^j$ ), which is the action-value obtained by the agent  $j$  when all agents follow the joint Nash equilibrium policy for infinite periods.

**Two-level  $Q$ -learning:** The TLQL algorithm [21] enables single-agent learning under the simultaneous presence of multiple advisors providing conflicting demonstrations. Here, the challenge is to determine which advisor to trust in a given state. In this regard, the TLQL contains two  $Q$ -tables, a high-level  $Q$ -table (abbreviated as high- $Q$ ) and a low-level  $Q$ -table (abbreviated as low- $Q$ ). The high- $Q$  stores the value of the  $\langle s, ad \rangle$  pair, where  $ad \in AD$  represents an advisor (with  $AD$  representing the set of all advisors). The high- $Q$  also stores the value of following the RL policy in addition to each advisor. The low- $Q$  maintains the value of each state-action pair.

At each time step, the agent observes the state and selects an advisor (or the RL policy) from the high- $Q$  using the  $\epsilon$ -greedy strategy. If the high- $Q$  returns an advisor, then the advisor's recommended action is performed. If the RL policy is returned, then an action is executed from the low- $Q$  based on the  $\epsilon$ -greedy strategy [38]. The low- $Q$  is updated using the vanilla  $Q$ -learning Bellman update [47]. Subsequently, the high- $Q$  is updated using a **synchronization** step. In this step, when an advisor's action is performed, the value of the advisor in the high- $Q$  is simply assigned the value of that action from the low- $Q$ . Finally, the high- $Q$  of the RL policy is updated using the relation  $highQ(s, RL) = \max_a lowQ(s, a)$ . This synchronization update of high- $Q$  preserves the convergence guarantees, due to the policy improvement guarantee in single-agent  $Q$ -learning [38].

There are two important limitations of TLQL. First, the high- $Q$  that represents the value of the advisors also depends on the RL policy through the synchronization step. This  $Q$  value represents the value of taking the action suggested by the advisor at the current state and then following the RL policy from the next state onward. This definition is problematic since at the beginning of training, the RL policy is sub-optimal, and the objective is to accelerate learning by relying on external advisors and avoid using the RL policy at all. As advisors are evaluated at each state using the RL policy, it is likely that the most effective advisor among the set of advisors is not being followed until the RL policy improves. At this stage, it might be possible to simply follow the RL policy itself, defeating the purpose of learning from advisors. Second, the advisors have not been evaluated at the beginning of learning. Hence, it is impossible to find the most suitable advisor to follow,

from the available advisors. While TLQL simply follows an  $\epsilon$ -greedy exploration strategy, this approach could take many data samples to figure out the right advisor. We address both these limitations.

#### 4 TWO-LEVEL ARCHITECTURE IN MARL

We consider a general-sum stochastic game, where there are a set of agents that are learning a policy with an objective of providing a best response to the other agents as well as the environment. Each agent  $j$  can access a finite set of (possibly sub-optimal) independent advisors  $AD^j$ . We use  $ad^j$  to represent an advisor of  $j$ , where  $ad^j \in \{ad_1^j, \dots, ad_{|AD^j|}^j\}$ . Each advisor  $ad^j$  can be queried by  $j$  to obtain an action recommendation at each state of the stochastic game. These online advisors provide real-time action advice to the agent, which helps in learning to dynamically adapt to opponents. We consider a centralized training setting and assume 1) the advisors are fixed and do not learn, 2) the communication between agents and their advisors is free, 3) there is no communication directly between learning agents, 4) the environment is fully observable (i.e., an agent can observe the global state, all actions, and all rewards), and 5) the state and action spaces are discrete. Though we require these assumptions for theoretical guarantees, we will show that it is possible to relax a number of these assumptions in practice.

To make TLQL applicable to multi-agent settings, we parameterize both the  $Q$ -functions with the joint actions, as is common in practice [22]. Also, we do not maintain the RL policy in the high- $Q$  table and do not perform a synchronization step. These steps are no longer needed to preserve the convergence results in multi-agent settings, since we do not have a policy improvement guarantee (unlike in single-agent settings) [40]. Instead, we choose to use the probabilistic policy reuse (PPR) technique [6], where a hyperparameter ( $\epsilon' \in [0, 1]$ ) decides the probability of following any advisor(s) (i.e., using the high- $Q$ ) or the agents' own policy (i.e., using the low- $Q$ ) for action selection, at each time step during training. This hyperparameter starts with a high value (maximum dependence on the available advisor(s)) at the beginning of training and is decayed (linearly) over time. After some finite time step during training, the value of this hyperparameter goes to 0 (no further dependence on any advisor(s)) and the agent only uses its low- $Q$  (own policy) for action selection. This helps in two ways: 1) in the time limit ( $t \rightarrow \infty$ ), a learning agent has the possibility of recovering from poor advising (by learning from the environment), and 2) eventually the trained agent can be independently deployed (with no requirement of having access to any advisor(s)).

The general structure of our proposed *Multi-Agent Two-Level Q-Learning* (MA-TLQL) algorithm is given in Figure 1. Since we are in a fully observable setting, like [13], we specify that each agent maintains copies of the  $Q$ -tables of other agents from which it can obtain the joint actions of other agents for the current state. If such copies cannot be maintained, agents could use the previously observed actions of other agents for the joint action as done in prior works [37, 50]. We use the two-level architecture, where each agent will maintain a high- $Q$  as well as a low- $Q$ . The high- $Q$  provides a value for the  $\langle s, \mathbf{a}^{-j}, ad^j \rangle$  tuples, where  $\mathbf{a}^{-j} = \{a^1, \dots, a^{j-1}, a^{j+1}, \dots, a^N\}$  is the joint action of all agents except the agent  $j$ . This high- $Q$  is a value estimate for the advisor  $ad^j$  as estimated by the agent  $j$  at the state  $s$  and joint action  $\mathbf{a}^{-j}$ . The

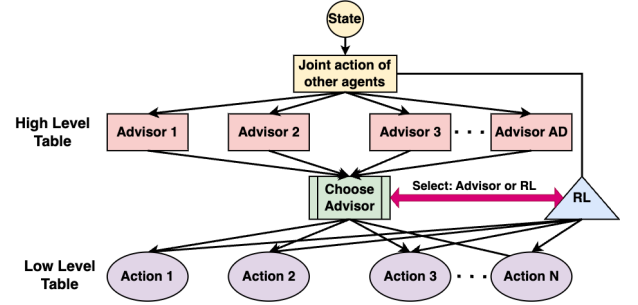


Figure 1: Structure of MA-TLQL, for a representative agent having access to a set of AD advisor(s)

high- $Q$  estimates are updated with an evaluation update given by

$$\begin{aligned} highQ_{t+1}^j(s, \mathbf{a}^{-j}, ad^j) &= highQ_t^j(s, \mathbf{a}^{-j}, ad^j) \\ &+ \alpha \left( (r_t^j + \gamma highQ_t^j(s', \mathbf{a}^{-j}, ad^j) - highQ_t^j(s, \mathbf{a}^{-j}, ad^j)) \right), \end{aligned} \quad (1)$$

where  $s$  and  $s'$  are the states at  $t$  and  $t + 1$ , and  $\alpha$  is the learning rate. Also,  $\mathbf{a}^{-j}$  and  $\mathbf{a}'^{-j}$  are joint actions at  $s$  and  $s'$ , respectively.

As described previously, a hyperparameter is used to decide between choosing to follow an advisor or the RL policy. If the agent follows an advisor, the high- $Q$  is used to select an advisor using an ensemble selection technique. Let us denote,  $Q^j = \{highQ^j(s, \mathbf{a}^{-j}, ad_1^j), \dots, highQ^j(s, \mathbf{a}^{-j}, ad_m^j)\}$ , to represent the high- $Q$  estimates of a set of  $M$  advisors (with  $|M| = m$ ) advising the same action  $a^j$  to an agent  $j$ . Here,  $ad_i^j$  represents an advisor  $i \in \{1, \dots, m\}$  of  $j$ . Then the value of vote for action  $a^j$ , at the state  $s$  and the joint action  $\mathbf{a}^{-j}$ , denoted by  $\mathcal{V}^j(s, \mathbf{a}^{-j}, a^j)$ , is calculated as

$$\begin{aligned} \mathcal{V}^j(s, \mathbf{a}^{-j}, a^j) &= \max_i Q^j \\ &+ \sum_{i=1, i \neq \arg \max_i highQ^j(s, \mathbf{a}^{-j}, ad_i^j)}^m \frac{1}{\mu(s)} highQ^j(s, \mathbf{a}^{-j}, ad_i^j). \end{aligned} \quad (2)$$

Here,  $\mu(s)$  is the number of times the agent has visited the state  $s$ . In Eq. 2, if an action is recommended by more than one advisor, the value of its vote is a weighted sum of all high- $Q$  estimates of advisors recommending that action. Each high- $Q$  estimate (except the best high- $Q$  estimate) is weighted by the reciprocal of the number of times the respective state is visited. In this way, when a state is visited many times, the advisor with the best high- $Q$  estimate is likely to be followed (wisdom of individual). When a state is visited only a few times, then the action suggested by a majority of advisors is likely to be selected (wisdom of crowd). From Eq. 2, if an action  $a^j$  is recommended by only one advisor, then the value of vote for  $a^j$  will be equal to the high- $Q$  estimate of that advisor. After the value of votes for all actions are calculated, the action with the maximum value of vote is executed, and the high- $Q$  estimate of the advisor recommending this action is updated by the agent  $j$  using Eq. 1.

If the agent decides to use its RL policy, it uses its low- $Q$ , which contains a value for the  $\langle s, \mathbf{a}^{-j}, a^j \rangle$  tuples (value for each action). At each step, the low- $Q$  is updated using a control update as follows:

$$\begin{aligned} lowQ_{t+1}^j(s, \mathbf{a}^{-j}, a^j) &= lowQ_t^j(s, \mathbf{a}^{-j}, a^j) \\ &+ \alpha (r_t^j + \gamma \max_{a'^j} lowQ_t^j(s', \mathbf{a}^{-j}, a'^j) - lowQ_t^j(s, \mathbf{a}^{-j}, a^j)). \end{aligned} \quad (3)$$

Now we describe how MA-TLQL addresses the two limitations of TLQL. The first was the dependence of high- $Q$  on the RL policy

in TLQL. Note, the high- $Q$  in MA-TLQL maintains the  $Q$  values of the advisor themselves, i.e., the value of following the advisor's policy from the current state onward (see Eq. 1). Thus, the coupling between the advisor values and the RL policy is removed (no synchronization). The second was the difficulty in picking the right advisor in TLQL. MA-TLQL uses an ensemble technique to choose the advisor during the early stages of learning. In later stages, it switches to following the best advisor according to the high- $Q$  estimates, which addresses this limitation of TLQL. In Appendix J, we present a toy example that illustrates the limitations of TLQL.

We provide the complete pseudocode for a tabular implementation of the MA-TLQL algorithm in Appendix A (Algorithm 1). Further, we extend this approach to large state-action environments using a neural network based implementation (Algorithm 2), which uses a target network and a replay buffer, as in the Deep  $Q$ -learning (DQN) algorithm [26]. We also provide an actor-critic implementation (Algorithm 3) which is suitable for CTDE [24]. We will refer to this algorithm as *multi-agent two-level actor-critic* (MA-TLAC). In this CTDE method, agents can obtain global information (including actions and rewards of other agents) during training, however, the agents only require access to its local observation during execution. This makes our method applicable to partially observable environments as in Lowe et al. [24]. MA-TLAC applies to continuous state space environments as well (refer to Appendix A for more details).

## 5 THEORETICAL RESULTS

We present a convergence guarantee for tabular MA-TLQL and characterize the convergence rate. For these results, we build on some prior works that provide several fundamental results on the nature of stochastic iterative functions [3, 5]. We apply these to MA-TLQL in general-sum stochastic games using three assumptions from Hu and Wellman [13], where the first two are standard [39].

**ASSUMPTION 1.** Every  $s \in \mathcal{S}$  and  $a^j \in A^j$ , for every agent  $j$  are visited infinitely often, and the reward function ( $\forall j$ ) stays bounded.

**ASSUMPTION 2.** For all  $s, t$ , and  $a$ ,  $0 \leq \alpha_t(s, a) < 1$ ,  $\sum_{t=0}^{\infty} \alpha_t(s, a) = \infty$ ,  $\sum_{t=0}^{\infty} [\alpha_t(s, a)]^2 < \infty$ .

**ASSUMPTION 3.** The Nash equilibrium is a global optimum or saddle point in every stage game of the stochastic game.

The third assumption is a restriction on the nature of the stochastic game. Several prior works note that this assumption is restrictive but needed to theoretically prove the convergence of  $Q$ -learning methods in general-sum stochastic games with two or more agents. In practice, however, it is still possible to observe convergence of  $Q$ -learning methods when this assumption is violated [13, 37, 50].

Now we prove our theoretical results. All theorem statements are provided here, while the proofs can be found in Appendices B – D. First, we provide the convergence guarantee for the low- $Q$ . Recall, the PPR technique guarantees that the MA-TLQL dependence on high- $Q$  is only until a finite time step during training. After this step, the agent only uses its low- $Q$  for action selection. As the convergence result in Theorem 1 is provided in the time limit ( $t \rightarrow \infty$ ), the influence of high- $Q$  can be neglected for this result.

**THEOREM 1.** Given Assumptions 1, 2, 3, the low- $Q$  values of an agent  $j$  converges to its Nash  $Q$  value in the limit ( $t \rightarrow \infty$ ).

Next, we provide sample complexity bounds for the MA-TLQL algorithm. Instead of explicitly considering the high- $Q$  values, we specify that the underlying joint policy has a covering time of  $L$ . The covering time specifies an upper bound on the number of time steps needed for all state-joint action pairs to be visited at least once starting from any state-joint action pair. Further, since the action selection is only based on the low- $Q$  values in the limit ( $t \rightarrow \infty$ ), we are most interested in the sample complexity of low- $Q$ , where the dependence on the high- $Q$  is effectively represented by  $L$ .

Regarding sample complexity, as is done in [5], we distinguish between two kinds of learning rates. Consider the following equation for the low- $Q$  (rewriting Eq. 3 and dropping *low* for simplicity),

$$Q_{t+1}^j(s_t, a_t) = (1 - \alpha_t^\omega(s_t, a_t))(Q_t^j(s_t, a_t) + \alpha_t^\omega(s_t, a_t)(r_t^j + \gamma \max_{a'} Q_t^j(s_{t+1}, a_{t+1}))). \quad (4)$$

The value of  $\alpha_t^\omega(s, a) = \frac{1}{\#(s, a, t)^\omega}$ , where  $\#(s, a, t)$  is the number of times until  $t$  that the joint action  $a$  is performed at  $s$ . Here, we consider  $\omega \in (1/2, 1]$ . The learning rate is linear if  $\omega = 1$ , and the learning rate is polynomial if  $\omega \in (1/2, 1)$ .

The next theorem provides a lower bound on the number of time steps needed for convergence in the case of a polynomial learning rate. From Assumption 1, let us specify that all rewards for the agent  $j$  are bounded by  $R_{\max}^j$ . We consider a variable  $Q_{\max}^j$ , which denotes the maximum possible low- $Q$  value for the agent  $j$ , which is bounded by  $Q_{\max}^j = R_{\max}^j/(1 - \gamma)$ . Additionally, we also use another variable  $\beta = (1 - \gamma)/2$  to present our upcoming results concisely.

**THEOREM 2.** Let us specify that with probability at least  $1 - \delta$ , for an agent  $j$ ,  $\|Q_T^j - Q_*^j\|_\infty \leq \epsilon$ . The bound on the rate of convergence of low- $Q$ ,  $Q_T^j$ , with a polynomial learning rate of factor  $\omega$  is given by (with  $Q_*^j$  as the Nash  $Q$ -value of the agent  $j$ )

$$T = \Omega\left(\left(\frac{L^{1+3\omega} Q_{\max}^{2,j} \ln\left(\frac{|S|\Pi_i|A_i|Q_{\max}^j}{\delta\beta\epsilon}\right)}{\beta^2\epsilon^2}\right)^{1-\omega} / L + \left(\left(\frac{L}{\beta} \ln \frac{Q_{\max}^j}{\epsilon} + 1\right)/2\right)^{\frac{1}{1-\omega}}\right). \quad (5)$$

Assuming the same action spaces for all agents (i.e.  $|A_1| = |A_2| = \dots = |A_N| = |A|$ ), we note that the dependence on the number of agents is  $\ln |A|^N = N \ln |A|$ . Overall this results in a sub-linear dependence on the number of agents based on the value of  $\omega$ , which is far superior to recent works that report an exponential dependence on the number of agents when learning in general-sum stochastic game environments (with an arbitrary number of agents) for convergence to a Nash equilibrium [23, 35]. Further, the dependence on the state space and action space in Theorem 2 is sub-linear ( $\ln |S|$ ), and the dependence on the covering time is  $\Omega(L^{2\omega-3\omega^2} + L^{1/1-\omega})$ , which is a polynomial dependence.

The next theorem considers the linear learning rate case.

**THEOREM 3.** Let us specify that with probability at least  $1 - \delta$ , for an agent  $j$ ,  $\|Q_T^j - Q_*^j\|_\infty \leq \epsilon$ . The bound on the rate of convergence of low- $Q$ ,  $Q_T^j$ , with a linear learning rate is given by

$$T = \Omega\left((L + \psi L + 1)^{\frac{1}{\beta}} \ln \frac{Q_{\max}^j}{\epsilon} \frac{Q_{\max}^{2,j} \ln\left(\frac{|S|\Pi_i|A_i|Q_{\max}^j}{\delta\beta\epsilon\psi}\right)}{\beta^2\epsilon^2\psi^2}\right), \quad (6)$$

where  $\psi$  is a small arbitrary positive constant satisfying  $\psi \leq 0.712$ .

Theorem 3 shows that the bound is linear in the number of agents and sub-linear in the state and action spaces. This linear dependence on the number of agents is also superior to prior results [23, 35]. Note, the dependence on the covering time in Theorem 3 could be much worse than that of Theorem 2, depending on the value of  $Q_{\max}^j$  and  $\epsilon$ . Since the value of  $\epsilon$  is small, the dependence is certainly worse than that obtained for the polynomial learning rate case. Also, the dependence on  $Q_{\max}^j$  is exponential as opposed to a polynomial dependence for Theorem 2. The last two theorems illustrate the performance benefit in using a polynomial learning rate as opposed to a linear learning rate in our algorithm.

## 6 EXPERIMENTS AND RESULTS

We consider three different experimental domains, one each for competitive, cooperative, and mixed settings, where each agent has access to a set of four advisors. We use neural network implementations of MA-TLQL and MA-TLAC, along with 5 other baselines: DQN [26], DQfD [12], CHAT [46], ADMIRAL-DM [37], and TLQL [21]. In Appendix F, we tabulate the characteristics of these baselines and provide further details regarding our choices. Since CHAT and ADMIRAL-DM assume the presence of a single advisor, we use a weighted random policy approach for implementing these two algorithms in the multiple-advisor setting, as in Li et al. [21]. If different advisors provide different actions at the same state, each action is weighted based on the number of advisors suggesting that action. For DQfD, during pre-training [12], we populate the replay buffer using advisor demonstrations from all the available advisors. For all our experiments, we will describe the critical details here, while the complete description is in Appendix K. All the experiments are repeated 30 times, with averages and standard deviations reported. For statistical significance we use the unpaired 2-sided t-test and report  $p$ -values, where  $p < 0.05$  is considered significant. The tests compare the highest performing algorithm (typically MA-TLQL) with the second-best baseline and best/average advisor performance. We conduct a total of seven experiments. The code for all experiments is open-sourced [36].

