
Near-optimal Policy Optimization Algorithms for Learning Adversarial Linear Mixture MDPs

Jiafan He
UCLA

Dongruo Zhou
UCLA

Quanquan Gu
UCLA

Abstract

Learning Markov decision processes (MDPs) in the presence of the adversary is a challenging problem in reinforcement learning (RL). In this paper, we study RL in episodic MDPs with adversarial reward and full information feedback, where the unknown transition probability function is a linear function of a given feature mapping, and the reward function can change arbitrarily episode by episode. We propose an optimistic policy optimization algorithm POWERS and show that it can achieve $\tilde{O}(dH\sqrt{T})$ regret, where H is the length of the episode, T is the number of interaction with the MDP, and d is the dimension of the feature mapping. Furthermore, we also prove a matching lower bound of $\tilde{\Omega}(dH\sqrt{T})$ up to logarithmic factors. Our key technical contributions are two-fold: (1) a new value function estimator based on importance weighting; and (2) a tighter confidence set for the transition kernel. They together lead to the nearly minimax optimal regret.

1 INTRODUCTION

The goal of reinforcement learning (RL) is to design a policy to maximize the reward through observation from interaction with the unknown environment. In reinforcement learning, the Markov decision process (MDP) (Puterman, 1994) is a typical model to describe the unknown environment and widely used to analyze the sequential dynamic environment. In this work, we consider episodic MDPs with a finite horizon. Traditional MDPs often assume the unknown

transition probability function is fixed and the reward function is stochastic, which means the reward of each state-action pair follows an unknown stationary distribution. Yet, in many real world models, the reward function is not fixed and may change over time. In order to capture the changed or even adversarial reward, Even-Dar et al. (2009) first introduced the concept of adversarial MDP model and proposed MDP-Expert (MDP-E) algorithm, which attains $\tilde{O}(\tau^2\sqrt{T})$ regret with τ being the mixing time of the MDP, for known transition probability function and full information of the reward function. In a concurrent work, Yu et al. (2009) proposed an algorithm in the same setting and obtained $\tilde{O}(T^{2/3})$ regret. There is a line of follow up work studying RL for adversarial MDPs (Neu et al., 2010, 2012; Zimin and Neu, 2013; Dekel and Hazan, 2013; Rosenberg and Mansour, 2019a; Efroni et al., 2020), which studies various settings depending on whether the transition probability function is known, and whether the feedback is full-information or bandit. Please see the related work section for a more detailed discussion.

However, most existing works on adversarial MDP are in the tabular MDP setting, where both the number of actions and states are finite, and the action-value function is represented by a table. In many real-world RL problems, the state and action spaces are large or even infinite. A widely used method to overcome the curse of large state and action spaces is function approximation, which reparameterizes the tabular action-value function as a function over some feature mapping that maps the state and action to a low-dimensional space. Learning adversarial MDPs with linear function approximation is still understudied. Some existing works (Tamar et al., 2014; Zhang et al., 2020; Wang and Zou, 2021) study learning optimal policies for robust MDPs with function approximation under specific MDP assumptions, which are not directly applicable to general adversarial MDPs. A notable existing work is Cai et al. (2020), which studies general adversarial MDPs with linear function approximation. In particular, Cai et al. (2020) proposed an optimistic variant of proximal policy optimization algorithm for the linear kernel

MDP (Jia et al., 2020; Ayoub et al., 2020; Zhou et al., 2021b) with unknown transition probability and full reward information in the adversarial setting, which achieves $\tilde{O}(\sqrt{d^2 H^3 T})$ regret. Here H is the length of the episode, T is the number of interactions with the MDP and d is the dimension of the feature mapping.

In this paper, we seek a computationally efficient and statistically optimal algorithm for learning adversarial MDPs. The focus of this work is the unknown transition and full information setting. We first propose an algorithm called optimistic Policy Optimization With BERNstein bonuS (POWERS) for adversarial linear mixture MDP (See Definition for more details) with full information feedback. At a high level, our algorithm POWERS is similar to Optimistic-PPO (OPPO) algorithm (Cai et al., 2020), which can also be seen as an extension of MDP-Expert (MDP-E) with linear function. More specifically, POWERS consists of two main steps in each round: (1) one-step least-square temporal difference (LSTD) learning along with exploration bonus for policy evaluation; and (2) mirror descent on the policy space for policy improvement. Our key algorithmic contributions include a weighted LSTD algorithm which takes into the variance of the Bellman residue into account, and a Bernstein-type bonus for exploration based on the principle of “optimism-in-the-face-of-uncertainty” (Abbasi-Yadkori et al., 2011). We prove that POWERS achieves $\tilde{O}(dH\sqrt{T})$ regret with high probability, where H is the length of the episode, T is the number of interactions with the MDP and d is the dimension of the feature mapping. We also prove an $\tilde{\Omega}(dH\sqrt{T})$ lower bound for adversarially learning linear kernel MDPs. Our upper bound matches the lower bound up to logarithmic factors, which suggests that our algorithm is nearly minimax optimal. To the best of our knowledge, our algorithm is the first computationally efficient and statistical (nearly) optimal algorithm for adversarial MDPs in the unknown transition and full reward information setting.

Notation. We use lower case letters to denote scalars, and use lower and upper case boldface letters to denote vectors and matrices respectively. For a vector $\mathbf{x} \in \mathbb{R}^d$ and matrix $\Sigma \in \mathbb{R}^{d \times d}$, we denote by $\|\mathbf{x}\|_2$ the Euclidean norm, $\|\mathbf{x}\|_1 = \sum_{i=1}^d |x_i|$, and $\|\mathbf{x}\|_\Sigma = \sqrt{\mathbf{x}^\top \Sigma \mathbf{x}}$. For two sequences $\{a_n\}$ and $\{b_n\}$, we write $a_n = O(b_n)$ if there exists an absolute constant C such that $a_n \leq Cb_n$, and we write $a_n = \Omega(b_n)$ if there exists an absolute constant C such that $a_n \geq Cb_n$. We use $\tilde{O}(\cdot)$ and $\tilde{\Omega}(\cdot)$ to further hide the logarithmic factors. For any $a \leq b \in \mathbb{R}$, $x \in \mathbb{R}$, let $[x]_{[a,b]}$ denote $a \cdot \mathbf{1}(x \leq a) + x \cdot \mathbf{1}(a \leq x \leq b) + b \cdot \mathbf{1}(b \leq x)$, where $\mathbf{1}(\cdot)$ is the indicator function. For a positive integer n , we use $[n] = \{1, 2, \dots, n\}$ to denote the set of integers from

1 to n .

2 RELATED WORK

RL with adversarial reward. There is a long line of research on learning adversarial MDPs, where the reward function is adversarially chosen at the beginning of each episode and can change arbitrarily across different episodes (Even-Dar et al., 2009; Yu et al., 2009; Gergely Neu et al., 2010; Neu et al., 2010; Zimin and Neu, 2013; Neu et al., 2012; Rosenberg and Mansour, 2019a,b; Wang et al., 2019; Cai et al., 2020; Efroni et al., 2020). The seminal works by Even-Dar et al. (2009); Yu et al. (2009) are in the known transition probability and full reward information setting. In the known transition and bandit feedback on the reward setting, Gergely Neu et al. (2010) proposed MDP-EXP3 algorithm and obtained $\tilde{O}(T^{2/3})$ regret. Neu et al. (2010) proposed Bandit O-SSP algorithm which achieves $\tilde{O}(\sqrt{T}/\alpha)$ regret with an additional assumption that all states are reachable with probability $\alpha > 0$ for any policy. Zimin and Neu (2013) further proposed Q-REPS algorithm, which improves the regret from $\tilde{O}(T^{2/3})$ to $\tilde{O}(\sqrt{T})$ without any additional assumption. In the unknown transition but full reward information setting, Neu et al. (2012) proposed FPOP algorithm that achieves $\tilde{O}(SA\sqrt{T})$ regret. Rosenberg and Mansour (2019a) proposed UC-O-REP algorithm and improved the regret to $\tilde{O}(S\sqrt{AT})$. In the most challenging unknown transition and bandit reward feedback setting, Rosenberg and Mansour (2019b) proposed Shifted Bandit UC-O-REPS algorithm which achieves $\tilde{O}(T^{3/4})$ regret. Rosenberg and Mansour (2019b) also proposed Bounded Bandit UC-O-REPS algorithm and obtained $\tilde{O}(\sqrt{T}/\alpha)$ regret under the assumption that all states are reachable with probability $\alpha > 0$ for any policy. Jin et al. (2020a) proposed UOB-REPS algorithm that achieves $\tilde{O}(\sqrt{T})$ regret without the additional assumption made by Rosenberg and Mansour (2019b). The focus of this paper is the unknown transition but full reward information setting.