Experiments 1–4 use the competitive, two-agent version of Pommerman [30]. The environment is complex, with each state containing roughly 200 elements related to agent position and special features (e.g., bombs). The reward function is sparse: agents only receive a terminal reward of  $\{-1, 0, +1\}$ . Experiments are conducted in two phases. In the first phase (training), our algorithms and the baselines train against a standard DQN opponent for 50,000 episodes, where we plot the cumulative rewards. During this phase, algorithms can use advisors to accelerate training. In the second phase (execution), we test the performance of the trained policies against DQN for 1000 episodes, where we plot the win rate (fraction of games won) for each algorithm. During this phase, agents cannot access advisors, take no exploratory actions, and do not learn. All advisors pertaining to these four experiments are rule-based agents.

**Experiment 1:** Our first experiment uses a set of four advisors ranked in terms of quality from Advisor 1 to Advisor 4. Here, Advisor 1 is the best advisor, capable of teaching the agent all skills needed to win the game of Pommerman, and Advisor 4 only suggests random actions. In Pommerman, there is a fixed set of six skills that an agent needs to master to be able to win [30]. Since

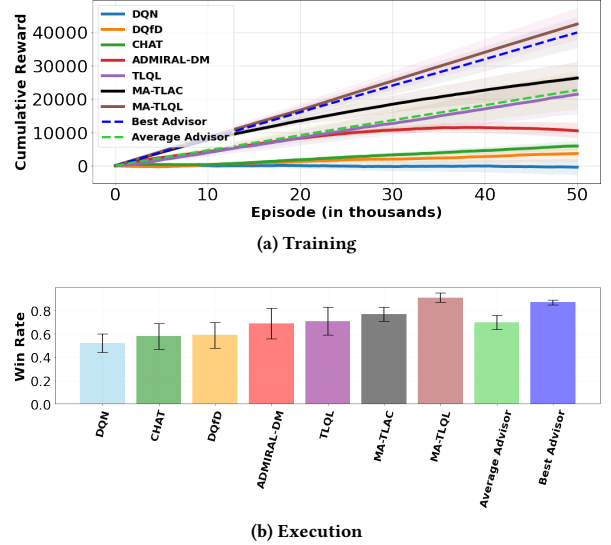
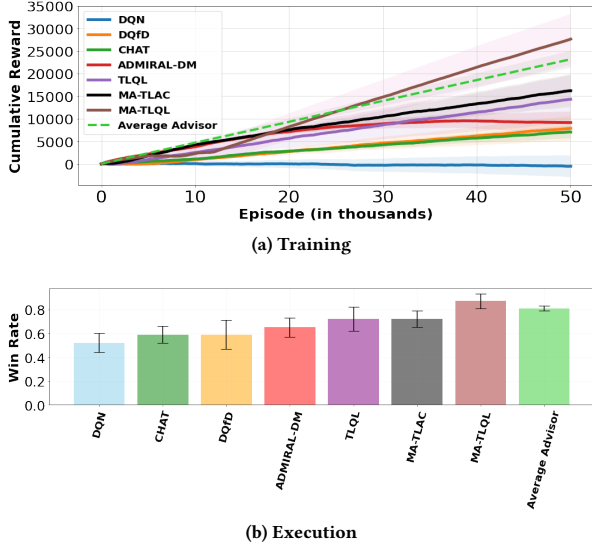


Figure 2: Two agent Pommerman with four sufficient advisors of different quality (Experiment 1)

this set of advisors can teach all these skills, we say the agent has access to a *sufficient set* of advisors. We plot the training and execution performances in Figure 2(a) and (b) respectively, including the performance of the best and average advisors (average of all Advisors 1–4) against DQN. MA-TLQL gives the best performance ( $p < 0.01$ ) and is the only algorithm providing a better performance than the best advisor ( $p < 0.11$ ) in both training and execution. MA-TLAC performs better than the average advisor ( $p < 0.04$ ). None of the others show better performances than the average advisor. CHAT and ADMIRAL-DM are not capable of leveraging and distinguishing amongst a set of advisors. DQfD uses pre-training, which is not very effective in the non-stationary multi-agent context. Learning from online advising is preferable in MARL. Also, DQfD and CHAT are independent techniques that are not actively tracking the opponent’s performance. While TLQL is capable of learning from multiple advisors, its independent nature in addition to coupling of advisor values with the RL policy reduces its effectiveness in multi-agent environments. MA-TLQL gives a better performance than MA-TLAC in both training and execution ( $p < 0.01$ ). As noted previously, the  $Q$ -learning family of algorithms tends to induce a positive bias while using the maximum action value, which leads to providing the best possible response [44]. This explains the superior performance of MA-TLQL. We conclude that MA-TLQL is capable of leveraging a set of good and bad advisors. Further, the training results in Figure 2(a) show that MA-TLQL is able to learn a better policy faster than the baselines by using advisors ( $p < 0.01$ ). The evaluation results in Figure 2(b) show that amongst all algorithms trained for the same number of episodes, MA-TLQL provides the best performance, when deployed without any advisors ( $p < 0.01$ ). Both observations point to better sample efficiency in MA-TLQL. Supplementary experiments in Appendix E show that MA-TLQL



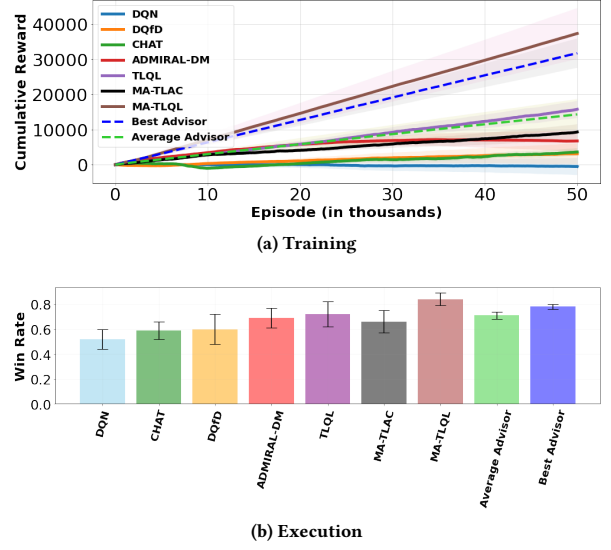
comes to relying more on good advisors than poor advisors, as compared to the baselines, illustrating its superiority.



**Figure 3: Two-agent Pommerman with four sufficient advisors of similar quality (Experiment 2)**

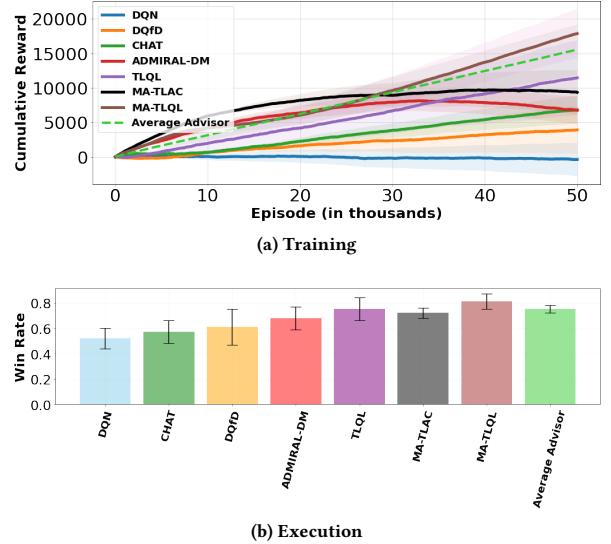
**Experiment 2:** We use the same domain as in Experiment 1, but with a different set of advisors. Now, all four advisors can teach strictly different Pommerman skills. For example, Advisor 1 can teach how to escape the enemy (and nothing else), and Advisor 2 can teach how to obtain necessary power-ups (and nothing else — full details are in Appendix K). These advisors provide pseudo-random action advice in states outside their expertise. This set of advisors is also a sufficient set. Now, learning agents must decide what advisor to listen to in the current state. From the training and execution results in Figure 3(a) and (b), we see that MA-TLQL gives the best overall performance ( $p < 0.02$ ), exceeding the average performance of the four advisors ( $p < 0.05$ ). Since all four advisors have similar quality, we only choose to use the average performance of the four advisors in this experiment for comparison. We conclude that MA-TLQL is capable of leveraging the combined knowledge of a set of advisors with different individual expertise, during learning.

**Experiment 3:** We use the same domain as in Experiment 1 but with a different set of four advisors. These advisors are similar to the set of advisors in our first experiment, where Advisor 1 gives the best advice throughout the domain, and Advisor 4 is random. However, this set of advisors is *not* capable of teaching all the strategies (i.e., Pommerman skills) needed to win in Pommerman, and compose an insufficient set (more details in Appendix K). It is critical for agents to learn from the environment in addition to the advisors. Training and execution results in Figure 4 shows the superior performance of MA-TLQL, the only algorithm that outperforms the best advisor ( $p < 0.05$ ) and all baselines ( $p < 0.02$ ). Surprisingly, TLQL performs better than MA-TLAC ( $p < 0.02$ ), likely due to the positive bias of  $Q$ -learning. This experiment reinforces the observation that MA-TLQL is capable of learning from good advisors



**Figure 4: Two-agent Pommerman with four insufficient advisors of different quality (Experiment 3)**

and avoids bad advisors (also see Appendix E). Since MA-TLQL outperforms the best advisor, this experiment demonstrates that MA-TLQL can learn from both, advisors and through direct interactions with the environment, hence having a much improved sample efficiency as compared to other algorithms that learn only from the environment. This is observed during both training and execution.



**Figure 5: Two-agent Pommerman with four insufficient advisors of similar quality (Experiment 4)**

**Experiment 4:** This is similar to the Experiment 2: four advisors have similar quality, but each understands a different Pommerman

skill. However, our set of advisors in this experiment are insufficient to teach all the skills in Pommerman, and the agent must also learn from the environment. The results in Figure 5 shows that MA-TLQL is capable of leveraging the combined expertise of the advisors and learning from the environment to obtain the best performance, as compared to the baselines ( $p < 0.04$ ) and advisors ( $p < 0.05$ ). This makes MA-TLQL more sample efficient than the prior algorithms.

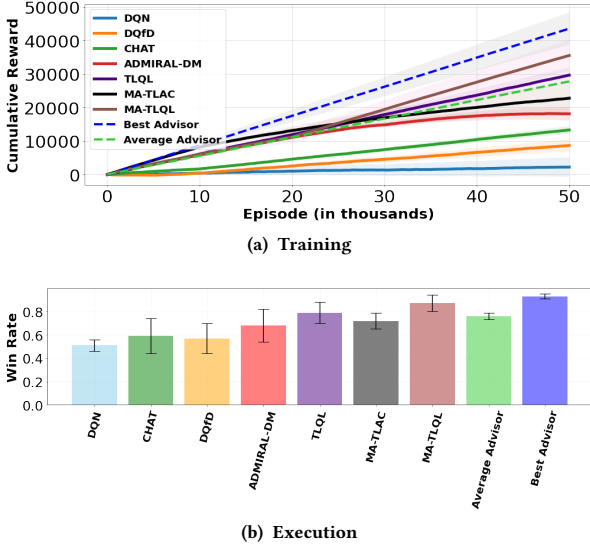


Figure 6: Team (mixed) Pommerman (Experiment 5)

**Experiment 5:** We now switch to a four-agent version of Pommerman, which is two vs. two. This is a mixed setting as agents need to learn cooperative as well as competitive skills. Overall, this is a more complex domain with a larger state space. We consider four sufficient advisors of different quality, similar to Experiment 1. We conduct two phases — training (for 50,000 episodes) and execution (for 1000 episodes). The training and execution results in Figure 6 show that MA-TLQL provides the best performance compared to the baselines ( $p < 0.04$ ) but does not perform better than the best available advisor. Since this is a more complex domain, MA-TLQL needs a larger training period for learning good policies. However, MA-TLQL still performs better than the average performance of the four advisors ( $p < 0.03$ ). We conclude that although MA-TLQL’s performance suffers in the more difficult mixed setting, it still outperforms all the other baselines and is capable of distinguishing between good and bad advisors (see also Appendix E). From both training and execution results in Figure 6, we note that MA-TLQL has a superior sample efficiency as compared to the other baselines.

**Experiment 6:** This experiment switches to the cooperative Pursuit domain [8]. There are eight pursuer learning agents that learn to capture a set of 30 randomly moving targets (evaders) (details in Appendix K). We use four pre-trained DQN networks as the advisors, learning for 500, 1000, 1500, and 2000 episodes, respectively. We again have two phases — training and execution. During training, all algorithms are trained for 2000 episodes. The trained networks are then used in the execution phase for 100 episodes with




Figure 7: Cooperative Pursuit setting (Experiment 6)

no further training or influence from advisors. Figure 7(a) plots the episodic rewards obtained during training and the Figure 7(b) plots the number of targets captured in the execution phase, where MA-TLQL shows the best performance ( $p < 0.03$ ). Hence, MA-TLQL can outperform all baselines in a cooperative environment as well.




Figure 8: Mixed Predator-Prey setting (Experiment 7)

**Experiment 7:** This final experiment considers a mixed cooperative competitive Predator-Prey environment which is a part of the Multi Particle Environment (MPE) suite [24]. Our implementation uses a discrete action space and a continuous state space (more details in Appendix K). There are a total of eight predators trying to capture eight prey (prey are not removed, but respawned upon capture). In our experiment, each algorithm trains the predators while the prey is trained using a standard DQN opponent. The experiments have two phases of training and execution, which is modelled as a CTDE setting. Here each agent obtains information about the actions and rewards of all other agents during training, but only has local observation during execution. Since this environment requires decentralization during execution, we omit the fully centralized MA-TLQL and ADMIRAL-DM. We also omit DQfD since it gave poor performances previously. As in Experiment 6, we use four pretrained DQN (predator) networks as advisors (trained for 1000, 2000, 7000, and 12000 episodes). Training is conducted for 12000 episodes and execution is conducted for 100 episodes. The training results in Figure 8(a) (plot of episodic rewards) show that MA-TLQL is the most sample efficient compared to other algorithms as it is able to leverage the available advisors better than others, thus outperforming them ( $p < 0.04$ ). The execution results in Figure 8(b) plots the average prey captured by each algorithm. MA-TLQL outperforms others during execution as well ( $p < 0.03$ ).




Figure 9: Ablation results using Experiment 1

From all the p-values across the seven experiments, we note that most of our observations are statistically significant. Despite observing MA-TLQL outperforming the best advisor in many of the experiments, some of these comparisons are not statistically significant (i.e.,  $p \geq 0.05$ ). While the main experiments of the paper consider fixed advisors, our algorithms can also be implemented with learning/changing advisors (see Appendix G). In Appendix I we study performances under different numbers of advisors. Also, our algorithms can be used along with opponent modelling techniques as done by prior works [9] (more details in Appendix H).

## 7 ABLATION STUDY

In this section, we run an ablation study on the three components of MA-TLQL that differ from the previously introduced TLQL algorithm by Li et al. [21]. To recall these three components are: i) joint action (JA) updates, ii) ensemble method (EM), and iii) advisor evaluation (AE). For this ablation study we will consider the two-agent version of Pommerman with four sufficient advisors having different (Experiment 1) and similar quality (Experiment 2).

The ablation results corresponding to Experiment 1 are given in Figure 9, where we plot the performances of TLQL and MA-TLQL in addition to TLQL with each of the three components. In Figure 9(a) and (b), the performance of TLQL with each of the three components is better than vanilla TLQL. TLQL using the ensemble method (i.e., TLQL+EM) is able to perform better than vanilla TLQL, since at the beginning of training the  $Q$ -values of the advisors are not accurate, and the ensemble technique chooses the advisor action that is agreed upon by most advisors in the given set (in line with our discussions in Section 4). Recall that the set of four different advisors had four advisors of decreasing quality, with the first three advisors capable of teaching some useful Pommerman skills and the last advisor being just random (see Appendix K). Using the ensemble prevents the use of the random advisor, as the first three advisors are more likely to agree upon an action, increasing the possibility of the agent choosing that action. Further, we see that TLQL highly benefits from using the joint action update (i.e., TLQL+JA) instead of an independent update seen in vanilla TLQL. The joint action update explicitly considers the strategies of other agent(s) and helps in providing stronger best responses as compared to an independent update in the multi-agent environments. Finally, TLQL using advisor evaluation in the high- $Q$  table (i.e., TLQL+AE) provides the best benefit compared to the




Figure 10: Ablation results using Experiment 2

other components. As discussed in Section 3 and Section 4, the high- $Q$  definition in vanilla TLQL is limiting since the advisor evaluation through the high- $Q$  is coupled with the inaccurate RL policy (and AE addresses this limitation). Further, from Figure 9, we see that MA-TLQL (integrating all the three components) shows the best performance as compared to vanilla TLQL and individual TLQL implementations with each of the three components ( $p < 0.05$ ). Thus, MA-TLQL is able to seamlessly integrate the advantages of each of the individual components of TLQL, demonstrating its superiority.

We also consider a similar ablation study using Experiment 2 (see Figure 10). As in Figure 9, we see that TLQL with each of the three components performs better than vanilla TLQL. Since we have four advisors of similar quality where each advisor is good at a different Pommerman skill, their agreement on an action is expected to be small. Hence, the ensemble technique (i.e., TLQL+EM) provides only a small improvement over vanilla TLQL. However, the other two components (i.e., TLQL+JA and TLQL+AE) provides a good performance benefit over TLQL. Finally, MA-TLQL, that integrates all the three components, provides the best performance ( $p < 0.03$ ).

## 8 CONCLUSION

This paper provided a principled approach for learning from multiple independent advisors in MARL. Inspired by Li et al. [21], we present a two-level architecture for multi-agent environments. We discuss two limitations in TLQL and address these limitations in our approach. Also, we provide a fixed point guarantee and sample complexity bounds regarding the learning of MA-TLQL. Additionally, we provided an actor-critic implementation that can work in the CTDE paradigm. Further, we performed an extensive experimental analysis of MA-TLQL and MA-TLAC in cooperative, competitive, and mixed settings, where we show that these algorithms are capable of suitably leveraging a set of advisors, and perform better than baselines. As future work, we would like to consider human advisors and further explore some avenues in the real-world context.

## ACKNOWLEDGEMENTS

Resources used in preparing this research were provided by the province of Ontario and the government of Canada through CIFAR, NSERC and companies sponsoring the Vector Institute. Part of this work has taken place in the IRL Lab at the University of Alberta, which is supported in part by research grants from the Alberta Machine Intelligence Institute (Amii); a Canada CIFAR AI Chair, Amii; Compute Canada; Huawei; Mitacs; and NSERC.