RL with linear function approximation. Recently, there emerges a large body of literature on solving MDP with linear function approximation. These works can be generally divided into three lines based on the specific assumption on the underlying MDP. The first line of work (Sun et al., 2019; Du et al., 2019) is based on the low Bellman rank assumption (Jiang et al., 2017), which assumes a low-rank factorization of the Bellman error matrix. The second line of work (Wang et al., 2019; He et al., 2021; Zanette et al., 2020a) focuses on the linear MDP (Yang and Wang, 2019a; Jin et al., 2020b), where the transition probability function and reward function are pa-

parameterized as a linear function of a feature mapping $\phi : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}^d$. Later, Zanette et al. (2020b) made a weaker assumption called low inherent Bellman error and proposed Eleanor algorithm. The last line of work (Cai et al., 2020; Yang and Wang, 2019b; He et al., 2021; Modi et al., 2019; Zhou et al., 2021a) is based on the linear mixture/kernel MDP (Jia et al., 2020; Ayoub et al., 2020; Zhou et al., 2021b; Wu et al., 2022), where the transition probability function can be parameterized as a linear function of a feature mapping $\phi : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow \mathbb{R}^d$. Note that none of the above work with linear function approximation can handle adversarially chosen reward with Cai et al. (2020) being a notable exception. Our paper also considers the linear kernel MDP but with an adversarial reward function.

RL with policy gradient. Our work is also related to policy optimization and policy gradient methods (Williams, 1992; Baxter and Bartlett, 2000; Sutton et al., 1999; Kakade, 2001; Kakade and Langford, 2002; Kakade et al., 2003; Bagnell and Schneider, 2003; Schulman et al., 2015, 2017; Abbasi-Yadkori et al., 2019a,b; Cai et al., 2020; Hao et al., 2020; Efroni et al., 2020), among which the most related works to ours are trust-region policy optimization (Schulman et al., 2015), proximal policy optimization (Schulman et al., 2017), Politex (Abbasi-Yadkori et al., 2019a), EE-Politex (Abbasi-Yadkori et al., 2019b), AAPI (Hao et al., 2020) and OPPO (Cai et al., 2020). More specifically, Cai et al. (2020) proposed the optimistic variant of the Proximal Policy Optimization algorithm for adversarial linear kernel MDP, which can be seen as an extension of MDP-E. Abbasi-Yadkori et al. (2019a) proposed Politex algorithm with least-squares policy evaluation for infinite-horizon average-reward MDPs, which can be seen a generalization of the MDP-E (Even-Dar et al., 2009). In fact, MDP-E is equivalent to TRPO/PPO (Schulman et al., 2015, 2017) as shown by Neu et al. (2017). Our algorithm can be seen as a nontrivial extension of OPPO and MDP-E.

3 PRELIMINARIES

Time-inhomogeneous, episodic adversarial MDPs. In this paper, we consider a time-inhomogeneous, episodic Markov decision process (MDP), which is denoted by a tuple $M = M(\mathcal{S}, \mathcal{A}, H, \{r_h^k\}_{h \in [H], k \in [K]}, \{\mathbb{P}_h\}_{h=1}^H)$. Here \mathcal{S} is the state space, \mathcal{A} is the action space, H is the length of the episode, $r_h^k : \mathcal{S} \times \mathcal{A} \rightarrow [0, 1]$ is the deterministic reward function at stage h of the k -th episode. $\mathbb{P}_h(s'|s, a)$ is the transition probability function which denotes the probability for state s to transfer to state s' given action a at stage h . For simplicity, we assume the reward function r_h^k is adversarially chosen by the environment at the beginning of the k -th episode and

known after the episode k . A policy $\pi = \{\pi_h\}_{h=1}^H$ is a collection of functions π_h , where each $\pi_h : \mathcal{S} \rightarrow \Delta(\mathcal{A})$ is a function which maps a state s to distributions over action set \mathcal{A} at stage h . For each state-action pair $(s, a) \in \mathcal{S} \times \mathcal{A}$, we denote the action-value function $Q_{k,h}^\pi$ and the value function $V_{k,h}^\pi$ as follows:

$$\begin{aligned} Q_{k,h}^\pi(s, a) &= r_h^k(s, a) \\ &\quad + \mathbb{E} \left[\sum_{h'=h+1}^H r_{h'}^k(s_{h'}, a_{h'}) \mid s_h = s, a_h = a \right], \\ V_{k,h}^\pi(s) &= \mathbb{E}_{a \sim \pi_h(\cdot|s)} [Q_{k,h}^\pi(s, a)], V_{k,H+1}^\pi(s) = 0. \end{aligned}$$

In the definition of $Q_{k,h}^\pi$, we denote by $\mathbb{E}[\cdot]$ the expectation over the state-action sequences $(s_h, a_h, s_{h+1}, a_{h+1}, \dots, s_H, a_H)$, where $s_h = s, a_h = a$ and $s_{h'+1} \sim \mathbb{P}_h(\cdot|s_{h'}, a_{h'})$, $a_{h'+1} \sim \pi_{h'+1}(\cdot|s_{h'+1})$ ($h' = h, h+1, \dots, H-1$). For simplicity, for any function $V : \mathcal{S} \rightarrow \mathbb{R}$, we denote

$$\begin{aligned} [\mathbb{P}_h V](s, a) &= \mathbb{E}_{s' \sim \mathbb{P}_h(\cdot|s, a)} V(s'), \\ [\mathbb{V}_h V](s, a) &= [\mathbb{P}_h V^2](s, a) - ([\mathbb{P}_h V](s, a))^2, \end{aligned} \quad (3.1)$$

where V^2 is a shorthand for the function whose value at state s is $(V(s))^2$. Using this notation, for policy π , we have the following Bellman equality $Q_{k,h}^\pi(s, a) = r_h^k(s, a) + [\mathbb{P}_h V_{k,h+1}^\pi](s, a)$.