## REFERENCES

- [1] Ofra Amir, Ece Kamar, Andrey Kolobov, and Barbara J. Grosz. 2016. Interactive Teaching Strategies for Agent Training. In *IJCAI*. IJCAI Press, New York, NY, USA, 9-15 July 2016, 804–811.
- [2] Samuel Barrett, Avi Rosenfeld, Sarit Kraus, and Peter Stone. 2017. Making friends on the fly: Cooperating with new teammates. *Artificial Intelligence* 242 (2017), 132–171.
- [3] Dimitri P. Bertsekas and John N. Tsitsiklis. 1996. *Neuro-dynamic programming*. Optimization and neural computation series, Vol. 3. Athena Scientific, Chestnut Street, USA.
- [4] Tim Brys, Anna Harutyunyan, Halit Bener Suay, Sonia Chernova, Matthew E. Taylor, and Ann Nowé. 2015. Reinforcement Learning from Demonstration through Shaping. In *IJCAI, July 25-31, 2015*. AAAI Press, Buenos Aires, Argentina, 3352–3358.
- [5] Eyal Even-Dar and Yishay Mansour. 2003. Learning Rates for Q-learning. *Journal of Machine Learning Research* 5 (2003), 1–25.
- [6] Fernando Fernández and Manuela M. Veloso. 2006. Probabilistic policy reuse in a reinforcement learning agent. In *5th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS 2006), May 8-12, 2006*. ACM, Hakodate, Japan, 720–727.
- [7] Yang Gao, Huazhe Xu, Ji Lin, Fisher Yu, Sergey Levine, and Trevor Darrell. 2018. Reinforcement Learning from Imperfect Demonstrations. In *ICLR, April 30 - May 3, 2018, Workshop Track Proceedings*. OpenReview.net, Vancouver, BC, Canada.
- [8] Jayesh K Gupta, Maxim Egorov, and Mykel Kochenderfer. 2017. Cooperative multi-agent control using deep reinforcement learning. In *AAMAS*. Springer, IFAAMAS, Sao Paulo, Brazil, 66–83.
- [9] He He, Jordan Boyd-Graber, Kevin Kwok, and Hal Daumé III. 2016. Opponent modeling in deep reinforcement learning. In *International conference on machine learning*. PMLR, New York City, US, 1804–1813.
- [10] Bassam Helou, Aditya Dusi, Anne Collin, Noushin Mehdipour, Zhiliang Chen, Cristhian Lizarazo, Calin Belta, Tichakorn Wongpiromsarn, Radboud Duintjer Tebbens, and Oscar Beijbom. 2021. The Reasonable Crowd: Towards evidence-based and interpretable models of driving behavior. In *IROS*. IEEE, Prague, Czech Republic, 6708–6715.
- [11] Pablo Hernandez-Leal, Bilal Kartal, and Matthew E. Taylor. 2019. A survey and critique of multiagent deep reinforcement learning. *Autonomous Agents and Multi-Agent Systems* 33, 6 (01 Nov 2019), 750–797. <https://doi.org/10.1007/s10458-019-09421-1>
- [12] Todd Hester, Matej Vecerík, Olivier Pietquin, Marc Lanctot, Tom Schaul, Bilal Piot, Dan Horgan, John Quan, Andrew Sendonaris, Ian Osband, Gabriel Dulac-Arnold, John P. Agapiou, Joel Z. Leibo, and Audrunas Gruslys. 2018. Deep Q-learning From Demonstrations. In *AAAI, February 2-7, 2018*. AAAI Press, New Orleans, Louisiana, USA.
- [13] Junling Hu and Michael P. Wellman. 2003. Nash Q-learning for general-sum stochastic games. *JMLR* 4, Nov (2003), 1039–1069.
- [14] Tommi Jaakkola, Michael I. Jordan, and Satinder P. Singh. 1994. On the Convergence of Stochastic Iterative Dynamic Programming Algorithms. *Neural Computation* 6, 6 (1994), 1185–1201. <https://doi.org/10.1162/neco.1994.6.6.1185>
- [15] Piyush Jain, Sean CP Coogan, Sriram Ganapathi Subramanian, Mark Crowley, Steve Taylor, and Mike D Flannigan. 2020. A review of machine learning applications in wildfire science and management. *Environmental Reviews* 28, 4 (2020), 478–505.
- [16] Mingxuan Jing, Xiaojian Ma, Wenbing Huang, Fuchun Sun, Chao Yang, Bin Fang, and Huaping Liu. 2020. Reinforcement Learning from Imperfect Demonstrations under Soft Expert Guidance. In *AAAI, February 7-12, 2020*. AAAI Press, New York, NY, USA, 5109–5116.
- [17] Beomjoon Kim, Amir-massoud Farahmand, Joelle Pineau, and Doina Precup. 2013. Learning from Limited Demonstrations. In *NeurIPS*. Morgan Kaufmann Publishers, Lake Tahoe, Nevada, United States, 2859–2867.
- [18] Dong-Ki Kim, Miao Liu, Shayegan Omidshafiei, Sebastian Lopez-Cot, Matthew Riemer, Golnaz Habibi, Gerald Tesauro, Sami Mourad, Murray Campbell, and Jonathan P. How. 2020. Learning Hierarchical Teaching Policies for Cooperative Agents. In *AAMAS, May 9-13, 2020*. IFAAMAS, Auckland, New Zealand, 620–628.
- [19] Vijay R. Konda and John N. Tsitsiklis. 1999. Actor-Critic Algorithms. In *NeurIPS*. The MIT Press, Denver, CO, USA.
- [20] Sergey Levine, Chelsea Finn, Trevor Darrell, and Pieter Abbeel. 2016. End-to-end training of deep visuomotor policies. *JMLR* 17, 1 (2016), 1334–1373.
- [21] Mao Li, Yi Wei, and Daniel Kudenko. 2019. Two-level Q-learning: learning from conflict demonstrations. *The Knowledge Engineering Review* 34 (2019).
- [22] Michael L. Littman. 1994. Markov Games as a Framework for Multi-Agent Reinforcement Learning. In *ICML, July 10-13, 1994*. Morgan Kaufmann, New Brunswick, NJ, USA, 157–163.
- [23] Qinghua Liu, Tiancheng Yu, Yu Bai, and Chi Jin. 2021. A Sharp Analysis of Model-based Reinforcement Learning with Self-Play. In *ICML (Proceedings of Machine Learning Research, Vol. 139)*. PMLR, Virtual Event, 7001–7010.
- [24] Ryan Lowe, Yi Wu, Aviv Tamar, Jean Harb, Pieter Abbeel, and Igor Mordatch. 2017. Multi-Agent Actor-Critic for Mixed Cooperative-Competitive Environments. In *NeurIPS, December 4-9, 2017*. Morgan Kaufmann Publishers, Long Beach, CA, USA, 6379–6390.
- [25] Laetitia Matignon, Guillaume J Laurent, and Nadine Le Fort-Piat. 2012. Independent reinforcement learners in cooperative markov games: a survey regarding coordination problems. *The Knowledge Engineering Review* 27, 1 (2012), 1–31.
- [26] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G Bellemare, Alex Graves, Martin Riedmiller, Andreas K Fidjeland, Georg Ostrovski, et al. 2015. Human-level control through deep reinforcement learning. *Nature* 518, 7540 (2015), 529–533.
- [27] Shayegan Omidshafiei, Dong-Ki Kim, Miao Liu, Gerald Tesauro, Matthew Riemer, Christopher Amato, Murray Campbell, and Jonathan P. How. 2019. Learning to Teach in Cooperative Multiagent Reinforcement Learning. In *AAAI, January 27 - February 1, 2019*. AAAI Press, Honolulu, Hawaii, USA, 6128–6136.
- [28] Bilal Piot, Matthieu Geist, and Olivier Pietquin. 2014. Boosted Bellman Residual Minimization Handling Expert Demonstrations. In *ECLM-PKDD, September 15-19, 2014*, Vol. 8725. Springer, Nancy, France, 549–564.
- [29] Tummalaipalli Sudhamsh Reddy, Vamsikrishna Gopikrishna, Gergely V. Zaruba, and Manfred Huber. 2012. Inverse reinforcement learning for decentralized non-cooperative multiagent systems. In *Proceedings of the IEEE International Conference on Systems, Man, and Cybernetics (SMC 2012), October 14-17, 2012*. IEEE, Seoul, Korea (South), 1930–1935.
- [30] Cinjon Resnick, Wes Eldridge, David Ha, Denny Britz, Jakob Foerster, Julian Togelius, Kyunghyun Cho, and Joan Bruna. 2018. Pommernan: A multi-agent playground. *arXiv preprint arXiv:1809.07124* (2018).
- [31] Lloyd S Shapley. 1953. Stochastic games. *Proceedings of the national academy of sciences* 39, 10 (1953), 1095–1100.
- [32] Felipe Leno Da Silva and Anna Helena Real Costa. 2019. A Survey on Transfer Learning for Multiagent Reinforcement Learning Systems. *JAIR* 64 (2019), 645–703.
- [33] Felipe Leno Da Silva, Ruben Glatt, and Anna Helena Real Costa. 2017. Simultaneously Learning and Advising in Multiagent Reinforcement Learning. In *AAMAS, May 8-12, 2017*. ACM, Sao Paulo, Brazil, 1100–1108.
- [34] David Silver, Aja Huang, Chris J Maddison, Arthur Guez, Laurent Sifre, George Van Den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, et al. 2016. Mastering the game of Go with deep neural networks and tree search. *Nature* 529, 7587 (2016), 484.
- [35] Ziang Song, Song Mei, and Yu Bai. 2021. When Can We Learn General-Sum Markov Games with a Large Number of Players Sample-Efficiently? *arXiv preprint arXiv:2110.04184* (2021).
- [36] Sriram Ganapathi Subramanian. [n.d.]. Learning from Multiple Independent Advisors in Multi-agent Reinforcement Learning. <https://github.com/Sriram94/matql>
- [37] Sriram Ganapathi Subramanian, Kate Larson, Matthew Taylor, and Mark Crowley. 2022. Multi-Agent Advisor Q-Learning. *Journal of Artificial Intelligence Research* 74 (2022), 1–74.
- [38] Richard S Sutton and Andrew G Barto. 1998. *Introduction to reinforcement learning*. Vol. 135. MIT press, Cambridge.
- [39] Csaba Szepesvari and Michael L Littman. 1999. A unified analysis of value-function-based reinforcement-learning algorithms. *Neural computation* 11, 8 (1999), 2017–2060.
- [40] Ming Tan. 1993. Multi-agent reinforcement learning: Independent vs. cooperative agents. In *ICML*. Cambridge University Press, Amherst, MA, USA, 330–337.
- [41] Matthew E. Taylor, Halit Bener Suay, and Sonia Chernova. 2011. Integrating reinforcement learning with human demonstrations of varying ability. In *AAMAS, May 2-6, 2011*. IFAAMAS, Taipei, Taiwan, 617–624.
- [42] Justin K Terry, Benjamin Black, Mario Jayakumar, Ananth Hari, Luis Santos, Clemens Dieffendahl, Niall L Williams, Yashas Lokesh, Ryan Sullivan, Caroline Horsch, and Praveen Ravi. 2020. PettingZoo: Gym for Multi-Agent Reinforcement Learning. (2020).
- [43] Lisa Torrey and Matthew E. Taylor. 2013. Teaching on a budget: agents advising agents in reinforcement learning. In *AAMAS, May 6-10, 2013*. IFAAMAS, Saint Paul, MN, USA, 1053–1060.
- [44] Hado van Hasselt. 2010. Double Q-learning. In *NeurIPS*. Curran Associates, Inc., Vancouver, British Columbia, Canada, 2613–2621.
- [45] Yixi Wang, Wenhuan Lu, Jianye Hao, Jianguo Wei, and Ho-fung Leung. 2018. Efficient Convention Emergence through Decoupled Reinforcement Social Learning with Teacher-Student Mechanism. In *AAMAS, July 10-15, 2018*. IFAAMAS / ACM, Stockholm, Sweden, 795–803.
- [46] Zhaodong Wang and Matthew E. Taylor. 2017. Improving Reinforcement Learning with Confidence-Based Demonstrations. In *IJCAI, August 19-25, 2017*. ICJAI, Melbourne, Australia, 3027–3033.
- [47] Christopher JH Watkins and Peter Dayan. 1992. Q-learning. *Machine Learning* 8, 3-4 (1992), 279–292.
- [48] Kevin Waugh, Brian D. Ziebart, and Drew Bagnell. 2011. Computational Rationalization: The Inverse Equilibrium Problem. In *ICML, June 28 - July 2, 2011*. Omnipress, Bellevue, Washington, USA, 1169–1176.
- [49] Tianpei Yang, Weixun Wang, Hongyao Tang, Jianye Hao, Zhao Peng Meng, Hangyu Mao, Dong Li, Wulong Liu, Yingfeng Chen, Yujing Hu, et al. 2021. An

Efficient Transfer Learning Framework for Multiagent Reinforcement Learning. *NeurIPS, Virtual Event* 34 (2021).

- [50] Yaodong Yang, Rui Luo, Minne Li, Ming Zhou, Weinan Zhang, and Jun Wang. 2018. Mean Field Multi-Agent Reinforcement Learning. In *ICML*, Vol. 80. PMLR, Stockholm Sweden, 5571–5580.
- [51] Dayong Ye, Tianqing Zhu, Zishuo Cheng, Wanlei Zhou, and S Yu Philip. 2020. Differential Advising in Multi-Agent Reinforcement Learning. *IEEE Transactions on Cybernetics* (2020).

## A ALGORITHM PSEUDOCODES

A complete pseudocode of a tabular implementation of our  $Q$ -learning based algorithm (MA-TLQL) is given in Algorithm 1. All agents initialize a low- $Q$  table and a high- $Q$  table in line 2. Then at each state, all agents choose to perform an action in lines 8–20. This action can come from the advisor or the RL policy as described in Section 4. Then the action is executed, and the next state and reward are observed in line 21. Finally, the  $Q$  values for the low- $Q$  as well as the high- $Q$  are updated (line 22 and line 23) according to equations presented in Section 4. The value of  $\epsilon'$  is linearly decayed from a high-value to a value close to zero during training (line 24).

To make Algorithm 1 applicable to high dimensional state and action spaces, we provide a function approximation-based implementation of MA-TLQL in Algorithm 2. Here neural networks are used as the function approximator, and the algorithm uses a separate target network and a replay buffer for training, as introduced in the well-known DQN algorithm [26]. The agent maintains a high- $Q$  network and two low- $Q$  networks (evaluation and target networks) and updates these networks using the temporal difference (T.D.) errors with the update equations presented in Section 4. If the full state of the stochastic game is not available, the agent can simply use its observation instead of the state, as applicable in most function approximation-based RL methods.

We also extend Algorithm 2 to an actor-critic implementation described in Algorithm 3. This algorithm is called multi-agent two-level actor-critic (MA-TLAC). This algorithm uses the policy as the actor and the  $Q$ -values as the critic, consistent with prior work [19]. We maintain two actors, and two critics to reflect the two-level (high and low) nature of our algorithm. The high-level actor determines an advisor and the high-level critic helps train the high-level actor, using the T.D. errors. Similarly, the low-level actor determines the appropriate action, with the low-level critic providing the T.D. errors for training. In MA-TLAC, since we use a separate actor network for advisor selection, we do not use the ensemble technique from Eq. 2. Instead, the high-level actor directly chooses one amongst the given advisors for the current state. The advantage of this algorithm is that it can be implemented using the popular CTDE paradigm [24], since the actors do not require the actions of other agents for action/advisor selection. In CTDE, global information (i.e., information from other agents) is available during training time but not available during execution. This CTDE based implementation also allows our method to be used in partially observable domains, since the actors can use the local observations for action/advisor selection while the critic can use the joint actions and states during training, as described in Lowe et al. [24]. Also, since the ensemble technique (Eq. 2) is not used in MA-TLAC, it is also applicable to continuous state space environments as well (unlike MA-TLQL, which is only applicable to environments with discrete state spaces).

All the algorithm pseudocode provided in this section assume that all agents are using the same algorithmic steps for learning where it can maintain copies of updates of other agents, as done in prior work [13]. If this is not possible, the agents would directly use the observed previous actions of other agents for its updates.

---

### Algorithm 1 MA-TLQL Tabular Method

---

```

1: Let  $Ad^j$  denote a set of advisors available to the agent  $j$ .
2: For all  $j \in 1, \dots, N$ ,  $s \in S$ , and  $a^j \in A^j$ :  $lowQ^j(s, a^j, \mathbf{a}^{-j}) \leftarrow 0$ 
   where  $\mathbf{a}^{-j} = [a^1, \dots, a^{j-1}, a^{j+1}, \dots, a^N]$ 
3: For all  $s \in S$ ,  $ad^j \in Ad^j$ , and for all  $j \in 1, \dots, N$ :
    $highQ^j(s, \mathbf{a}^{-j}, ad^j) \leftarrow 0$ 
4: Initialize a value for hyperparameters  $\epsilon$  and  $\epsilon'$  and  $\eta$ 
5: while training is not finished do
6:   For each agent  $j$ , get the current state  $s$ 
7:   For each agent  $j$ , get the joint actions of other agents  $\mathbf{a}^{-j}$ 
   at state  $s$  using the respective copies and previous actions of
   all agents
8:   For each agent  $j$ , let  $u$  be a uniform random number between
   0 and 1
9:   if  $u < \epsilon'$  then
10:     Let  $u'$  be a uniform random number between 0 and 1
11:     if  $u' < \eta$  then
12:       Choose an advisor using the high- $Q$  values of agent
        $j$  for the current state and joint action of other agents from
       Eq. 2 and use its action as the current action  $a_t^j$ 
13:     else
14:       Set the advisor  $ad^j$  as a random advisor from  $Ad^j$ 
       and use its action as the current action  $a_t^j$ .
15:     end if
16:   else if  $u > \epsilon'$  and  $u < \epsilon$  then
17:     Set the action  $a_t^j$  as a random action from the action
     space  $A^j$ 
18:   else
19:     Choose a greedy action  $a_t^j$  from the low- $Q$  value using
      $s$  and the joint action of other agents
20:   end if
21:   Execute the joint action  $\mathbf{a}$ , observe joint reward  $r$  and the
   next state  $s'$ , where  $\mathbf{a} = [a^1, \dots, a^N]$  and  $r = [r^1, \dots, r^N]$ 
22:   Update value of low- $Q$  for the agent  $j$  using Eq. 3. Obtain the
   next actions for other agents  $\mathbf{a}'^{-j}$  from the respective copies
   and previous actions of other agents
23:   If an advisor was chosen, update value of high- $Q$  of the
   advisor for the agent  $j$  using Eq. 1
24:   At the end of each episode, linearly decay  $\epsilon'$ 
25: end while

```

---

---

**Algorithm 2** MA-TLQL Neural Network Method

---

```
1: Let  $Ad^j$  denote a set of advisors available to the agent  $j$ 
2: Initialize  $Q_{\phi^j}, Q_{\theta^j}$  for all  $j \in 1, \dots, N$  (to denote low- $Q$ ). Initialize  $Q_{\theta^j}, Q_{\phi^j}$  for all  $j \in 1, \dots, N$  (to denote high- $Q$ )
3: Initialize a value for hyperparameters  $\epsilon$  and  $\epsilon'$  and  $\eta$ 
4: while training is not finished do
5:   For each agent  $j$ , get the current state  $s$ 
6:   For each agent  $j$ , get the joint actions of other agents  $\mathbf{a}^{-j}$  at state  $s$  using the respective copies and previous actions of all agents
7:   For each agent  $j$ , let  $u$  be a uniform random number between 0 and 1
8:   if  $u < \epsilon'$  then
9:     Let  $u'$  be a uniform random number between 0 and 1
10:    if  $u' < \eta$  then
11:      Choose an advisor using the high- $Q$  values of agent  $j$  from Eq. 2 for the current state and joint action of other agents using the high- $Q$ ,  $Q_{\theta^j}$ , and use its action as the current action  $a_t^j$ 
12:    else
13:      Set the advisor  $ad^j$  as a random advisor from  $Ad^j$  and use its action as the current action  $a_t^j$ .
14:    end if
15:    else if  $u > \epsilon'$  and  $u < \epsilon$  then
16:      Set the action  $a_t^j$  as a random action from the action space  $A^j$ 
17:    else
18:      Choose a greedy action  $a_t^j$  from the low- $Q$ ,  $Q_{\phi^j}$ , using  $s$  and the joint action of other agents
19:    end if
20:    Execute the joint action  $\mathbf{a}$ , observe joint reward  $\mathbf{r}$  and the next state  $s'$ , where  $\mathbf{a} = [a^1, \dots, a^N]$  and  $\mathbf{r} = [r^1, \dots, r^N]$ 
21:    For each agent  $j$ , store  $\langle s, \mathbf{a}, \mathbf{r}, s', \mathbf{a}' \rangle$  in replay buffer  $\mathcal{D}^j$ , where  $\mathbf{a} = [a^1, \dots, a^N]$ ,  $\mathbf{a}' = [a'^1, \dots, a'^N]$ . Obtain the next actions for other agents  $\mathbf{a}'^{-j}$  from the respective copies and previous actions of other agents
22:    If an advisor was used, for each agent  $j$ , store  $\langle s, \mathbf{a}, \mathbf{r}, s', \mathbf{a}', ad^j \rangle$  in replay buffer  $\mathcal{D}^{j'}$ , where  $ad^j$  is the advisor
23:    Set the next state  $s'$  as the current state  $s$ 
24:    At the end of each episode, linearly decay  $\epsilon'$ 
25:    while  $j = 1$  to  $N$  do
26:      Sample a minibatch of  $K$  experiences  $\langle s, \mathbf{a}, \mathbf{r}, s', \mathbf{a}' \rangle$  from  $\mathcal{D}^j$ 
27:      Set  $y^j = r^j + \gamma \max_{a'^j} Q_{\phi^j}(s', \mathbf{a}'^{-j}, a'^j)$  according to Eq. 3
28:      Update the  $Q$ -network  $\phi^j$  by minimizing the loss  $\mathcal{L}(\phi^j) = \frac{1}{K} \sum (y^j - Q_{\phi^j}(s, \mathbf{a}^{-j}, a^j))^2$ 
29:      Sample a minibatch of  $K$  experiences  $\langle s, \mathbf{a}, \mathbf{r}, s', \mathbf{a}', ad^j \rangle$  from  $\mathcal{D}^{j'}$ 
30:      Set  $y^j = r^j + \gamma Q_{\theta^j}(s', \mathbf{a}'^{-j}, ad^j)$  according to Eq. 1
31:      Update the  $Q$ -network  $\theta^j$  by minimizing the loss  $\mathcal{L}(\theta^j) = \frac{1}{K} \sum (y^j - Q_{\theta^j}(s, \mathbf{a}^{-j}, ad^j))^2$ 
32:    end while
33:    Update the parameters of the target network for each agent by copying over the evaluation network every  $\mathcal{T}$  steps:  $\phi_{\theta}^j \leftarrow \phi^j$  and  $\theta_{\phi}^j \leftarrow \theta^j$ 
34: end while
```

---

---

**Algorithm 3** MA-TLAC

---

```
1: Let  $Ad^j$  denote a set of advisors available to the agent  $j$ 
2: Initialize  $Q_{\phi^j}, \pi_{\phi^j}$ , the low-level critic and actor networks for all  $j \in \{1, \dots, n\}$ 
3: Initialize  $Q_{\theta^j}, \pi_{\theta^j}$ , the high-level critic and actor networks for all  $j \in \{1, \dots, n\}$ 
4: Initialize a value for hyperparameters  $\epsilon$  and  $\epsilon'$  and  $\eta$ 
5: while training is not finished do
6:   For each agent  $j$ , get the current state  $s$ 
7:   For each agent  $j$ , let  $u$  be a uniform random number between 0 and 1
8:   if  $u < \epsilon'$  then
9:     Let  $u'$  be a uniform random number between 0 and 1
10:    if  $u' < \eta$  then
11:      Choose an advisor  $ad^j$  using the high-level actor  $\pi_{\theta^j}$ , for the agent  $j$ , for the current state  $s$ , and use its action as the current action  $a^j$ 
12:    else
13:      Set the advisor  $ad^j$  as a random advisor from  $Ad^j$  and use its action as the current action  $a^j$ 
14:    end if
15:    else if  $u > \epsilon'$  and  $u < \epsilon$  then
16:      Set the action  $a^j$  as a random action from the action space  $A^j$ 
17:    else
18:      Choose a greedy action  $a^j$  from the low-level actor,  $\pi_{\phi^j}$ , using  $s$ 
19:    end if
20:    Execute the joint action  $\mathbf{a}$ , observe joint reward  $\mathbf{r}$  and the next state  $s'$ , where  $\mathbf{a} = [a^1, \dots, a^N]$  and  $\mathbf{r} = [r^1, \dots, r^N]$ 
21:    For each agent  $j$ , obtain the joint actions of other agents  $\mathbf{a}^{-j}$  (current observed actions of other agents) at state  $s$ 
22:    Set  $y^j = r^j + \gamma \max_{a'^j} Q_{\phi^j}(s', \mathbf{a}'^{-j}, a'^j)$  according to Eq. 3
23:    For each  $j$ , update the low-level critic by minimizing the loss  $\mathcal{L}(\phi^j) = (y^j - Q_{\phi^j}(s, \mathbf{a}^{-j}, a^j))^2$ 
24:    For each  $j$ , calculate the advantage estimate using the relation  $A(s, \mathbf{a}^{-j}, a^j) = y^j - \sum_{a^j} \pi_{\phi^j}(a^j|s) Q_{\phi^j}(s, \mathbf{a}^{-j}, a^j)$ 
25:    For each  $j$ , update the low-level actor using the log loss  $\mathcal{J}(\phi_{\pi}^j) = \log \pi_{\phi_{\pi}^j}(a^j|s) A(s, \mathbf{a}^{-j}, a^j)$ 
26:    If an agent  $j$  used an advisor  $ad^j$ , then update the advisor's  $Q$ -estimate.
27:    For each  $j$ , set  $y^j = r^j + \gamma Q_{\theta^j}(s', \mathbf{a}'^{-j}, ad^j)$  according to Eq. 1
28:    Obtain the next actions for other agents  $\mathbf{a}'^{-j}$  from the respective copies
29:    For each  $j$ , update the high-level critic by minimizing the loss  $\mathcal{L}(\theta^j) = (y^j - Q_{\theta^j}(s, \mathbf{a}^{-j}, ad^j))^2$  where  $ad^j$  is the advisor chosen by the agent  $j$ 
30:    For each  $j$ , calculate the advantage estimate using the relation  $A(s, \mathbf{a}^{-j}, ad^j) = y^j - \sum_{ad^j} \pi_{\theta^j}(ad^j|s) Q_{\theta^j}(s, \mathbf{a}^{-j}, ad^j)$ 
31:    For each  $j$ , update the high-level actor using the log loss  $\mathcal{J}(\theta_{\pi}^j) = \log \pi_{\theta_{\pi}^j}(a^j|s) A(s, \mathbf{a}^{-j}, ad^j)$ 
32:    Set the next state as the current state  $s = s'$ 
33:    At the end of each episode, linearly decay  $\epsilon'$ 
34: end while
```

---

## B PROOF OF THEOREM 1

THEOREM 1. *Given Assumptions 1, 2, 3, the low- $Q$  values of an agent  $j$  converges to its Nash  $Q$  value in the limit ( $t \rightarrow \infty$ ).*

PROOF. Our proof will be along the lines of Theorem 3 in Subramanian et al. [37].

Let us consider a lemma from prior work.

LEMMA 1. *A random iterative process*

$$\Delta_{t+1}(x) = (1 - \alpha_t(x))\Delta_t(x) + \alpha_t(x)F_t(x) \quad (7)$$

where  $x \in X$ ,  $t = 0, 1, \dots, \infty$ , converges to zero with probability one (w. p. 1) if the following properties hold:

1. The set of possible states  $X$  is finite.
2.  $0 \leq \alpha_t(x) \leq 1$ ,  $\sum_t \alpha_t(x) = \infty$ ,  $\sum_t \alpha_t^2(x) < \infty$  w. p. 1, where the probability is over the learning rates  $\alpha_t$ .
3.  $\|\mathbb{E}\{F_t(x)|\mathcal{P}_t\}\|_W \leq \mathcal{K}\|\Delta_t\|_W + c_t$ , where  $\mathcal{K} \in [0, 1)$  and  $c_t$  converges to zero w. p. 1.
4.  $\text{var}\{F_t(x)|\mathcal{P}_t\} \leq K(1 + \|\Delta_t\|_W)^2$ , where  $K$  is some constant. Here  $\mathcal{P}_t$  is an increasing sequence of  $\sigma$ -fields that includes the past of the process. In particular, we assume that  $\alpha_t, \Delta_t, F_{t-1} \in \mathcal{P}_t$ . The notation  $\|\cdot\|_W$  refers to some (fixed) weighted maximum norm and the notation  $\text{var}$  refers to the variance.

PROOF. Refer to Theorem 1 in Jaakola et al. [14] for proof.  $\square$

Across this section, since we are only focusing on the low- $Q$  values, with a small abuse of notation, we will use  $Q$  to denote the low- $Q$  values. Now, we define a Nash operator  $P_t$ , using the following equation,

$$P_t Q^k(s, \mathbf{a}) = \mathbb{E}_{s' \sim p}[r_t^k(s, \mathbf{a}) + \gamma \pi_*^1(s') \cdots \pi_*^n(s') Q^k(s')] \quad (8)$$

where  $s'$  is the state at time  $t + 1$ ,  $(\pi_*^1(s'), \dots, \pi_*^n(s'))$  is the Nash equilibrium solution for the stage game  $(Q^1(s'), \dots, Q^n(s'))$ , and  $p$  is the transition function.  $Q^k$  denotes the  $Q$ -value of a representative agent  $k$ .

LEMMA 2. *Under Assumption 3, the Nash operator as defined in Eq. 8 forms a contraction mapping with the fixed point being the Nash  $Q$ -value of the game.*

PROOF. See Theorem 17 of Hu and Wellman [13].  $\square$

Now, since the  $P_t$  operator forms a contraction mapping,  $\|P_t Q^j - P_t Q_*^j\| \leq \gamma \|Q^j - Q_*^j\|$ , is satisfied for some  $\gamma \in [0, 1)$  and all  $Q^j$ . Here  $Q_*^j$  is the Nash  $Q$ -value of the agent  $j$ .

The objective is to apply Lemma 1 to show that the low- $Q$  in MATLQL converge to the Nash  $Q$  values.

The first two conditions of Lemma 1 are satisfied from the Assumption 1 and Assumption 2. Now, comparing Eq. 7 and Eq. 3 we get that  $x$  can be associated with the state joint action pairs  $(s, \mathbf{a})$  and  $\Delta_t(s_t, \mathbf{a}_t)$  can be associated with  $Q_t^j(s, \mathbf{a}) - Q_*^j(s, \mathbf{a})$ . Here,  $Q_*^j(s, \mathbf{a})$  is the Nash  $Q$  value of the agent  $j$ .

Now we get

$$\Delta_{t+1}(x) = (1 - \alpha_t(x))\Delta_t(x) + \alpha_t(x)F_t(x), \quad (9)$$

where

$$\begin{aligned} F_t(x) &= r_t^j + \gamma v^{Nash,j}(s_{t+1}) - Q_*^j(s_t, \mathbf{a}_t) \\ &\quad + \gamma [\max_{\mathbf{a}'} Q_t^j(s_{t+1}, \mathbf{a}_{t+1}) - v^{Nash,j}(s_{t+1})] \\ &\triangleq r_t^j + \gamma v^{Nash,j}(s_{t+1}) - Q_*^j(s_t, \mathbf{a}_t) + C_t^j(s_t, \mathbf{a}_t) \\ &\triangleq F_t^{Q,j}(s_t, \mathbf{a}_t) + C_t^j(s_t, \mathbf{a}_t) \end{aligned} \quad (10)$$

The Nash value function  $v^{Nash,j}(s)$  of an agent  $j$  is defined as the expected cumulative discounted future rewards obtained by the agent  $j$ , given that all agents follow the Nash policy  $\pi_*$ .

Here, we set,  $F_t(s_t, \mathbf{a}_t) = F_t^{Q,j}(s_t, \mathbf{a}_t) = C_t^j(s_t, \mathbf{a}_t) = 0$  if  $(s, \mathbf{a}) \neq (s_t, \mathbf{a}_t)$ . Let the  $\sigma$ -field generated by all the random variables provided by  $(s_t, \alpha_t, a_t^1, \dots, a_t^n, r_{t-1}, \dots, s_1, \alpha_1, a_1, Q_0)$ , be represented by  $\mathcal{P}_t$ . Now, all the  $Q$ -values are  $\mathcal{P}_t$  measurable which makes  $\Delta_t$  and  $F_t$ ,  $\mathcal{P}_t$  measurable and this satisfies the measurability condition of Lemma 1.

Hu and Wellman [13] proved that the result,  $v^{Nash,j}(s_{t+1}) \triangleq v^j(s', \pi_*^1, \dots, \pi_*^n) = \pi_*^1(s') \cdots \pi_*^n(s') Q_*^j(s')$  holds (see the proof in Lemma 10 of [13]). Hence, from Lemma 2, we can show that the  $\mathbb{E}[F_t^Q]$  forms a contraction mapping. This can be done using the fact that  $\mathbb{E}(P_t Q_*) = Q_*$  (refer to Lemma 11 in [13]). Here, the norm is the maximum norm on the joint action.

Now, we have the following for all  $t$ ,

$$\|\mathbb{E}[F_t^{Q,j}(s_t, \mathbf{a}_t)|\mathcal{P}_t]\| \leq \gamma \|Q_t^j(s_t, \mathbf{a}_t) - Q_*^j(s_t, \mathbf{a}_t)\| = \gamma \|\Delta_t\| \quad (11)$$

Now from Eq. 10,

$$\begin{aligned} &\|\mathbb{E}[F_t(s_t, \mathbf{a}_t)|\mathcal{P}_t]\| \\ &\leq \|\mathbb{E}[F_t^{Q,j}(s_t, \mathbf{a}_t)|\mathcal{P}_t]\| + \|\mathbb{E}[C_t^j(s_t, \mathbf{a}_t)|\mathcal{P}_t]\| \\ &\leq \gamma \|\Delta_t\| + \|\mathbb{E}[C_t^j(s_t, \mathbf{a}_t)|\mathcal{P}_t]\| \end{aligned} \quad (12)$$

This satisfies the third condition of Lemma 1, provided that  $c_t = \|\mathbb{E}[C_t^j(s_t, \mathbf{a}_t)|\mathcal{P}_t]\|$  converges to 0 with probability 1 (w. p. 1). From the definition of  $C_t^j$  and the Assumption 3, it can be shown that the value of  $C_t^j$  converges to 0 in the limit of time (see Theorem 3 in Subramanian et al. [37]).

Thus, it follows from Lemma 1 that the process  $\Delta_t$  converges to 0 and hence, low- $Q$  value for an agent  $j$ , converges to Nash  $Q$  value  $Q_*^j$ .  $\square$

## C PROOF OF THEOREM 2

In this section, we give the proof for Theorem 2. For providing this bound we use the notion of covering time  $L$ . The covering time means that within  $L$  steps from any start state, all state-joint action pairs are performed at least once by all agents. Similar to Dar and Mansour [5], we clarify that we do not need to assume that the state-joint action pairs are being generated by any particular strategy. Same as in Appendix B, across this section, we will use  $Q$  to denote the low- $Q$  values. For a representative agent  $j$ , we will focus on the value of  $r_t^j = \|Q_t^j - Q_*^j\|$  and the aim is to bound the time until  $r_t^j \leq \epsilon$ . Here the norm denotes the maximum difference of the  $Q$  values across all states and joint actions. The proof of this



theorem follows the Theorem 4 in Dar and Mansour [5]. While the work of Dar and Mansour was restricted to single-agent MDPs, our result extends the analysis of Dar and Mansour to the general-sum stochastic game setting.

In line with the Eq. 4, let us consider a stochastic iterative process of the form,

$$X_{t+1}^j(i) = (1 - \alpha_t(i))X_t^j(i) + \alpha_t(i)((H_t X_t^j)(i) + w_t^j(i)). \quad (13)$$

As mentioned in Section 5, let us specify  $\beta = \frac{1-\gamma}{2}$ . Also, consider a constant  $Q_{\max}^j$  where the  $Q_{\max}^j$  denotes the maximum low- $Q$  value possible to be obtained in the stochastic game by the agent  $j$ . Hence, the relation  $\|X_0\| \leq Q_{\max}^j$  holds. Further, let us consider a sequence  $D_k^j$ , with  $D_1^j = Q_{\max}^j$  and  $D_{k+1}^j = (1 - \beta)D_k^j$  for all  $k \geq 1$ . By the nature of this construction, the sequence  $D_k^j$  is guaranteed to converge to 0, since at each step the value of  $D_k^j$  is being continuously multiplied by a fractional value. Now we can prove the following result.

LEMMA 3. For every  $k$ , there exists a time  $\tau_k$  such that, for any  $t \geq \tau_k$  we have  $\|X_t^j\| \leq D_k^j$ .

PROOF. See Theorem 8 in Dar and Mansour [5].  $\square$

The Lemma 3 guarantees that at time  $t \geq \tau_k$ , for any  $i$  the value of  $\|X_t^j(i)\|$  is in the interval  $[-D_k^j, D_k^j]$ .

We can state the following lemma, where we bound the number of iterations until  $D_i^j \leq \epsilon$ .

LEMMA 4. For  $m \geq \frac{1}{\beta} \ln(Q_{\max}^j/\epsilon)$  we have  $D_m^j \leq \epsilon$ .

PROOF. Since we have that  $D_1^j = Q_{\max}^j$  and  $D_i^j = (1 - \beta)D_{i-1}^j$ , we will need a  $m$  that satisfies  $D_m^j = Q_{\max}^j(1 - \beta)^m \leq \epsilon$ . By taking a logarithm on both sides, we get

$$\begin{aligned} \ln((Q_{\max}^j)(1 - \beta)^m) &\leq \ln(\epsilon) \\ \ln(Q_{\max}^j) + m \ln(1 - \beta) &\leq \ln(\epsilon) \\ \implies \ln(Q_{\max}^j) - \ln(\epsilon) &\leq -m \ln(1 - \beta) \\ \implies \ln(Q_{\max}^j/\epsilon) &\leq m \sum_{k=0}^{\infty} \beta^k / k \\ \implies 1/\beta \ln(Q_{\max}^j/\epsilon) &\leq m. \end{aligned} \quad (14)$$

In the last step we omit the higher powers since  $\beta$  is a small fraction. This proves the result.  $\square$

Now, we define a sequence of times  $\tau_k$ , with reference to the MA-TLQL low-level  $Q$ -updates with a polynomial learning rate. Let us define  $\tau_{k+1} = \tau_k + L\tau_k^\omega$ . Here the term  $\omega$  denotes the decay factor of the learning rate with  $\alpha_t^\omega = \frac{1}{(t+1)^\omega}$  for  $\omega \in (1/2, 1)$ . The term  $L\tau_k^\omega$  specifies the number of steps needed to update each state joint action pair at least  $\tau_k^\omega$  times. The time between the  $\tau_k$  and  $\tau_{k+1}$  is denoted as the  $k^{\text{th}}$  iteration. Now we provide a definition for the number of times a state joint action pair is visited.

DEFINITION 1. Let  $n(s, \mathbf{a}, t_1, t_2)$  be the number of times that the state joint action pair  $(s, \mathbf{a})$  was performed in the time step interval  $[t_1, t_2]$ .