In the *online learning setting*, for each episode $k \geq 1$, at the beginning of the episode k , the agent determines a policy π^k to be followed in this episode and we assume that the initial state s_1^k is fixed across all episodes¹. At each stage $h \in [H]$, the agent observe the state s_h^k , choose an action following the policy $a_h^k \sim \pi_h^k(\cdot|s_h^k)$ and observe the next state with $s_{h+1}^k \sim \mathbb{P}_h(\cdot|s_h^k, a_h^k)$. For the adversarial environment, we focus on the expected regret, which is the expected loss of the algorithm relative to the best-fixed policy in hindsight (Cesa-Bianchi and Lugosi, 2006):

$$\text{Regret}(M, K) = \sup_{\pi} \sum_{k=1}^K (V_{k,1}^\pi(s_1^k) - V_{k,1}^{\pi^k}(s_1^k)).$$

For simplicity, we denote the optimal policy π^* as $\pi^* = \sup_{\pi} \sum_{k=1}^K V_{k,1}^\pi(s_1^k)$. Therefore, we have the following Bellman optimally equation $Q_{k,h}^*(s, a) = r_h^k(s, a) + [\mathbb{P}_h V_{k,h}^*](s, a)$, where $Q_{k,h}^*(s, a), V_{k,h}^*(s, a)$ are the corresponding optimal action-value function and value function. For any two policies π and π' , we define the Kullback–Leibler divergence between them as follows $D_{KL}(\pi \parallel \pi') = \sum_{a \in \mathcal{A}} \pi(a) \log(\pi(a)/\pi'(a))$.

Linear Mixture MDPs. In this work, we focus on a special class of MDPs called *linear mixture MDPs* (Jia et al., 2020; Ayoub et al., 2020; Zhou et al., 2021b),

¹While we study fixed initial state, our results readily can generalize to the case that the initial state s_1^k is random chosen from a fixed distribution across all episodes.

where the transition probability function is a linear function of a given feature mapping $\phi : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow \mathbb{R}^d$. The formal definition of a linear kernel MDP is as follows:

Definition 3.1. $M(\mathcal{S}, \mathcal{A}, H, \{r_h^k\}_{h \in [H], k \in [K]}, \{\mathbb{P}_h\}_{h=1}^H)$ is called a inhomogenous, episode B -bounded linear mixture MDP if there exist a *known* feature mapping $\phi(s'|s, a) : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow \mathbb{R}^d$ and an *unknown* vector $\theta_h \in \mathbb{R}^d$ with $\|\theta\|_2 \leq B$ ², such that

- For any state-action-next-state triplet $(s, a, s') \in \mathcal{S} \times \mathcal{A} \times \mathcal{S}$, we have $\mathbb{P}_h(s'|s, a) = \langle \phi(s'|s, a), \theta_h \rangle$;
- For any bounded function $V : \mathcal{S} \rightarrow [0, 1]$ and any tuple $(s, a) \in \mathcal{S} \times \mathcal{A}$, we have $\|\phi_V(s, a)\|_2 \leq 1$, where $\phi_V(s, a) = \sum_{s' \in \mathcal{S}} \phi(s'|s, a)V(s') \in \mathbb{R}^d$.

Based on Definition 3.1, we can see that for any linear mixture MDP M and function $V : \mathcal{S} \rightarrow \mathbb{R}$, we have the following properties:

$$\begin{aligned} [\mathbb{P}_h V](s, a) &= \sum_{s' \in \mathcal{S}} \mathbb{P}_h(s'|s, a)V(s') \\ &= \sum_{s' \in \mathcal{S}} \langle \phi(s'|s, a), \theta_h \rangle V(s') \\ &= \langle \phi_V(s, a), \theta_h \rangle, \end{aligned} \quad (3.2)$$

and

$$\begin{aligned} [\mathbb{V}_h V](s, a) &= \sum_{s' \in \mathcal{S}} \mathbb{P}_h(s'|s, a)V^2(s') \\ &\quad - [\sum_{s' \in \mathcal{S}} \mathbb{P}_h(\cdot|s, a)V(s')]^2 \\ &= \langle \phi_{V^2}(s, a), \theta_h \rangle - [\langle \phi_V(s, a), \theta_h \rangle]^2. \end{aligned} \quad (3.3)$$

(3.2) and (3.3) suggest that both the conditional expectation and the variance of a function V can be calculated based on certain linear functions of different feature mappings, i.e., ϕ_V and ϕ_{V^2} . Therefore, we can estimate them by estimating the corresponding linear functions.