Before providing the suitable bounds, we would like to provide some equations that relate the stochastic iterative technique given in Eq. 13 with the  $Q$ -update in Eq. 4.

First we provide a formal definition of a stochastic game, that will be useful for our further analysis. Consider a stochastic game that can be defined as follows:

DEFINITION 2. A stochastic game is defined as  $\langle S, N, \mathbf{A}, P, \mathbf{R}, \gamma \rangle$  where  $S$  is a finite set of states,  $N$  is the finite set of agents,  $|N| = n$ , and  $\mathbf{A} = A^1 \times \dots \times A^n$  is the set of joint actions, where  $A^j$  is the finite action set of an agent  $j$ , and  $\mathbf{a} = (a^1, \dots, a^n) \in \mathbf{A}$  is the joint action where an agent  $j$  takes action  $a^j \in A^j$ . Furthermore,  $P_{i,k}(\mathbf{a}) : S \times \mathbf{A} \times S \mapsto [0, 1]$  is the transition function that provides the probability of reaching state  $k$  from state  $i$  when all agents are performing the joint action  $\mathbf{a} \in \mathbf{A}$  in state  $i$ ,  $\mathbf{R}(s, \mathbf{a}) = \{R^1(s, \mathbf{a}), \dots, R^n(s, \mathbf{a})\}$  is the set of reward functions, where  $R^j(s, \mathbf{a}) : S \times \mathbf{A} \mapsto \mathbb{R}^n$  is the reward function of the agent  $j$ , and  $\gamma$  is the discount factor satisfying  $0 \leq \gamma < 1$ .

Towards the same, we define an operator  $H$  that can be represented as

$$(HQ^j)(i, \mathbf{a}) = \sum_{k=0}^{|S|} P_{ik}(\mathbf{a})(R^j(i, \mathbf{a}) + \gamma \max_{b^j \in A^j} Q^j(k, \mathbf{b})). \quad (15)$$

Here the  $\mathbf{b} = \{b^j, b^{-j}\}$ , where  $b^{-j}$  denotes the joint action of all agents except the agent  $j$ .

Rewriting the  $Q$ -function with  $H$  we get,

$$Q_{t+1}^j(i, \mathbf{a}) = (1 - \alpha_t(i, \mathbf{a}))Q_t^j(i, \mathbf{a}) + \alpha_t(i, \mathbf{a})(H Q_t^j)(i, \mathbf{a}) + w_t^j(i, \mathbf{a}). \quad (16)$$

Let  $\bar{i}$  be the state that is reached by performing joint action  $\mathbf{a}$  at time  $t$  in state  $i$  and  $r^j(i, \mathbf{a})$  be the reward observed by the agent  $j$  at state  $i$ ; then

$$\begin{aligned} w_t^j(i, \mathbf{a}) &= r^j(i, \mathbf{a}) + \gamma \max_{b^j \in A^j} Q_t^j(\bar{i}, \mathbf{b}) \\ &\quad - \sum_{k=0}^{|S|} P_{ik}^j(\mathbf{a})(R^j(i, \mathbf{a}) + \gamma \max_{b^j \in A^j} Q_t^j(k, \mathbf{b})). \end{aligned} \quad (17)$$

From this construction  $w_t^j$  is bounded by  $Q_{\max}^j$  for all  $t$  and has zero expectation. Further we will define two other sequences  $W_{t;\tau}^j$  and  $Y_{t;\tau}^j$  where  $\tau$  represents some initial time. These are given by the following equations.

$$W_{t+1;\tau_k}^j(s, \mathbf{a}) = (1 - \alpha_t^\omega(i))W_{t;\tau_k}^j(s, \mathbf{a}) + \alpha_t^\omega(i)w_t^j(s, \mathbf{a}) \quad (18)$$

where  $W_{\tau_k;\tau_k}^j(s, \mathbf{a}) = 0$ . The value of  $W_{t;\tau_k}^j$  bounds the contributions of  $w_t^j(s, \mathbf{a})$ , to the value of  $Q_t^j$ , starting from an arbitrary  $\tau_k$ . Now we have

$$Y_{t+1;\tau_k}^j(s, \mathbf{a}) = (1 - \alpha_t^\omega(i))Y_{t;\tau_k}^j(s, \mathbf{a}) + \alpha_t^\omega(s, \mathbf{a})\gamma D_k^j \quad (19)$$

where  $Y_{\tau_k;\tau_k}^j = D_k^j$ .

Next, we can state a lemma that will bound the  $Q$ -functions w.r.t the sequences  $W_{t;\tau_k}^j$  and  $Y_{t;\tau_k}^j$ .

LEMMA 5. For every state  $s$  and joint action  $\mathbf{a}$  and time  $\tau_k$ , we have

$$\begin{aligned} & -Y_{t;\tau_k}^j(s, \mathbf{a}) + W_{t;\tau_k}^j(s, \mathbf{a}) \\ & \leq Q_{\star}^j(s, \mathbf{a}) - Q_t^j(s, \mathbf{a}) \\ & \leq Y_{t;\tau_k}^j(s, \mathbf{a}) + W_{t;\tau_k}^j(s, \mathbf{a}) \end{aligned} \quad (20)$$

PROOF. See Lemma 4.4 in Bertsekas and Tsitsiklis [3].  $\square$

From the Lemma 5 we see that a bound on the difference  $r_t^j$  depends on the bound for  $Y_{t;\tau_k}^j$  and  $W_{t;\tau_k}^j$ . So we can bound  $Y_{t;\tau_k}^j$  and  $W_{t;\tau_k}^j$  separately and two bounds together will provide a bound for  $r_t^j$ .

We first provide a result on the nature of the sequence  $Y_{\tau_k;\tau_k}^j$ .

LEMMA 6. The sequence  $Y_{t;\tau_k}^j$  is a monotonically decreasing sequence.

PROOF. From Eq. 19 we can write (subtract  $\gamma D_k^j$  from both sides),

$$\begin{aligned} & (Y_{t+1;\tau_k}^j(s, \mathbf{a}) - \gamma D_k^j) \\ & = (1 - \alpha_t^\omega(i))Y_{t;\tau_k}^j(s, \mathbf{a}) + (1 - \alpha_t^\omega(s, \mathbf{a}))\gamma D_k^j \\ & = (1 - \alpha_t^\omega(i))(Y_{t;\tau_k}^j(s, \mathbf{a}) + \gamma D_k^j). \end{aligned} \quad (21)$$

Now, the convergence of  $\|Y_{t+1;\tau_k}^j(s, \mathbf{a}) - \gamma D_k^j\|$  follows since  $\lim_{n \rightarrow \infty} \prod_{t=k}^n (1 - \alpha_t^\omega(s, \mathbf{a})) = 0$ . This shows that the sequence  $(Y_{t+1;\tau_k}^j(s, \mathbf{a}) - \gamma D_k^j)$  monotonically decreases to 0 and hence the sequence  $Y_{t;\tau_k}^j$  monotonically decreases to  $\gamma D_k^j$ . This proves our result.  $\square$

Next, we provide a bound on the value of  $Y_{t;\tau_k}^j$ .

LEMMA 7. Consider the low- $Q$  update given in Eq. 4, with a polynomial learning rate and assume that for any  $t \geq \tau_k$  we have  $Y_{t;\tau_k}^j(s, \mathbf{a}) \leq D_k$ . Then for any  $t \geq \tau_k + L\tau_k^\omega = \tau_{k+1}$  we have  $Y_{t;\tau_k}^j(s, \mathbf{a}) \leq D^j(\gamma + \frac{2}{e}\beta)$ .

PROOF. For each state-joint action pair  $s, \mathbf{a}$  we are assured that  $n(s, \mathbf{a}, \tau_k, \tau_{k+1}) \geq \tau_k^\omega$ , since the covering time is  $L$  and the underlying policy has made  $L\tau_k^\omega$  steps (since we have the relation  $\tau_{k+1} = \tau_k + L\tau_k^\omega$ ).

Let  $Y_{\tau_k;\tau_k}^j(s, \mathbf{a}) = \gamma D_k^j + \rho_{\tau_k}^j$ , where  $\rho_{\tau_k} = (1 - \gamma)D_k^j$ . Now we have the following expression,

$$Y_{t+1;\tau_k}^j(s, \mathbf{a}) = (1 - \alpha_t^\omega)Y_{t;\tau_k}^j(s, \mathbf{a}) + \alpha_t^\omega \gamma D_k^j = \gamma D_k^j + (1 - \alpha_t^\omega)\rho_t^j \quad (22)$$

where  $\rho_{t+1}^j = \rho_t^j(1 - \alpha_t^\omega)$ . We aim to show that after time  $\tau_{k+1} = \tau_k + \tau_k^\omega$ , for any  $t \geq \tau_k^\omega$  for any  $t \geq \tau_{k+1}$  we have  $\rho_t^j \leq \frac{2}{e}\beta D_k^j$ . By definition, we can rewrite  $\rho_t^j$  as,

$$\begin{aligned} \rho_t^j & = (1 - \gamma)D_k^j \prod_{l=1}^{t-\tau_k} (1 - \alpha_{l+\tau_k}^\omega) \\ & = 2\beta D_k^j \prod_{l=1}^{t-\tau_k} (1 - \alpha_{l+\tau_k}^\omega) \\ & = 2\beta D_k^j \prod_{l=1}^{t-\tau_k} (1 - \frac{1}{|n(s, \mathbf{a}, \tau_{k+1}, l)|^\omega}) \end{aligned} \quad (23)$$

the last identity follows from the definition of  $\alpha_t^\omega$ .

Since the  $\tau_k^\omega$ 's are monotonically decreasing,

$$\rho_t^j \leq 2\beta D_k^j (1 - \frac{1}{\tau_k^\omega})^{t-\tau_k}. \quad (24)$$

For  $t \geq \tau_k + \tau_k^\omega$  we have,

$$\rho_t^j \leq 2\beta D_k^j \left(1 - \frac{1}{\tau_k^\omega}\right)^{\tau_k^\omega} \leq \frac{2}{e}\beta D_k^j. \quad (25)$$

The last step is obtained from the fact that  $\lim_{x \rightarrow \infty} (1 - (1/x))^x = 1/e$ .

Hence,  $Y_{t;\tau_k}(s, \mathbf{a}) \leq (\gamma + \frac{2}{e}\beta)D_k$ .  $\square$

From Lemma 7 we have provided a bound for the term  $Y_{t;\tau_k}^j$  for  $t = \tau_{k+1}$  which will automatically hold for all  $t \geq \tau_{k+1}$ , since the term  $Y_{t;\tau_k}^j$  is monotonically decreasing (from Lemma 6) and deterministic (from Eq. 19 and Lemma 5).

Next we bound the term  $W_{t;\tau_k}^j$  by  $(1 - \frac{2}{e})\beta D_k^j$ . The sum of the bounds for  $W_{t;\tau_k}^j(s, \mathbf{a})$  and  $Y_{t;\tau_k}^j(s, \mathbf{a})$ , would be  $(\gamma + \beta)D_k^j = (1 - \beta)D_k^j = D_{k+1}^j$ , as desired.

Now we state a definition for a sequence  $\eta_i^{k,t,j}(s, \mathbf{a})$  and  $W_{t;\tau_k}^{l,j}$ .

DEFINITION 3. Let

$$\begin{aligned} W_{t;\tau_k}^j(s, \mathbf{a}) & = (1 - \alpha_t^\omega(s, \mathbf{a}))W_{t-1;\tau_k}^j(s, \mathbf{a}) + \alpha_t^\omega(s, \mathbf{a})w_t^j(s, \mathbf{a}) \\ & = \sum_{i=1}^{t-\tau_k} \eta_{i+\tau_k}^{k,t,j}(s, \mathbf{a})w_{i+\tau_k}^j(s, \mathbf{a}), \end{aligned}$$

where

$$\eta_{i+\tau_k}^{k,t,j}(s, \mathbf{a}) = \alpha_{i+\tau_k}^\omega(s, \mathbf{a}) \prod_{l=\tau_k+i+1}^t (1 - \alpha_l^\omega(s, \mathbf{a})). \quad (26)$$

For bounding the sequence  $W_{t;\tau_k}^j$  we consider the interval  $t \in [\tau_{k+1}, \tau_{k+2}]$ . First we provide a lemma that bounds the coefficients in this interval and bounds the influence of  $w_t^j(s, \mathbf{a})$  in this interval.

LEMMA 8. Let  $\tilde{w}_{i+\tau_k}^{k,t,j}(s, \mathbf{a}) = \eta_{i+\tau_k}^{k,t,j}(s, \mathbf{a})w_{i+\tau_k}^j(s, \mathbf{a})$ , then for any  $t \in [\tau_{k+1}, \tau_{k+2}]$  the random variable  $\tilde{w}_{i+\tau_k}^{k,t,j}(s, \mathbf{a})$  has zero mean and bounded by  $(L/\tau_k)^\omega Q_{\max}^j$ .

PROOF. Since  $\eta_{i+\tau_k}^{k,t,j}(s, \mathbf{a}) = \alpha_{i+\tau_k}^\omega(s, \mathbf{a}) \prod_{l=\tau_k+i+1}^t (1 - \alpha_l^\omega(s, \mathbf{a}))$ , we can divide  $\eta_{i+\tau_k}^{k,t,j}$  into two parts, the first  $\alpha_{i+\tau_k}^\omega$  and the second  $\mu = \prod_{l=\tau_k+i+1}^t (1 - \alpha_l^\omega)$ . Since,  $\mu$  is bounded from above by 1, we have,

$$\begin{aligned} \eta_{i+\tau_k}^{k,t,j}(s, \mathbf{a}) & \leq \alpha_{i+\tau_k}^\omega(s, \mathbf{a}) \\ & = \frac{1}{|n(s, \mathbf{a}, i+\tau_k)|} \leq \left(\frac{L}{i+\tau_k}\right)^\omega \leq \left(\frac{L}{\tau_k}\right)^\omega. \end{aligned} \quad (27)$$

Here, the  $*$  is from the fact that in a time interval of  $\tau$ , each state-joint action pair is performed at-least  $\tau/L$  times by the definition of covering time.

Hence, we get the relation that  $\eta_{i+\tau_k}^{k,t,j}(s, \mathbf{a}) \leq (L/\tau_k)^\omega$ .

Now, consider the expectation of  $\tilde{w}_{i+\tau_k}(s, \mathbf{a})$ . By definition, we have that,  $w_{\tau_k+i}^j(s, \mathbf{a})$  has zero mean and is bounded by  $Q_{\max}^j$  for any history and state-joint action pair, hence,

$$\begin{aligned} & \mathbb{E}[\tilde{w}_{i+\tau_k}(s, \mathbf{a})] \\ &= \mathbb{E}[\eta_{i+\tau_k}^{k,t,j}(s, \mathbf{a}) w_{i+\tau_k}(s, \mathbf{a})] \\ &= \eta_{i+\tau_k}^{k,t,j}(s, \mathbf{a}) \mathbb{E}[w_{i+\tau_k}(s, \mathbf{a})] = 0. \end{aligned} \quad (28)$$

Next we can prove that it is bounded as well,

$$\begin{aligned} & |\tilde{w}_{i+\tau_k}(s, \mathbf{a})| \\ &= |\eta_{i+\tau_k}^{k,t,j}(s, \mathbf{a}) w_{i+\tau_k}(s, \mathbf{a})| \\ &\leq |\eta_{i+\tau_k}^{k,t,j}(s, \mathbf{a})| Q_{\max}^j \\ &\leq (L/\tau_k)^\omega Q_{\max}^j \end{aligned} \quad (29)$$

□

Now, let us define  $W_{t;\tau_k}^{l,j}(s, \mathbf{a}) = \sum_{i=1}^l \tilde{w}_{i+\tau_k}^t(s, \mathbf{a})$ . The objective is to prove that this is a martingale difference sequence having bounded differences.

LEMMA 9. For any  $t \in [\tau_{k+1}, \tau_{k+2}]$  and  $1 \leq l \leq t$  we have that  $W_{t;\tau_k}^{l,j}(s, \mathbf{a})$  is a martingale sequence that satisfies,

$$|W_{t;\tau_k}^{l,j}(s, \mathbf{a}) - W_{t;\tau_k}^{l-1,j}(s, \mathbf{a})| \leq (L/\tau_k)^\omega Q_{\max}^j \quad (30)$$

PROOF. We first note that the term  $W_{t;\tau_k}^{l,j}(s, \mathbf{a})$  is a martingale sequence, since

$$\begin{aligned} & \mathbb{E}[W_{t;\tau_k}^{l,j}(s, \mathbf{a}) - W_{t;\tau_k}^{l-1,j}(s, \mathbf{a}) | F_{\tau_{k+1}-1}] \\ &= \mathbb{E}[\tilde{w}_{l+\tau_k}^{j,t}(s, \mathbf{a}) | F_{\tau_{k+1}-1}] = 0 \end{aligned} \quad (31)$$

where the variable  $F_{\tau_{k+1}-1}$  denotes all previous values of  $W^l$ .

Also, by Lemma 8 we can show that  $\tilde{w}_{l+\tau_k}^j(s, \mathbf{a})$  is bounded by  $(L/\tau_k)^\omega Q_{\max}^j$ , thus

$$\begin{aligned} & |W_{t;\tau_k}^{l,j}(s, \mathbf{a}) - W_{t;\tau_k}^{l-1,j}(s, \mathbf{a})| \\ &= \tilde{w}_{l+\tau_k}^j(s, \mathbf{a}) \leq (L/\tau_k)^\omega Q_{\max}^j \end{aligned} \quad (32)$$

□

The next lemma bounds the term  $W_{t;\tau_k}^j$ .

LEMMA 10. Consider the low- $Q$  update given in Eq. 4, with a polynomial learning rate. With probability at least  $1 - \frac{\delta}{m}$  we have that, for every state-joint action pair  $|W_{t;\tau_k}^j(s, \mathbf{a})| \leq (1 - \frac{2}{e})\gamma D_k$  for any  $t \in [\tau_{k+1}, \tau_{k+2}]$ , i.e.

$$\Pr\left[\forall s, \mathbf{a} \forall t \in [\tau_{k+1}, \tau_{k+2}] : |W_{t;\tau_k}(s, \mathbf{a})| \leq (1 - \frac{2}{e})\beta D_k\right] \geq 1 - \frac{\delta}{m} \quad (33)$$

given that

$$\tau_k = \Theta\left(\left(\frac{L^{1+3\omega} Q_{\max}^{2,j} \ln(Q_{\max}^j |S| |\Pi_i| |A| m / (\delta \beta D_k))}{\beta^2 D_k^2}\right)^{1/\omega}\right) \quad (34)$$

PROOF. For each state-joint action pair comparing  $W_{t;\tau_k}^{l,j}(s, \mathbf{a})$  and  $W_{t;\tau_k}^j(s, \mathbf{a})$  we note that  $W_{t;\tau_k}^j(s, \mathbf{a}) = W_{t;\tau_k}^{t-\tau_k+1,j}(s, \mathbf{a})$ .

Let  $l = n(s, \mathbf{a}, \tau_k, t)$ , then for any  $t \in [\tau_{k+1}, \tau_{k+2}]$  we have that

$$\begin{aligned} l &\leq \tau_{k+2} - \tau_k = \tau_{k+1} + L\tau_{k+1}^\omega - \tau_k \\ &= \tau_k + L\tau_k^\omega + L(\tau_k + L\tau_k^\omega) - \tau_k \\ &= L\tau_k^\omega + L\tau_k + L^2\tau_k^\omega \leq \Theta(L^{1+\omega}\tau_k^\omega). \end{aligned} \quad (35)$$

By Lemma 9 we can apply Azuma's inequality to  $W_{t;\tau_k}^{t-\tau_k+1,j}(s, \mathbf{a})$  with  $c_i = (L/\tau_k)^\omega Q_{\max}^j$ . Therefore, we can derive that

$$\begin{aligned} & \Pr\left[|W_{t;\tau_k}(s, \mathbf{a})| \geq \tilde{\epsilon} | t \in [\tau_{k+1}, \tau_{k+2}] \right] \\ &\leq 2e^{-\frac{\tilde{\epsilon}^2}{2 \sum_{i=\tau_k+1}^t c_i^2}} \leq 2e^{-\frac{\tilde{\epsilon}^2}{2lc_i^2}} \\ &\leq 2e^{-\frac{\tilde{\epsilon}^2 \tau_k^{2\omega}}{1L^{2\omega} Q_{\max}^{2,j}}} \\ &\leq 2e^{-\frac{\tilde{\epsilon}^2 \tau_k^\omega}{L^{1+3\omega} Q_{\max}^{2,j}}} \end{aligned} \quad (36)$$

for some constant  $c > 0$ . We can set  $\tilde{\delta}_k = 2e^{-\frac{c\tau_k^\omega \tilde{\epsilon}^2}{L^{1+3\omega} Q_{\max}^{2,j}}}$ , which holds for  $\tau_k^\omega = \Theta(\ln(1/\tilde{\delta}_k) L^{1+3\omega} Q_{\max}^{2,j} / \tilde{\epsilon}^2)$ .

Using the union bound we get,

$$\begin{aligned} & \Pr[\forall s, \forall \mathbf{a}, \forall t \in [\tau_{k+1}, \tau_{k+2}] : |W_{t;\tau_k}^j(s, \mathbf{a})| \leq \tilde{\epsilon}] \\ &\geq 1 - \sum_{t=\tau_{k+1}}^{\tau_{k+2}} \Pr[\forall s, \forall \mathbf{a}, |W_{t;\tau_k}^j(s, \mathbf{a})| \geq \tilde{\epsilon}] \\ &\geq 1 - \sum_{t=\tau_{k+1}}^{\tau_{k+2}} \Pr[\forall s, \forall \mathbf{a}, |W_{t;\tau_k}^j(s, \mathbf{a})| \geq \tilde{\epsilon}] \\ &\geq 1 - \sum_{t=\tau_{k+1}}^{\tau_{k+2}} \tilde{\delta}_k |S| |\Pi_i| |A|_i \\ &\geq 1 - (\tau_{k+2} - \tau_{k+1}) \tilde{\delta}_k |S| |\Pi_i| |A|_i \end{aligned} \quad (37)$$

We would like to set a level of probability of  $1 - \frac{\delta}{m}$ , for each state-joint action pair. From the above equation we get,

$$\begin{aligned} & 1 - (\tau_{k+2} - \tau_{k+1}) \tilde{\delta}_k |S| |\Pi_i| |A|_i = 1 - \delta/m \\ &\implies (\tau_{k+2} - \tau_{k+1}) \tilde{\delta}_k |S| |\Pi_i| |A|_i = \delta/m. \\ &\implies \tilde{\delta}_k = \delta / (\tau_{k+2} - \tau_{k+1}) m |S| |\Pi_i| |A|_i. \end{aligned} \quad (38)$$

Thus taking  $\tilde{\delta}_k = \frac{\delta}{m(\tau_{k+2} - \tau_{k+1}) |S| |\Pi_i| |A|_i}$  assures  $1 - \frac{\delta}{m}$  for each state-joint action pair. As a result we have,

$$\begin{aligned} \tau_k^\omega &= \Theta\left(\frac{L^{1+3\omega} Q_{\max}^{2,j} \ln(|S| |\Pi_i| |A| m \tau_k^\omega / \delta)}{\tilde{\epsilon}^2}\right) \\ &= \Theta\left(\frac{L^{1+3\omega} Q_{\max}^{2,j} \ln(|S| |\Pi_i| |A| m Q_{\max}^j / (\delta \tilde{\epsilon}))}{\tilde{\epsilon}^2}\right) \end{aligned} \quad (39)$$

Setting  $\tilde{\epsilon} = (1 - 2/e)\beta D_k$  gives the desired bound.  $\square$

Now that we have bounded for each iteration the time needed to achieve the desired probability  $1 - \frac{\delta}{m}$ . The next lemma provides a bound for error in all iterations.

LEMMA 11. Consider the low- $Q$  update given in Eq. 4, with a polynomial learning rate. With probability  $1 - \delta$ , for every iteration  $k \in [1, m]$  and time  $t \in [\tau_{k+1}, \tau_{k+2}]$  we have  $|W_{t;\tau_k}^j(s, \mathbf{a})| \leq (1 - \frac{2}{e})\beta D_k^j$ , i.e.,

$$\Pr\left[\forall k \in [1, m], \forall t \in [\tau_{k+1}, \tau_{k+2}], \forall s, \mathbf{a} : |W_{t;\tau_k}^j(s, \mathbf{a})| \leq (1 - \frac{2}{e})\beta D_k^j\right] \geq 1 - \delta \quad (40)$$

given that

$$\tau_0 = \Theta\left(\left(\frac{L^{1+3\omega} Q_{\max}^{2,j} \ln(Q_{\max}^j |S| |\Pi_i| |A|_i m / (\delta \beta \epsilon))}{\beta^2 \epsilon^2}\right)^{1/\omega}\right) \quad (41)$$

PROOF. From Lemma 10 we have that

$$\Pr\left[\forall t \in [\tau_{k+1}, \tau_{k+2}] : |W_{t;\tau_k}^j| \geq (1 - \frac{2}{e})\beta D_k^j\right] \leq \frac{\delta}{m} \quad (42)$$

Using the union bound we have that

$$\begin{aligned} \Pr[\forall k \leq m, \forall t \in [\tau_{k+1}, \tau_{k+2}] |W_{t;\tau_k}^j| \geq \tilde{\epsilon}] \\ \leq \sum_{k=1}^m \Pr[\forall t \in [\tau_{k+1}, \tau_{k+2}] |W_{t;\tau_k}^j| \geq \tilde{\epsilon}] \leq \delta \end{aligned} \quad (43)$$

where  $\tilde{\epsilon} = (1 - \frac{2}{e})\beta D_k$ .  $\square$

The next lemma solves the recurrence  $\sum_{i=0}^{m-1} L\tau_i^\omega + \tau_0$  and derives the time complexity.

LEMMA 12. Let

$$a_{k+1} = a_k + L a_k^\omega = a_0 + \sum_{i=0}^k L a_i^\omega \quad (44)$$

Then for any constant  $\omega \in (0, 1)$ ,  $a_k = \Omega(a_0^{1-\omega}/L + L^{\frac{1}{1-\omega}}((k+1)/2)^{\frac{1}{1-\omega}})$ .

PROOF. Let us define the following series

$$b_{k+1} = \sum_{i=0}^k L b_i^\omega + b_0 \quad (45)$$

with an initial condition  $b_0 = L^{\frac{1}{1-\omega}}$ .

Now we lower bound  $b_k$  by  $(L(k+1)/2)^{\frac{1}{1-\omega}}$ . We use induction to prove this hypothesis. For  $k = 0$ ,

$$b_0 = L^{\frac{1}{1-\omega}} \geq \left(\frac{L}{2}\right)^{\frac{1}{1-\omega}} \quad (46)$$

Assume that the induction hypothesis holds for  $k-1$  and prove for  $k$ ,

$$\begin{aligned} b_k &= b_{k-1} + L b_{k-1}^\omega \\ &= (Lk/2)^{\frac{1}{1-\omega}} + L(Lk/2)^{\frac{\omega}{1-\omega}} \\ &= L^{\frac{1}{1-\omega}}((k/2)^{\frac{1}{1-\omega}} + (k/2)^{\frac{\omega}{1-\omega}}) \\ &\geq L^{\frac{1}{1-\omega}}((k+1)/2)^{\frac{1}{1-\omega}}. \end{aligned} \quad (47)$$

For  $a_0 \geq L^{\frac{1}{1-\omega}}$  we can view the series as starting at  $b_k = a_0$ . Since the first time step is 1, we can see that the start point has moved  $\Theta(a_0^{1-\omega}/L)$ . Therefore, we have a total complexity of  $\Omega(a_0^{1-\omega}/L + L^{\frac{1}{1-\omega}}((k+1)/2)^{\frac{1}{1-\omega}})$ .  $\square$

THEOREM 2. Let us specify that with probability at least  $1 - \delta$ , for an agent  $j$ ,  $\|Q_T^j - Q_*^j\|_\infty \leq \epsilon$ . The bound on the rate of convergence of low- $Q$ ,  $Q_T^j$ , with a polynomial learning rate of factor  $\omega$  is given by (with  $Q_*^j$  as the Nash  $Q$ -value of  $j$ )

$$T = \Omega\left(\left(\frac{L^{1+3\omega} Q_{\max}^{2,j} \ln\left(\frac{|S| |\Pi_i| |A|_i Q_{\max}^j}{\delta \beta \epsilon}\right)}{\beta^2 \epsilon^2}\right)^{1-\omega} / L + \left(\left(\frac{L}{\beta} \ln \frac{Q_{\max}^j}{\epsilon} + 1\right)/2\right)^{\frac{1}{1-\omega}}\right).$$

PROOF. The proof of Theorem 2 follows from Lemmas 7, 11, 4, and 12.

Specifically, from the relation in Lemma 12 substitute the value of  $a_0$  from the Lemma 11 (value of  $\tau_0$ ), and value of  $k$  from Lemma 4 (lower bound for  $m$ ). From the Lemma 7 and Lemma 11, we see that the condition required in Lemma 5 is satisfied to provide a lower bound.  $\square$

## D PROOF OF THEOREM 3

In this section, we aim to show that the size of the  $k$ th iteration is  $L(1+\psi)\tau_k$  for some positive constant  $\psi \leq 0.712$ . The covering time property guarantees that in  $(1+\psi)L\tau_k$  steps, each pair of state-joint actions are performed at least  $(1+\psi)\tau_k$  times. The sequence of times in this case is  $\tau_{k+1} = \tau_k + (1+\psi)L\tau_k$ . As in the last section, we will first bound  $Y_{t;\tau_k}^j$  and then bound  $W_{t;\tau_k}^j$ . As in Appendix B, we will use  $Q$  to denote the low- $Q$  values across this section as well. The proof of this theorem follows the Theorem 5 in Dar and Mansour [5]. While the work of Dar and Mansour was restricted to single-agent MDPs, our result extends the analysis of Dar and Mansour to the general-sum stochastic game setting.

LEMMA 13. Consider the low- $Q$  update given in Eq. 4, with a linear learning rate. Assume that for  $t \geq \tau_k$  we have that  $Y_{t;\tau_k}^j(s, \mathbf{a}) \leq D_k^j$ . Then for any  $t \geq \tau_k + (1+\psi)L\tau_k = \tau_{k+1}$  we have that  $Y_{t;\tau_k}^j(s, \mathbf{a}) \leq (\gamma + \frac{2}{2+\psi}\beta)D_k^j$ .

PROOF. For each state-joint action pair, we are assured that  $n(s, \mathbf{a}, \tau_k, \tau_{k+1}) \geq (1+\psi)\tau_k$ , since in an interval of  $(1+\psi)L\tau_k$  steps, each state-joint action pair is visited at least  $(1+\psi)\tau_k$  times by the definition of covering time.

Let  $Y_{\tau_k, \tau_k}^j(s, \mathbf{a}) = \gamma D_k^j + \rho_{\tau_k}^j$ , where  $\rho_{\tau_k}^j = (1-\gamma)D_k^j$ . We now have,

$$\begin{aligned} \rho_t^j &= (1-\gamma)\prod_{l=1}^{t-\tau_k}(1-\alpha_{l+\tau_k}) \\ &= 2\beta D_k^j \prod_{l=1}^{t-\tau_k}(1-\alpha_{l+\tau_k}) \\ &= 2\beta D_k^j \prod_{l=1}^{t-\tau_k}\left(1 - \frac{1}{|n(s, \mathbf{a}, \tau_{k+1}, l)|}\right), \end{aligned} \quad (48)$$

where the last identity follows from the fact that  $\alpha_t = \frac{1}{n(s, \mathbf{a}, 0, t)}$ . Since the  $\tau_k$ 's are monotonically decreasing, using  $t = \tau_k + (1 + \psi)L\tau_k$ , we get (using  $\psi < 0.712$ ),

$$\begin{aligned} \rho_t &\leq 2D_k^j \beta (1 - \frac{1}{(1+\psi)\tau_k})^{t-\tau_k} \\ &\leq 2D_k^j \beta (1 - \frac{1}{(1+\psi)\tau_k})^{1+\psi L\tau_k} \\ &\leq 2D_k^j \beta (1 - \frac{1}{(1+\psi)\tau_k})^{1+\psi\tau_k} \\ &\leq \frac{2D_k^j \beta}{e} \\ &\leq \frac{2D_k^j \beta}{2+\psi} \end{aligned} \quad (49)$$

Hence,  $Y_{t;\tau_k}^j(s, \mathbf{a}) \leq (\gamma + \frac{2}{2+\psi}\beta)D_k^j$ .  $\square$

The following lemma enables the use of Azuma's inequality.

LEMMA 14. For any  $t \geq \tau_k$  and  $1 \leq l \leq t$  we have that  $W_{t;\tau_k}^{l,j}(s, \mathbf{a})$  is a martingale sequence, which satisfies,

$$|W_{t;\tau_k}^{l,j}(s, \mathbf{a}) - W_{t;\tau_k}^{l-1,j}(s, \mathbf{a})| \leq \frac{Q_{\max}^j}{n(s, \mathbf{a}, 0, t)} \quad (50)$$

PROOF.  $W_{t;\tau_k}^{l,j}(s, \mathbf{a})$  is a martingale difference sequence since,

$$\begin{aligned} &\mathbb{E}[W_{t;\tau_k}^{l,j}(s, \mathbf{a}) - W_{t;\tau_k}^{l-1,j}(s, \mathbf{a}) | F_{\tau_k+l-1}] \\ &= \mathbb{E}[\eta_{\tau_k+l}^{k,t,j}(s, \mathbf{a}) w_{\tau_k+l}^j(s, \mathbf{a}) | F_{\tau_k+l-1}] \\ &= \eta_{\tau_k+l}^{k,t,j}(s, \mathbf{a}) \mathbb{E}[w_{\tau_k+l}^j(s, \mathbf{a}) | F_{\tau_k+l-1}] = 0 \end{aligned} \quad (51)$$

For linear learning rate we have that  $\eta_{\tau_k+l}^{k,t,j}(s, \mathbf{a}) \leq \alpha_{l+\tau_k} = 1/n(s, \mathbf{a}, 0, t)$ , thus

$$\begin{aligned} &|W_{t;\tau_k}^{l,j}(s, \mathbf{a}) - W_{t;\tau_k}^{l-1,j}(s, \mathbf{a})| \\ &= \eta_{\tau_k+l}^{k,t,j}(s, \mathbf{a}) |w_{\tau_k+l}^j(s, \mathbf{a})| \\ &\leq \frac{Q_{\max}^j}{n(s, \mathbf{a}, 0, t)}. \end{aligned} \quad (52)$$

$\square$

The following lemma provides a bound for the stochastic term  $W_{t;\tau_k}^j$ .

LEMMA 15. Consider the low- $Q$  update given in Eq. 4, with a linear learning rate. With probability at least  $1 - \frac{\delta}{m}$  we have that for every state-joint action pair  $|W_{t;\tau_k}^j(s, \mathbf{a})| \leq \frac{\psi}{2+\psi} \beta D_k^j$ , for any  $t > \tau_{k+1}$  and any positive constant  $\psi \leq 0.712$ , i.e.,

$$Pr\left[\forall t \in [\tau_{k+1}, \tau_{k+2}] : |W_{t;\tau_k}^j(s, \mathbf{a})| \leq \frac{\psi}{2+\psi} \beta D_k^j\right] \geq 1 - \frac{\delta}{m} \quad (53)$$

given that  $\tau_k \geq \Theta(\frac{Q_{\max}^{2,j} \ln(Q_{\max}^j |S| |\Pi_i| |A|_i m)}{\psi^2 \beta^2 D_k^2})$ .

PROOF. By Lemma 14 we can apply Azuma's inequality on the term,  $W_{t;\tau_k}^{t-\tau_k+1,j}$  (note that  $W_{t;\tau_k}^{t-\tau_k+1,j} = W_{t;\tau_k}$ ), and with the expression  $c_i = \Theta(\frac{Q_{\max}^j}{n(s, \mathbf{a}, 0, t)})$  for any  $t \geq \tau_{k+1}$ . Therefore, we derive that,

$$Pr[|W_{t;\tau_k}^j| \geq \tilde{\epsilon}] \leq 2e^{-\frac{2e\tilde{\epsilon}^2}{\sum_{i=\tau_k}^t c_i^2}} \leq 2e^{-c \frac{\tilde{\epsilon}^2 n(s, \mathbf{a}, \tau_k, t)}{Q_{\max}^{2,j}}} \quad (54)$$

for some positive constant  $c$ .

Let us define,

$$\begin{aligned} \zeta_t(s, \mathbf{a}) &= 1, \text{ if } \alpha_t(s, \mathbf{a}) \neq 0 \\ \zeta_t(s, \mathbf{a}) &= 0, \text{ otherwise.} \end{aligned} \quad (55)$$

Using the union bound and the property that, in an interval of length  $(1 + \psi)L\tau_k$ , each state-joint action pair is visited at least  $(1 + \psi)\tau_k$  times, we get

$$\begin{aligned} &Pr\left[\forall t \in [\tau_{k+1}, \tau_{k+2}] : |W_{t;\tau_k}^j(s, \mathbf{a})| \geq \tilde{\epsilon}\right] \\ &\leq Pr[\forall t \geq ((1 + \psi)L + 1)\tau_k : |W_{t;\tau_k}^j(s, \mathbf{a})| \geq \tilde{\epsilon}] \\ &\leq \sum_{t=((1+\psi)L+1)\tau_k}^{\infty} Pr\left[|W_{t;\tau_k}^j(s, \mathbf{a})| \geq \tilde{\epsilon}\right] \\ &\leq \sum_{t=((1+\psi)L+1)\tau_k}^{\infty} \zeta_t(s, \mathbf{a}) 2e^{-c \frac{\tilde{\epsilon}^2 n(s, \mathbf{a}, 0, t)}{Q_{\max}^{2,j}}} \\ &\leq 2e^{-c \frac{\tilde{\epsilon}^2 ((1+\psi)\tau_k)}{Q_{\max}^{2,j}}} \sum_{t=0}^{\infty} e^{-\frac{t\tilde{\epsilon}^2}{Q_{\max}^{2,j}}} \\ &= \frac{2e^{-c \frac{\tilde{\epsilon}^2 ((1+\psi)\tau_k)}{Q_{\max}^{2,j}}}}{1 - e^{-\frac{\tilde{\epsilon}^2}{Q_{\max}^{2,j}}}} \\ &= \Theta\left(\frac{Q_{\max}^{2,j} e^{-\frac{c'\tau_k \tilde{\epsilon}^2}{Q_{\max}^{2,j}}}}{\tilde{\epsilon}^2}\right) \end{aligned} \quad (56)$$

for some positive constant  $c'$ . Setting  $\frac{\delta}{m|S|\Pi_i|A|_i} = \Theta\left(\frac{e^{-\frac{c'\tau_k \tilde{\epsilon}^2}{Q_{\max}^{2,j}}}}{\frac{Q_{\max}^{2,j}}{\tilde{\epsilon}^2}}\right)$ ,

which holds for  $\tau_k = \Theta(\frac{Q_{\max}^{2,j} \ln(Q_{\max}^j |S| |\Pi_i| |A|_i m)}{\tilde{\epsilon}^2})$ , and the expression,  $\tilde{\epsilon} = \frac{\psi}{2+\psi} \beta D_k^j$  assures us that for every  $t \geq \tau_{k+1}$  (and as a result for any  $t \in [\tau_{k+1}, \tau_{k+2}]$ ), with probability at least  $1 - \frac{\delta}{m}$  the statement holds at every state-joint action pair.  $\square$

We have bounded for each iteration the time needed to achieve the desired probability of  $1 - \frac{\delta}{m}$ . The following lemma provides a bound for the error in all the iterations.