4 THE PROPOSED ALGORITHM

In this section, we propose an algorithm POWERS to learn the episodic linear mixture MDP (see Definition 3.1) with adversarial rewards, which is illustrated in Algorithm 1. At a high level, POWERS is an improved version of the Optimistic-PPO (OPPO) algorithm (Cai et al., 2020) with a refined estimate of the action-value function $Q_{k,h}(s, a)$. The POWERS can be divided into two phases: (1) policy improvement phase and (2) policy evaluation phase.

Policy improvement phase (Line 5 to Line 10):

In the policy improvement phase, POWERS calculates

²For any MDP M , parameter B cannot be arbitrarily small. More specifically, parameter B should satisfy $B \geq 1$.

its policy π^k for the current episode, based on its previous policy π^{k-1} using the proximal policy optimization (PPO) method (Schulman et al., 2017). In detail, let s_1^k be the starting state at the k -th episode, then following PPO, we update π^k as a solution to the following optimization problem:

$$\pi^k \leftarrow \operatorname{argmax}_{\pi} [L_{k-1}(\pi) - \alpha^{-1} \tilde{D}_{KL}(\pi, \pi^{k-1})], \quad (4.1)$$

where

$$L_{k-1}(\pi) = \mathbb{E}_{\pi^{k-1}} \left[\sum_{h=1}^H \langle Q_{k-1,h}(s_h, \cdot), \pi_h^k(\cdot|s_h) \rangle \Big| s_1 = s_1^k \right]$$

is proportional to the first-order Taylor approximation of $V_{k-1,h}^{\pi^{k-1}}$ at π^{k-1} , and replaces the action-value function $Q_{k-1,h}^{\pi^{k-1}}(\cdot, \cdot)$ by the estimated one $Q_{k-1,h}(\cdot, \cdot)$, and

$$\begin{aligned} \tilde{D}_{KL}(\pi, \pi^{k-1}) &= \mathbb{E}_{\pi^{k-1}} \left[\sum_{h=1}^H D_{KL}(\pi_h(\cdot|s_h), \pi_h^k(\cdot|s_h)) \Big| s_1 = s_1^k \right] \end{aligned}$$

is the sum of KL-divergences between π_h and π_h^{k-1} , which encourages π^k to stay close to π^{k-1} to ensure the above first-order Taylor approximation is accurate enough. The closed-form solution to (4.1) is in Line 7. Here $\alpha > 0$ is the step size of the exponential update. Note that the update rule in Line 7 is also the same as the MDP-E algorithm (Even-Dar et al., 2009). After obtaining π^k , POWERS chooses action a_h^k based on the new policy π_h^k and the current state s_h^k . It then observes the next state s_{h+1}^k and the adversarial reward function $r_h^k(\cdot, \cdot)$.

Policy evaluation phase (Line 11 to Line 23):

In the policy evaluation phase, POWERS evaluates the policy π^k by constructing the action-value function $Q_{k,h}$ and the value function $V_{k,h}$ for policy π^k based on the observed data, which are optimistic estimates of the action-value function $Q_{k,h}^{\pi^k}$ and the value function $V_{k,h}^{\pi^k}$ respectively.

Specifically, for each episode $k \in [K]$ and each stage $h \in [H]$, POWERS maintains an estimator $\hat{\theta}_{k,h}$ and an uncentered covariance matrix $\hat{\Sigma}_{k,h}$ based on the observed data before the k -th episode. Then POWERS recursively computes the optimistic $Q_{k,h}, V_{k,h}$ as follows:

$$\begin{aligned} Q_{k,h}(s, a) &= \left[r_h^k(s, a) + \langle \hat{\theta}_{k,h}, \phi_{V_{k,h+1}}(s, a) \rangle \right. \\ &\quad \left. + \hat{\beta}_k \|\hat{\Sigma}_{k,h}^{-1/2} \phi_{V_{k,h+1}}(s, a)\|_2 \right]_{[0, H-h+1]}, \\ V_{k,h}(s) &= \mathbb{E}_{a \sim \pi_h^k(\cdot|s)} [Q_{k,h}(s, a)], \end{aligned}$$

Algorithm 1 POWERS

Require: Regularization parameter λ , learning rate α .