LEMMA 16. Consider the low- $Q$  update given in Eq. 4, with a linear learning rate. With probability  $1 - \delta$ , for every iteration  $k \in [1, m]$ , time  $t \in [\tau_{k+1}, \tau_{k+2}]$ , and any positive constant  $\psi \leq 0.712$ , we have  $|W_{t;\tau_k}^j| \leq \frac{\psi \beta D_k^j}{2+\psi}$ , i.e.,

$$Pr\left[\forall k \in [1, m], \forall t \in [\tau_{k+1}, \tau_{k+2}] : |W_{t;\tau_k}^j| \leq \frac{\psi \beta D_k^j}{2+\psi}\right] \geq 1 - \delta \quad (57)$$



given that  $\tau_0 = \Theta\left(\frac{Q_{\max}^{2,j} \ln(Q_{\max}^j |S| |A| m / (\delta \beta \epsilon \psi))}{\psi^2 \beta^2 \epsilon^2}\right)$

PROOF. From Lemma 15, we know that

$$\Pr\left[\forall t \in [\tau_{k+1}, \tau_{k+2}] : |W_{t;\tau_k}^j| \geq \frac{\psi \beta D_k^j}{2+\psi}\right] \leq \frac{\delta}{m}. \quad (58)$$

Using the union bound, we have that,

$$\begin{aligned} & \Pr\left[\forall k \leq m, \forall t \in [\tau_{k+1}, \tau_{k+2}] : |W_{t;\tau_k}^j| \geq \frac{\psi \beta D_k^j}{2+\psi}\right] \\ & \leq \sum_{k=1}^m \Pr\left[\forall t \in [\tau_{k+1}, \tau_{k+2}] : |W_{t;\tau_k}^j| \geq \frac{\psi \beta D_k^j}{2+\psi}\right] \\ & \leq \delta. \end{aligned} \quad (59)$$

□

**THEOREM 3.** *Let us specify that with probability at least  $1 - \delta$ , we have for an agent  $j$ ,  $\|Q_T^j - Q_*^j\|_\infty \leq \epsilon$ . The bound on the rate of convergence of low- $Q$ ,  $Q_T^j$ , with a linear learning rate is given by*

$$T = \Omega\left((L + \psi L + 1)^{\frac{1}{\beta}} \ln \frac{Q_{\max}^j}{\epsilon} \frac{Q_{\max}^{2,j} \ln\left(\frac{|S| |\Pi_L| |A| Q_{\max}^j}{\delta \beta \epsilon \psi}\right)}{\beta^2 \epsilon^2 \psi^2}\right),$$

where  $\psi$  is a small arbitrary positive constant satisfying  $\psi \leq 0.712$

PROOF. The Theorem 3 follows from Lemmas 16, 13, 4, and the fact that  $a_{k+1} = a_k + (1 + \psi)La_k = a_0((1 + \psi)L + 1)^k$ .


Specifically, substitute the value of  $k$  from Lemma 4 (lower bound for  $m$ ), value of  $a_0$  from Lemma 16 (value of  $\tau_0$ ), and see that Lemma 13 and Lemma 16 satisfy the condition for the lower bound in Lemma 5.

□


## E FREQUENCY OF LISTENING TO ADVISORS

This section plots the frequency of listening to each advisor in some of our experiments. We would like to show that the MA-TLQL listens more to the good advisor and avoids the bad advisor more than other related baselines. For these experiments, we consider the TLQL [21] and ADMIRAL-DM [46] algorithms for comparison. Since the CHAT implementation uses the same method to choose advisors as ADMIRAL-DM (weighted random policy approach), we will omit CHAT for these results (performance is similar to ADMIRAL-DM). DQfD uses pretraining and does not choose advisors in an online fashion; hence we omit DQfD for these experiments as well.

We consider our first experiment in Section 6 (Experiment 1), where we had a set of four sufficient advisors of different quality. The first advisor (Advisor 1) had a better quality than the others, and the agents must listen more to this advisor. On the other hand, Advisor 4 only suggested random actions and the agents are expected to avoid listening to this advisor. In Figure 11(a), we plot a curve that corresponds to the percentage of time steps an algorithm listened to Advisor 1 out of all the time steps the algorithm had an opportunity to listen to one of the available advisors. From the plots, we see that MA-TLQL listens more (compared to the other baselines) to this advisor (Advisor 1) from the beginning until the end of training. Since the MA-TLQL uses an ensemble technique to choose an advisor, this gives it a distinct advantage in the early



(a) Good Advisor (Advisor 1)




(b) Bad Advisor (Advisor 4)


**Figure 11: Frequency of listening to advisors in the two-agent Pommerman experiment with four sufficient advisors of different quality (Experiment 1)**

stages of training. Further, since MA-TLQL performs an explicit evaluation of the advisors independent of the RL policy, it manages to listen more to the correct advisor as compared to other baselines. MA-TLAC also listens more to the good advisor as compared to the other baselines (TLQL, ADMIRAL-DM). Since TLQL couples the advisor evaluation with the RL policy, it listens a lot less to the good advisor as compared to MA-TLQL. Also, TLQL considers the RL policy as part of the high-level table, which makes it less reliant on advisors. This could be a problem when good advisors are available. In Figure 11(b), we plot the percentage of each algorithm listening to the bad advisor (Advisor 4). We see that MA-TLQL has the least dependence on this advisor as compared to all other algorithms. This reinforces our observation that MA-TLQL is most likely to choose to listen to the correct advisors in this multi-agent setting.

We plot the percentage of listening to Advisor 1 and Advisor 4 in Experiment 3 from Section 6 that had an insufficient set of four advisors of decreasing quality. The results in Figure 12(a) and (b)



(a) Good Advisor (Advisor 1)



(b) Bad Advisor (Advisor 4)


**Figure 12: Frequency of listening to advisors in the two-agent Pommerman experiment with four insufficient advisors of different quality (Experiment 3)**

show that MA-TLQL listens more to the good advisor and less to the bad advisor, same as our observations for Experiment 1.


Similarly, we plot the percentages of listening to the good and bad advisor in the Pommerman team environment (mixed setting) used in Experiment 5 in Figure 13(a) and (b). Here we plot the results for one of the two pommerman agents playing in the same team (the other agent’s results are similar). Once again, we note that MA-TLQL listens more to the correct advisor than the other algorithms and better avoids the bad advisors compared to the other algorithms.

## F NATURE OF ALGORITHMS CONSIDERED

In Table 1, we tabulate all the algorithms considered in this paper. The differences between the algorithms stem from the nature of the algorithm (independent or multi-agent), ability to naturally support learning from conflict demonstrations (i.e., more than one advisor), and type of demonstrations that they naturally support (offline



(a) Good Advisor (Advisor 1)



(b) Bad Advisor (Advisor 4)

**Figure 13: Frequency of listening to advisors in the team Pommerman experiment with four sufficient advisors of different quality (Experiment 5)**

vs. online). Offline demonstrations are demonstrations that are collected (in a memory buffer) from an advisor before the “training phase” where the algorithm is trained using interactions with the environment. These demonstrations are typically used to train the algorithm in a “pre-training” phase before regular training, as done in several prior works [7, 12, 17]. Alternatively, online demonstrations are obtained in real-time during the training phase (and not pre-collected). Here, an agent can actively obtain action recommendations directly from an advisor for the current game context.


In general, in all the experiments considered in this paper, we found that algorithms that consider actions of other agents to provide best-responses (non-independent) performed better than independent algorithms which consider all other agents to be part of the environment. One reason for this behaviour could be the fact that independent algorithms break the Markovian assumptions in reinforcement learning methods [40]. Additionally, we found that

Algorithm	Nature of Algorithm	Number of Advisors	Type of Demonstrations
DQN [26]	Independent	Does not learn from advising	Not applicable
DQfD [12]	Independent	One	Offline
CHAT [46]	Independent	One	Online
ADMIRAL-DM [37]	Multi-agent	One	Online
TLQL [21]	Independent	More than one	Online
MA-TLQL (ours)	Multi-agent	More than one	Online
MA-TLAC (ours)	Multi-agent	More than one	Online


**Table 1: Description of all algorithms considered in this paper**

algorithms that learn from online advising perform better as compared to algorithms that learn from offline advising. If algorithms learn from online advising, then these algorithms can exploit the knowledge of advisors in response to dynamically changing other agent(s), in real-time. Since other agent(s)/opponent(s) are not typically available before training in multi-agent environments, offline demonstrations are not very successful in multi-agent training, as opposed to single-agent training where they were quite successful [12]. In this paper, we considered environments where the advising come from multiple sources of independent knowledge. In many states, the different advisors provide conflicting recommendations. Hence, algorithms that can effectively resolve conflicting information from the different advisors were successful as compared to other algorithms that are not capable of naturally supporting learning under multiple conflicting advisors. As seen from Table 1, the only two algorithms that have all the three desirable properties (i.e., support multi-agent update, learn from conflicting advisors, support offline demonstrations) are MA-TLQL and MA-TLAC which gave the best performance in most of the seven experiments considered in Section 6.

In our paper, we do not consider HMAT [18] and LeCTR [27] as baselines, though these can also be classified as action advising methods. We do not consider these algorithms as appropriate benchmarks due to two important reasons. First, they are both restricted to two-agent cooperative settings while we are interested in more general domains including those with more than two agents and both competitive and mixed-motive environments. Second, we are interested in independent learning from a set of external advisors, while both HMAT and LeCTR focus on peer-to-peer learning (partly



(a) Pursuit Environment



(b) Predator-Prey Environment

**Figure 14: Performances of MA-TLQL and MA-TLAC under changing (learning) advisors and fixed advisors. Here (a) corresponds to the setting in Experiment 6 and (b) corresponds to the setting in Experiment 7, described in Section 6.**

due to their focus on teaching teams, as opposed to our work which is on more general multi-agent learning problems).

## G MA-TLQL WITH LEARNING ADVISORS

In all the experiments in this paper we considered advisors that are fixed and non-changing, which we specified in Section 4. If the advisors are fixed, their  $Q$ -values can be determined using the high- $Q$  updates in MA-TLQL. The fixed nature of advisors allows us to provide theoretical guarantees of convergence of the high- $Q$  values (under the assumption of infinite updates in the limit) using similar arguments as in Theorem 1. However, experimentally we can still consider other types of advisors which are actively learning during the training stage and hence are updating the policies actively (though using such advisors do not have any theoretical guarantees of convergence). In this section, we will revisit Experiment 6 and Experiment 7 from Section 6 and study the performances of MA-TLQL and MA-TLAC under learning advisors. For statistical

significance we use the unpaired 2-sided t-test and report  $p$ -values, as in the earlier experiments.


First we consider Experiment 6, where we used the Pursuit cooperative environment along with four pre-trained networks of DQN as the advisor. Now, we will consider the same set of four advisors and label them as “fixed advisors”. Additionally, we will let the same four pre-trained networks of DQN continue training while it is being actively used for action advising. We label this set of four advisors as “changing advisors”. In Figure 14(a) we plot the performances (plot of episodic rewards) of MA-TLQL and MA-TLAC along with the fixed as well as the changing advisor set. While using the changing advisor set, we see that MA-TLQL and MA-TLAC outperform their counterparts using the fixed advisors ( $p < 0.04$ ). As the advisors are actively learning and changing their strategies during the training stage, they are able to provide better action recommendations to MA-TLQL and MA-TLAC at the different parts of the environment. This is reflected in their superior performances.

Similarly, we revisit Experiment 7 with the Predator-Prey environment and consider two different sets of advisors. In Experiment 7, we used four pre-trained DQN networks as advisors. We will reuse the same set of four advisors and label them as the “fixed advisors”. A continuously training counterpart is labelled as the “changing advisors”. We plot the performances of MA-TLAC in the Predator-Prey experiment in Figure 14(b), and once again we notice that MA-TLAC using the changing advisor set outperforms MA-TLAC using the fixed advisor set ( $p < 0.03$ ).


Hence, experimentally we see that MA-TLQL and MA-TLAC can still be used with changing advisors which provide good empirical performances. However, the non-stationary nature of these advisors makes it impossible to provide theoretical guarantees of convergence. This is similar to using independent algorithms in multi-agent environments, which do not have any theoretical guarantees of convergence to either a local or global optimum, yet empirically, several prior works have noted that these algorithms perform well in various multi-agent environments [11, 25].

## H MA-TLQL WITH OPPONENT MODELLING

In all the experiments in this paper, the MA-TLQL used a multi-agent update where it simply tracked opponent actions (by considering the previous action) and did not perform any active opponent modelling. The core focus of this paper is learning from advisors, and we chose not to perform any particular opponent modelling technique to keep the algorithm simple. However, in this section, we will implement MA-TLQL along with an opponent modelling technique that uses a separate neural network (2 Relu layers of 50 neurons and an output layer) to predict the action of the opponent. The network uses the state and previous action of the opponent as the input and predicts the action of the opponent. This predicted action of the opponent is used by the agent to calculate best responses. Finally, the actual observed action of the opponent is used to define a cross-entropy loss function that is used to train the network. We revisit Experiment 3 and Experiment 4 in Section 6 where we used four insufficient advisors of different and similar quality. Now, we will use the same experimental procedure (with the same set of advisors) and consider the performances of MA-TLQL with and



(a) One vs. One Pommerman: Insufficient different quality advisors



(b) One vs. One Pommerman: Insufficient similar quality advisors

**Figure 15: Performances of MA-TLQL with and without opponent modelling using four insufficient advisors. Here (a) corresponds to the setting in Experiment 3 and (b) corresponds to the setting in Experiment 4, described in Section 6.**

without active opponent modelling. For statistical significance we use the unpaired 2-sided t-test and report  $p$ -values.

We plot the performances in Figure 15. From both Figure 15(a) and (b) we see that MA-TLQL with opponent modelling performs better than the MA-TLQL algorithm that does not perform any active opponent modelling. In the Figure 15(a) we see that opponent modelling only provides a small increase in performance ( $p < 0.2$ ). One reason for this is the fact that MA-TLQL without opponent modelling is already showing a high performance in this experiment by learning efficiently from the advisors. Alternatively, Figure 15(b) shows a considerable increase in performance while MA-TLQL is using opponent modelling ( $p < 0.04$ ). Here, the performance of MA-TLQL without opponent modelling is not as good as the performance in Figure 15(a) and opponent modelling shows a marked improvement in performance.






Figure 16: MA-TLQL with different numbers of advisors in the Pursuit environment (setting used in Experiment 6, described in Section 6)

## I MA-TLQL WITH DIFFERENT NUMBERS OF ADVISORS

In this section we train MA-TLQL with different number of advisors in the Pursuit environment considered in Experiment 6 of Section 6. Recall that in Experiment 6 we considered four pre-trained networks of DQN as the advisor (pre-trained for 500, 1000, 1500, and 2000 episodes). Labelling these advisors, we will denote Advisor 1 as the advisor pre-trained for 500 episodes, Advisor 2 as the advisor pre-trained for 1000 episodes, Advisor 3 as the advisor pre-trained for 1500 episodes, and Advisor 4 as the advisor pre-trained for 2000 episodes. In this experiment, we initially train MA-TLQL with no advisor (denoted as MA-QL), and subsequently train MA-TLQL with the addition of one advisor from the set of advisors. MA-QL always chooses actions from the low- $Q$  (since there are no advisors, the high- $Q$  is not maintained). The objective is to study the performance of MA-TLQL under the presence of different numbers of advisors. For statistical significance we use the unpaired 2-sided  $t$ -test and report  $p$ -values.

The results are given in Figure 16. We see that MA-QL trained without any advisor gives the least performance. MA-TLQL trained using one advisor (Advisor 1) gives a better performance than MA-QL ( $p < 0.06$ ). Next, we train MA-TLQL with two advisors (Advisor 1 and Advisor 2), which gives a better performance than MA-QL and MA-TLQL trained with one advisor ( $p < 0.1$ ). Similarly, MA-TLQL with three advisors (Advisor 1, Advisor 2, and Advisor 3) gives a better performance than the case with two advisors ( $p < 0.09$ ) and the best performance is given by MA-TLQL with all four advisors ( $p < 0.04$ ). This result shows that MA-TLQL performance can keep improving with the addition of better advisors (than those available in the current set), which shows that MA-TLQL is capable of identifying the right advisor from the available set and exploiting the expertise of different available advisors in a multi-agent environment. From the  $p$ -values, we note that the observation of the best performance of MA-TLQL with all four advisors is statistically significant. While we observe constantly improving

S1 Reward = 0	S2 Reward = 0	G Reward = 1
S3 Reward = -1	S4 Reward = -1	S5 Reward = 0

Figure 17: Toy environment to compare updates of TLQL and MA-TLQL

performances with each advisor, some of these comparisons are not statistically significant as provided by the  $p$ -values.

## J ILLUSTRATIVE EXAMPLE

In this section we would like to show a toy example where the TLQL updates as provided by Li et al. [21] takes a longer time to figure out the best advisor from a set of advisors, as compared to the MA-TLQL updates, we introduced in this paper. We will just use a single agent based grid-world environment as given in Figure 17, instead of multi-agent environments. The TLQL updates will use the Bellman update for the low- $Q$  updates and a subsequent synchronization step to update its high- $Q$  values as proposed in Li et al. [21]. In MA-TLQL, the low- $Q$  values will use the control update as given in Eq. 3 albeit in a single-agent fashion (joint actions need not be considered), since this environment only contains one agent. The high- $Q$  will use the evaluation update as given in Eq. 1 in a single-agent fashion. Also, for simplicity, we do not use the ensemble selection strategy in Eq. 2 for the MA-TLQL action selection in this example. Rather we only use the high- $Q$  and low- $Q$  updates for advisor and action evaluation at the given state. Here we have a set of 6 states  $\{S1, S2, S3, S4, S5, G\}$  with the agent starting at state  $S1$  and trying to reach the goal state  $G$ . The states  $S3, S4$  and  $G$  are the terminal states where the agent receives a reward of -1 in states  $S3$  and  $S4$ , and a reward of +1 in state  $G$ . The state  $S5$  only exists for symmetry (cannot be reached in practice). The agent can take one of the two actions  $\{R, D\}$  (to denote right and down respectively) at each state. In this environment, it can be seen that the agent needs to take action  $R$  in states  $S1$  and  $S2$  to obtain the maximum rewards. The agent has access to two advisors  $A1$  and  $A2$ , where  $A1$  is an optimal advisor providing the correct action (right) at every state and  $A2$  is a sub-optimal advisor which provides action  $R$  with probability 0.5 and the action  $D$  with probability 0.5. All transitions in this domain are deterministic. Also, we specify that the learning rate ( $\alpha$ ) is 0.1 and the discount factor ( $\gamma$ ) is 0.9.

Now we consider the  $Q$ -updates (high-level and low-level) pertaining to both MA-TLQL and TLQL. At the beginning, we initialize all the  $Q$ -values to 0 arbitrarily.

At the initial time step ( $t = 0$ ) let us assume that the agent starts at the initial state  $S1$ . Let both the advisors suggest action  $R$ . Then the agent takes this action and first updates its low- $Q$  using the equation,



(ST./AC.)	Right	Down
S1	0	0
S2	0	0

(a) Low- $Q$  values at time  $t = 1$ 

(ST./AC.)	Adv. 1	Adv. 2
S1	0	0
S2	0	0

(b) TLQL high- $Q$  at time  $t = 1$ 

(ST./AC.)	Adv. 1	Adv. 2
S1	0	0
S2	0	0

(c) MA-TLQL high- $Q$  at time  $t = 1$ **Table 2: TLQL and MA-TLQL updates at time  $t = 1$ . The columns refer to the actions and the rows refer to the states.**

(ST./AC.)	Right	Down
S1	0	0
S2	0.1	0

(a) Low- $Q$  values at time  $t = 2$ 

(ST./AC.)	Adv. 1	Adv. 2
S1	0	0
S2	0.1	0.1

(b) TLQL high- $Q$  at time  $t = 2$ 

(ST./AC.)	Adv. 1	Adv. 2
S1	0	0
S2	0.1	0.1

(c) MA-TLQL high- $Q$  at time  $t = 2$ **Table 3: TLQL and MA-TLQL updates at time  $t = 2$** 

(ST./AC.)	Right	Down
S1	0.009	0
S2	0.1	0

(a) Low- $Q$  values at time  $t = 3$ 

(ST./AC.)	Adv. 1	Adv. 2
S1	0.009	0.009
S2	0.1	0.1

(b) TLQL high- $Q$  at time  $t = 3$ 

(ST./AC.)	Adv. 1	Adv. 2
S1	0.009	0.009
S2	0.1	0.1

(c) MA-TLQL high- $Q$  at time  $t = 3$ **Table 4: TLQL and MA-TLQL updates at time  $t = 3$** 

$$\begin{aligned} lowQ_1(S1, R) \\ = lowQ_0(S1, R) + \alpha(r + \gamma \max_a lowQ_0(S2, a) - lowQ_0(S1, R)) \end{aligned}$$

$$lowQ_1(S1, R) = 0 + 0.1(0 + 0.9 \times 0 - 0)$$

$$lowQ_1(S1, R) = 0. \quad (60)$$

Now, TLQL will set the value of the high- $Q$  of both advisors to be 0 (synchronization). The MA-TLQL updates will also yield a value of 0 for both the advisors. In Table 2 we tabulate the  $Q$ -values for the non-terminal states (S1 and S2).

At the next time step  $t = 2$ , the agent is at state S2. Again let both advisors suggest the right action. Now the low- $Q$  is updated as,

$$\begin{aligned} lowQ_2(S2, R) \\ = lowQ_1(S2, R) + \alpha(r + \gamma \max_a lowQ_1(G, a) - lowQ_1(S2, R)) \\ lowQ_2(S2, R) = 0 + 0.1(1 + 0.9 \times 0 - 0) \\ lowQ_2(S2, R) = 0.1 \end{aligned} \quad (61)$$

We specify that when the next state is terminal, the temporal difference (T.D.) target is the reward itself. Both the algorithms, TLQL and MA-TLQL, will have the same high- $Q$  values in this case as well. The  $Q$ -values for time  $t = 2$  are tabulated in Table 3.

Now, we move to time  $t = 3$ . Since the goal has been reached at the previous time step, the agent resets back to the initial state S1. Now, let us assume that both advisors specify the right action again at state S1. The low- $Q$  values are updated as,

$$\begin{aligned} lowQ_3(S1, R) \\ = lowQ_2(S1, R) + \alpha(r + \gamma \max_a lowQ_2(S2, a) - lowQ_2(S1, R)) \\ lowQ_3(S1, R) = 0 + 0.1(0 + 0.9 \times 0.1 - 0) \\ lowQ_3(S1, R) = 0.009 \end{aligned} \quad (62)$$

Again, high- $Q$  values for MA-TLQL and TLQL will be the same as the low- $Q$  values and are tabulated in Table 4.

Next, the agent moves to state S2. We are at time  $t = 4$ . Let the advisor A2 specify action  $D$  at this state (Advisor A1 always specifies  $R$ ). Also let us assume that the agent chooses to listen to A2 at this state ( $Q$ -values of both advisors are the same, so the agent is indifferent between the two advisors) and hence it performs action  $D$  (down) from A2. Now the  $Q$ -value for the low- $Q$  can be updated as,

$$\begin{aligned} lowQ_4(S2, D) &= lowQ_3(S2, D) \\ &\quad + \alpha(r + \gamma \max_a lowQ_4(S4, a) - lowQ_4(S2, D)) \\ lowQ_4(S2, D) &= -0.1 \end{aligned} \quad (63)$$

The high- $Q$  values for the MA-TLQL update is given by,

(ST./AC.)	Right	Down
S1	0.009	0
S2	0.1	-0.1

(a) Low- $Q$  values at time  $t = 4$ 

(ST./AC.)	Adv. 1	Adv. 2
S1	0.009	0.009
S2	0.1	-0.1

(b) TLQL high- $Q$  at time  $t = 4$ 

(ST./AC.)	Adv. 1	Adv. 2
S1	0.009	0.009
S2	0.1	-0.01

(c) MA-TLQL high- $Q$  at time  $t = 4$ **Table 5: TLQL and MA-TLQL updates at time  $t = 4$** 

$$\begin{aligned}
highQ_4(S2, A2) &= highQ_3(S2, A2) \\
&+ \alpha(r + \gamma highQ_3(S4, A2) - highQ_3(S2, A2)) \\
highQ_4(S2, A2) &= 0.1 + 0.1(-1 - 0.1)
\end{aligned} \tag{64}$$

$$highQ_4(S2, A2) = -0.01$$

All  $Q$ -values for time  $t = 4$  are tabulated in Table 5.

Since the state  $S4$  was a terminal state, the agent is back to state  $S1$ . We are at time  $t = 5$ . At this state, let us assume that both the advisors specify action  $R$ . Now the agent chooses to perform this action and updates its low- $Q$  using,

$$\begin{aligned}
lowQ_5(S1, R) &= lowQ_4(S1, R) \\
&+ \alpha(r + \gamma \max_a lowQ_4(S2, a) - lowQ_4(S1, R)) \\
lowQ_5(S1, R) &= 0.009 + 0.1(0 + 0.9 \times 0.1 - 0.009)
\end{aligned} \tag{65}$$

$$lowQ_5(S1, R) = 0.0171.$$

Since both the advisors specified action  $R$ , both advisors are assigned the same  $Q$ -values in the high- $Q$  in the TLQL update. Now, the high- $Q$  estimate of advisor  $A1$  will be the same as the low- $Q$  value in the MA-TLQL update as well. However, the high- $Q$  estimate of the advisor  $A2$  of the MA-TLQL update will be as follows,

$$\begin{aligned}
highQ_5(S1, A2) &= highQ_4(S1, A2) \\
&+ \alpha(r + \gamma highQ_4(S2, A2) - highQ_4(S1, A2)) \\
highQ_5(S1, A2) &= 0.009 + 0.1(0 + 0.9 \times -0.01 - 0.009) \\
highQ_5(S1, A2) &= 0.009 - 0.0018 = 0.0072
\end{aligned} \tag{66}$$

At this stage (time  $t = 5$ ) all the  $Q$ -values are tabulated in Table 6. Comparing the high- $Q$  estimate of TLQL and MA-TLQL, we see that in the TLQL updates (Table 6b), the agent is indifferent between following advisor  $A1$  or advisor  $A2$  in state  $S1$  (same  $Q$ -values), while it would decide to choose advisor  $A1$  at the state  $S2$ . Even after 5 update steps TLQL has not been able to determine the right advisor (as advisor  $A1$  is better than  $A2$ ) for both states. In contrast,

(ST./AC.)	Right	Down
S1	0.0171	0
S2	0.1	-0.1

(a) Low- $Q$  values at time  $t = 5$ 

(ST./AC.)	Adv. 1	Adv. 2
S1	0.0171	0.0171
S2	0.1	-0.1

(b) TLQL high- $Q$  at time  $t = 5$ 

(ST./AC.)	Adv. 1	Adv. 2
S1	0.0171	0.0072
S2	0.1	-0.01

(c) MA-TLQL high- $Q$  at time  $t = 5$ **Table 6: TLQL and MA-TLQL updates at time  $t = 5$** 

the MA-TLQL updates found in Table 6c clearly show a higher  $Q$ -value for the advisor  $A1$  than the advisor  $A2$  for both the states. This example presents a situation where the MA-TLQL distinguishes between a good and a bad advisor faster than the vanilla TLQL update as introduced by Li et al. [21].

## K EXPERIMENTAL DETAILS

This section provides the complete details of all of our experimental domains, including details about the reward function and the advisors used. For the Pommerman and Pursuit environments, we assume that all the actions of other agents are either directly observable (fully observable), shared amongst agents, or provided by the game engine to perform centralized updates. For the MPE environment, the actions of other agents are observable only during training and not during execution (CTDE).

Table 7 contains a summary of all of our experimental settings along with the associated configuration of advisors.

### K.1 Pommerman

In Pommerman, the complete set of skills needed to be learned in order to win games include 1) escaping from the enemy, 2) obtaining power ups (bombs/life), 3) killing the enemy, 4) blasting walls to open routes, and 5) coordinating with a teammate (in the team version) [30].

In our experiments, we consider two Pommerman domains. The first four experiments use the two-agent version of Pommerman and the fifth experiment uses the four-agent team version of Pommerman. Each episode in our training and execution experiments corresponds to a full Pommerman game with a randomized board and a maximum of 800 steps before completion. The game ends either when all the steps are completed or when one of the two Pommerman agents dies (two-agent version). In the team version, the game ends either when all the steps are complete (800 steps maximum) or when one of the two teams completely dies.

The first experiment (Experiment 1) uses four sufficient advisors of varying quality. The first advisor (Advisor 1) can teach all the strategies (i.e., various Pommerman skills as mentioned in Section 6) needed to win the Pommerman game. Advisor 2 can teach moves associated with killing the opponent if the opponent is very close to the agent and defensive strategies that help avoid the enemy. The third advisor (Advisor 3) can only teach defensive strategies that

Exp.	Domain	Type	Advisors	# of training agents
1	Two-agent Pommerman	Competitive	4 sufficient advisors with different quality	2
2	Two-agent Pommerman	Competitive	4 sufficient advisors with similar quality	2
3	Two-agent Pommerman	Competitive	4 insufficient advisors with different quality	2
4	Two-agent Pommerman	Competitive	4 insufficient advisors with similar quality	2
5	Four-agent Pommerman	Mixed	4 sufficient advisors with different quality	4
6	Pursuit SISL	Cooperative	4 insufficient advisors with different quality	8
7	Predator-Prey MPE	Mixed (CTDE)	4 insufficient advisors with different quality	8

**Table 7: Description of experimental settings**

help in escaping the enemy, and cannot teach aggressive strategies needed to kill the enemy. Finally, the last advisor (Advisor 4) only provides random actions. Hence, all the agents should follow the first advisor as far as possible, and the fourth advisor must be avoided entirely.

In the second experiment (Experiment 2), we use a new set of four advisors. Here, the first advisor teaches only defensive skills (escaping the enemy), the second advisor can only teach aggressive skills (killing the enemy), the third advisor can only teach strategies that enable obtaining the power ups, and the fourth advisor teaches ways to seek and blast open wooden walls which opens up various paths in the game. It can be seen here that no one advisor is can teach all the strategies needed to win in Pommerman. However,

together, all four advisors can teach the requisite strategies, and they need to be leveraged appropriately.

In Experiment 3, we use a set of four advisors of decreasing quality as in the first experiment; however, none of the advisors are can teach strategies that seek enemies and kill them (insufficient set). Also, none of the advisors are capable of teaching the skills needed to seek wooden walls to blast open. Hence, this set of advisors is insufficient for winning the Pommerman game. The first advisor (Advisor 1) can teach strategies to escape from an enemy, kill the enemy if the enemy is right next to the agent, and obtaining the power-ups. The second advisor (Advisor 2) can only teach strategies that help in escaping the enemy or killing an enemy close to the agent. Advisor 3 can only teach strategies that pertain to escaping the enemy, and Advisor 4 only provides random suggestions.

In Experiment 4, we have a set of advisors similar to the second experiment; however, the advisors are incapable of teaching sufficient skills needed to win in Pommerman. The first advisor (Advisor 1) teaches only defensive skills to escape an enemy. The second advisor (Advisor 2) helps in learning a strategy that can kill an enemy right next to the agent. The third advisor (Advisor 3) can teach strategies that lead to obtaining the power-ups, and the fourth advisor (Advisor 4) can teach skills needed to blast open wooden walls, if the agent is very close to the wall.

The fifth experiment (Experiment 5) with the team domain uses the same set of advisors as Experiment 1.

The Pommerman environment was released by Resnick et al. [30] under the Apache2 license.

## K.2 Pursuit domain

The pursuit domain was first introduced by Gupta et al. [8], and we use the implementation provided by the Petting Zoo environment [42] (released under MIT license). The game has a set of 8 pursuer agents cooperating with each other to capture a set of 30 evaders in the environment. We use the same reward function and environmental parameters as in Terry et al. [42] with some minor modifications. In our setting, the agents get a reward of +1 for hitting (tagging) the evaders and a reward of +30 for catching an evader. An evader is decided to be caught if it is surrounded by a group of at least two pursuer agents. The captured evaders are removed from the environment. All agents get an urgency penalty of -0.1 at each time step of the game. Each episode has a maximum of 500 steps. The episode terminates either when all the evaders are captured or when all the steps are completed. All the pursuers receive the same reward at all time steps (global reward structure), where the rewards are distributed amongst all agents. Each pursuer observes a  $7 \times 7$  grid around itself (the entire grid is  $16 \times 16$ ), which means that the pursuer can get full information about other pursuers and evaders within this observable grid and no information outside this grid.

## K.3 Multi Particle Environment (Predator-Prey)

The Multi Particle Environment (MPE) was originally released by Lowe et al. [24] as a set of testbeds for the purpose of testing algorithms that pertain to cooperative and mixed cooperative-competitive settings with characteristics of communication and

obstacle interactions. From this suite of testbeds we use the Simple Tag environment that pertains to a Predator-Prey setting where a set of three predators try to capture a single prey. We use the environment defined as a part of the Petting Zoo library [42] (released under MIT license). Here all agents have a discrete action space with a set of 5 actions (four cardinal directions and one action that signifies no movement). Both the predators and prey have a continuous observation space that corresponds to the velocity and position of all agents including the agent itself. The predators get a reward of +10 for hitting (colliding/tagging) prey and the prey get a punishment of -10 for being hit by any predator (the prey is not removed from the environment). The prey also receives a small additional penalty for exiting the field of play (see Terry et al. [42] for more details). All the predators receive the same reward at all time steps (global reward structure), where the rewards are distributed amongst all predators. Our environment contains eight predators and eight prey, in addition to five obstacles that block the path of the predators and prey. Each game contains 500 steps of training or execution. We model this domain as a CTDE setting, where the actions taken, and the rewards obtained by all agents are available to each agent during training, but not available during execution.

## L HYPERPARAMETERS AND IMPLEMENTATION DETAILS

All the hyperparameters in our implementation of baselines are either the same or closely match the values recommended by the respective papers that introduced these algorithms. Our algorithms also use similar hyperparameters as the baseline algorithms, with a few exceptions (for performance and computational efficiency reasons).

The DQN [26], CHAT [46], TLQL [21], ADMIRAL-DM [37], and MA-TLQL implementations use almost the same hyperparameters. These algorithms use a learning rate of 0.01, a discount factor of 0.9, a replay memory size of  $2 \times 10^6$ , and a fixed exploration rate of 0.9. The target network is replaced every 10 learning iterations using the hard replacement strategy. The evaluation and target networks use 3 fully connected layers (2 ReLU layers of 50 neurons and an output layer). We use a batch size of 32.

For the CHAT implementation, we use the neural network based confidence variant (NNHAT) from Wang and Taylor [46]. We set a confidence threshold of 0.6 and use 3 fully connected layers (2 layers using ReLU as the activation function with 50 neurons and an additional output layer). The advisors are directly used in CHAT instead of preparing decision based rules from classifier models as done in [46] due to performance reasons and also because the advisors used in our experiments are either rule-based agents or pretrained networks and not human advisors as designed in CHAT. Since these advisors are considered to be extracted rules in CHAT (not actual demonstrators/advisors), we allow CHAT dependence on advisors in the execution phase as well.

We use the PPR technique in TLQL for our experiments, though this is not used in Li et al. [21]. This is done due to two reasons. First, unlike in single-agent settings, independent  $Q$ -learning methods do not have a policy improvement guarantee in multi-agent settings (as discussed in Section 4) [40], hence, the vanilla TLQL is not

guaranteed to stop depending on advisors as motivated by Li et al. [21]. Second, without PPR, the experimental TLQL training performance is very good (since it has unlimited dependence on advisors), while the execution performances are very poor (since advisors are not available during execution). The best execution performances for TLQL is obtained while using PPR.

Regarding the implementation of the PPR technique, the TLQL, ADMIRAL-DM, MA-TLAC, and MA-TLQL implementations start with a value of  $\epsilon' = 1$ , which is linearly decayed to 0 at the end of training. There is no influence of advisors during execution, and hence,  $\epsilon' = 0$  during execution.

To stay consistent with the description in Li et al. [21], the TLQL implementation uses a total of 3 networks (evaluation and target networks for low- $Q$  in addition to high- $Q$ ). The high- $Q$  and low- $Q$  networks use fully connected layers with the same architecture as described for the DQN. The MA-TLQL implementation uses four networks (two evaluation and target networks for the low- $Q$  and high- $Q$  respectively), also having the same configuration as described for the DQN. Further, the MA-TLQL uses a second replay buffer for the high- $Q$  in addition to the replay buffer of the low- $Q$ . Both buffers use the same memory size of  $2 \times 10^6$ .

Regarding DQfD [12], we set  $1 \times 10^6$  as the demo buffer size and perform 50,000 mini-batch updates for pretraining. The replay buffer size is twice the size of the demo buffer. The  $N$ -step return weight is 1.0, the supervised loss weight is 1.0 and the L2 regularization weight is  $10^{-5}$ . The epsilon greedy exploration is 0.9. The discount factor is 0.99 and the learning rate is 0.002. The network architecture uses 3 fully connected layers (2 ReLU layers of 24 neurons and an output layer). The pretraining for DQfD comes from a data buffer related to a series of games where the advisors compete against each other.

The MA-TLAC uses two actor networks and two critic networks. The network architecture is the same as described for MA-TLQL. The actor networks use a learning rate of  $10^{-6}$ , and the critic networks use a learning rate of  $10^{-3}$ .

For all training experiments, we use a set of 30 random seeds (1 – 30). We use a new set of 30 random seeds (31 – 60) for the execution experiments.

## M WALL CLOCK TIMES

All the training for the experiments were conducted on a virtual machine having 2 Nvidia A100 GPUs with a GPU memory of 40 GB. The CPUs use the AMD EPYC processors with a memory of 125 GB. The Pommerman experiments took an average of 12 hours wall clock time to complete, and the Pursuit experiments took an average of 15 hours wall clock time to complete.