- 1: Set initial policy $\{\pi_h^0(\cdot|\cdot)\}_{h=1}^H$ as uniform distribution on the action set \mathcal{A}
- 2: For $h \in [H+1]$, set the initial value functions $Q_{0,h}(\cdot, \cdot) \leftarrow 0, V_{0,h}(\cdot) \leftarrow 0$
- 3: For $h \in [H]$, set $\widehat{\Sigma}_{1,h}, \widetilde{\Sigma}_{1,h} \leftarrow \lambda \mathbf{I}, \widehat{\mathbf{b}}_{1,h}, \widetilde{\mathbf{b}}_{1,h} \leftarrow \mathbf{0}, \widehat{\boldsymbol{\theta}}_{1,h}, \widetilde{\boldsymbol{\theta}}_{1,h} \leftarrow \mathbf{0}$
- 4: **for** $k = 1, \dots, K$ **do**
- 5: Receive state s_1^k
- 6: **for** $h = 1, \dots, H$ **do**
- 7: Update the policy by $\pi_h^k(\cdot|\cdot) \propto \pi_h^{k-1}(\cdot|\cdot) \exp\{\alpha Q_{k-1,h}(\cdot, \cdot)\}$
- 8: Take action $a_h^k \sim \pi_h^k(\cdot|s_h^k)$ and receive next state $s_{h+1}^k \sim \mathbb{P}_h(\cdot|s_h^k, a_h^k)$
- 9: Observe the adversarial reward function $r_h^k(\cdot, \cdot)$
- 10: **end for**
- 11: Set $V_{k,H+1}(\cdot) \leftarrow 0$
- 12: **for** $h = H, \dots, 1$ **do**
- 13: Set $Q_{k,h}(\cdot, \cdot) \leftarrow \left[r_h^k(\cdot, \cdot) + \langle \widehat{\boldsymbol{\theta}}_{k,h}, \boldsymbol{\phi}_{V_{k,h+1}}(\cdot, \cdot) \rangle + \widehat{\beta}_k \|\widehat{\Sigma}_{k,h}^{-1/2} \boldsymbol{\phi}_{V_{k,h+1}}(\cdot, \cdot)\|_2 \right]_{[0, H-h+1]}$
- 14: Set $V_{k,h}(\cdot) \leftarrow \mathbb{E}_{a \sim \pi_h^k(\cdot|\cdot)}[Q_{k,h}(\cdot, a)]$
- 15: Set the estimated variance $[\widehat{\mathbb{V}}_{k,h} V_{k,h+1}](s_h^k, a_h^k)$ as in (4.2)
- 16: Set the bonus term $E_{k,h}$ as in (4.3)
- 17: $\bar{\sigma}_{k,h} \leftarrow \sqrt{\max\{H^2/d, [\widehat{\mathbb{V}}_{k,h} V_{k,h+1}](s_h^k, a_h^k) + E_{k,h}\}}$
- 18: $\widehat{\Sigma}_{k+1,h} \leftarrow \widehat{\Sigma}_{k,h} + \bar{\sigma}_{k,h}^{-2} \boldsymbol{\phi}_{V_{k,h+1}}(s_h^k, a_h^k) \boldsymbol{\phi}_{V_{k,h+1}}(s_h^k, a_h^k)^\top$
- 19: $\widehat{\mathbf{b}}_{k+1,h} \leftarrow \widehat{\mathbf{b}}_{k,h} + \bar{\sigma}_{k,h}^{-2} \boldsymbol{\phi}_{V_{k,h+1}}(s_h^k, a_h^k) V_{k,h+1}(s_{h+1}^k)$
- 20: $\widetilde{\Sigma}_{k+1,h} \leftarrow \widetilde{\Sigma}_{k,h} + \boldsymbol{\phi}_{V_{k,h+1}^2}(s_h^k, a_h^k) \boldsymbol{\phi}_{V_{k,h+1}^2}(s_h^k, a_h^k)^\top$
- 21: $\widetilde{\mathbf{b}}_{k+1,h} \leftarrow \widetilde{\mathbf{b}}_{k,h} + \boldsymbol{\phi}_{V_{k,h+1}^2}(s_h^k, a_h^k) V_{k,h+1}^2(s_{h+1}^k)$
- 22: $\widehat{\boldsymbol{\theta}}_{k+1,h} \leftarrow \widehat{\Sigma}_{k+1,h}^{-1} \widehat{\mathbf{b}}_{k+1,h}, \widetilde{\boldsymbol{\theta}}_{k+1,h} \leftarrow \widetilde{\Sigma}_{k+1,h}^{-1} \widetilde{\mathbf{b}}_{k+1,h}$
- 23: **end for**
- 24: **end for**

where $\widehat{\beta}_k$ is the radius of the confidence ball defined as:

$$\widehat{\beta}_k = 8\sqrt{d \log(1+k/\lambda) \log(4k^2 H/\delta)} + 4\sqrt{d} \log(4k^2 H/\delta) + \sqrt{\lambda} B.$$

Now we illustrate how to construct the estimator $\widehat{\boldsymbol{\theta}}_{k,h}$ and the covariance matrix $\widehat{\Sigma}_{k,h}$, which is the key difference compared with OPPO proposed in Cai et al. (2020). Recall (3.2), we know that the expectation of the random variables $V_{k,h}(s_{k,h+1})$ can be written as a linear function with weight vector $\boldsymbol{\theta}_h$. Therefore, a natural way to estimate $\boldsymbol{\theta}_h$ is to consider it as the unknown weight vector of a stochastic linear bandits problem with context $\boldsymbol{\phi}_{V_{k,h}}(s_h^k, a_h^k)$ and target $V_{k,h}(s_{k,h+1})$, and apply algorithms for linear bandits such as OFUL (Abbasi-Yadkori et al., 2011), to obtain the estimator $\widehat{\boldsymbol{\theta}}_{k,h}$. Such an approach is adopted by Cai et al. (2020).

However, OFUL uses the vanilla linear regression to construct $\widehat{\boldsymbol{\theta}}_{k,h}$, which is limited to the *homoscedastic* noises case. For the linear mixture MDP, the noises are actually *heteroscedastic* as each target enjoys different noises. Thus the vanilla linear regression is known as

statistically inefficient (Kirschner and Krause, 2018). Inspired by Kirschner and Krause (2018); Zhou et al. (2021a), we adopt the *weighted linear regression* to construct $\widehat{\boldsymbol{\theta}}_{k,h}$, which is the solution to the following weighted regression problem:

$$\widehat{\boldsymbol{\theta}}_{k,h} \leftarrow \arg \min_{\boldsymbol{\theta} \in \mathbb{R}^d} \lambda \|\boldsymbol{\theta}\|_2^2 + \sum_{i=1}^{k-1} [\langle \boldsymbol{\phi}_{V_{i,h+1}}(s_h^i, a_h^i), \boldsymbol{\theta} \rangle - V_{i,h+1}(s_{h+1}^i)] / \bar{\sigma}_{i,h}^2,$$

where $\bar{\sigma}_{i,h}^2$ is the upper confidence bound of the variance $[\mathbb{V}_h V_{i,h+1}(s_h^i, a_h^i)]$. $\widehat{\Sigma}_{k,h}$ is the weighted ‘‘covariance’’ matrix of $\boldsymbol{\phi}_{V_{i,h+1}}(s_h^i, a_h^i)$ weighted by $1/\bar{\sigma}_{i,h}^2$. The online update rules for $\widehat{\boldsymbol{\theta}}_{k,h}$ and $\widehat{\Sigma}_{k,h}$ are shown in Lines 18 and 22.

Next we show how to construct the variance upper bounds $\bar{\sigma}_{i,h}^2$. Due to (3.3), it suffices to estimate $\langle \boldsymbol{\phi}_{V_{k,h+1}}(s_h^k, a_h^k), \boldsymbol{\theta}_h \rangle$ and $\langle \boldsymbol{\phi}_{V_{k,h+1}^2}(s_h^k, a_h^k), \boldsymbol{\theta}_h \rangle$. For the first one, we use $\langle \boldsymbol{\phi}_{V_{k,h+1}}(s_h^k, a_h^k), \widehat{\boldsymbol{\theta}}_{k,h} \rangle$ to estimate it. For the second one, we use $\langle \boldsymbol{\phi}_{V_{k,h+1}^2}(s_h^k, a_h^k), \widetilde{\boldsymbol{\theta}}_{k,h} \rangle$, where $\widetilde{\boldsymbol{\theta}}_{k,h}$ is the linear regression estimator with contexts $\boldsymbol{\phi}_{V_{i,h+1}^2}(s_h^i, a_h^i)$ and targets $V_{i,h+1}^2(s_{h+1}^i)$. Its update rule is shown in Line 22. Notice that the stochas-

tic noise in the linear regression only comes from the stochastic transition probability \mathbb{P}_h rather than the adversarial reward. Since \mathbb{P}_h is fixed across different episodes, we can bound the estimation error and specify the choice of $\bar{\sigma}_{k,h}^2$ in the following lemma.

Lemma 4.1. We define the the estimated variance $[\bar{\mathbb{V}}_{k,h} V_{k,h+1}](s_h^k, a_h^k)$ as

$$\begin{aligned} [\bar{\mathbb{V}}_{k,h} V_{k,h+1}](s_h^k, a_h^k) &= \left[\langle \phi_{V_{k,h+1}^2}(s_h^k, a_h^k), \tilde{\theta}_{k,h} \rangle \right]_{[0, H^2]} \\ &\quad - \left[\langle \phi_{V_{k,h+1}}(s_h^k, a_h^k), \hat{\theta}_{k,h} \rangle \right]_{[0, H]}^2, \end{aligned} \quad (4.2)$$

then with probability at least $1 - 3\delta$, for all $k \in [K], h \in [H]$, we have

$$\left| [\bar{\mathbb{V}}_h V_{k,h+1}](s_h^k, a_h^k) - [\mathbb{V}_h V_{k,h+1}](s_h^k, a_h^k) \right| \leq E_{k,h},$$

where $E_{k,h}$ is defined as

$$\begin{aligned} E_{k,h} &= \min \left\{ \tilde{\beta}_k \left\| \tilde{\Sigma}_{k,h}^{-1/2} \phi_{V_{k,h+1}^2}(s_h^k, a_h^k) \right\|_2, H^2 \right\} \\ &\quad + \min \left\{ 2H \tilde{\beta}_k \left\| \tilde{\Sigma}_{k,h}^{-1/2} \phi_{V_{k,h+1}}(s_h^k, a_h^k) \right\|_2, H^2 \right\}, \\ \tilde{\beta}_k &= 8H^2 \sqrt{d \log(1 + kH^4/(d\lambda)) \log(4k^2 H/\delta)} \\ &\quad + 4H^2 \log(4k^2 H/\delta) + \sqrt{\lambda} B, \\ \hat{\beta}_k &= 8d \sqrt{\log(1 + k/\lambda) \log(4k^2 H/\delta)} \\ &\quad + 4\sqrt{d} \log(4k^2 H/\delta) + \sqrt{\lambda} B. \end{aligned} \quad (4.3)$$

For the estimator $\hat{\theta}_{k,h}$, we have

$$\theta_h \in \mathcal{C}_{k,h} = \left\{ \theta : \left\| \hat{\Sigma}_{k,h}^{1/2} (\theta - \hat{\theta}_{k,h}) \right\| \leq \hat{\beta}_k \right\}. \quad (4.4)$$

By Lemma 4.1, we know that in order to guarantee $\bar{\sigma}_{k,h}^2$ is an upper bound of the variance, it suffices to set it as $[\bar{\mathbb{V}}_{k,h} V_{k,h+1}](s_h^k, a_h^k) + E_{k,h}$. Finally, due to the technical reason, we set $\bar{\sigma}_{k,h}^2$ as

$$\bar{\sigma}_{k,h} = \sqrt{\max \left\{ H^2/d, [\bar{\mathbb{V}}_{k,h} V_{k,h+1}](s_h^k, a_h^k) + E_{k,h} \right\}}.$$

Furthermore, according to (4.4), we have

$$\begin{aligned} Q_{k,h}(s, a) &= \left[r_h^k(s, a) + \max_{\theta \in \mathcal{C}_{k,h}} \langle \phi_{V_{k,h+1}}, \theta \rangle \right]_{[0, H-h+1]} \\ &\geq \left[r_h^k(s, a) + \langle \phi_{V_{k,h+1}}, \theta_h \rangle \right]_{[0, H-h+1]}, \end{aligned}$$

and by (3.2), it is easy to show that the optimistic action-value function $Q_{k,h}(s, a)$ and the optimistic value function $V_{k,h}(s)$ are indeed upper bounds of the true action-value function $Q_{k,h}^{\pi^k}$ and the true value function $V_{k,h}^{\pi^k}$, respectively.

4.1 Computational complexity

The computational complexity of POWERS is related to the property of the given feature mapping $\phi(s'|s, a)$ and we consider a special class of linear mixture MDPs studied by Yang and Wang (2019b); Zhou et al. (2021b,a). For this special class of linear mixture MDPs, we have

$$[\phi(s'|s, a)]_i = [\psi(s')]_i \cdot [\mu(s, a)]_i, \forall i \in [d],$$

where $\psi : \mathcal{S} \rightarrow \mathbb{R}^d$ and $\mu : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}^d$. Under this setting, for each function V , the vector $\phi_V(s, a)$ can be written as the product of $\mu(s, a)$ and $\sum_{s' \in \mathcal{S}} \psi(s') V(s')$. Furthermore, the term $\sum_{s' \in \mathcal{S}} \psi(s') V(s')$ can be estimated by Monte Carlo method and in this work, we assume an access to the oracle \mathcal{O} which can compute the term $\sum_{s' \in \mathcal{S}} \psi(s') V(s')$. We also assume the size of action space is finite ($|\mathcal{A}| < \infty$) and analyze the computational complexity of POWERS in the sequel.

Recall that POWERS can be divided into two phases: (1) policy improvement; and (2) policy evaluation. For the policy evaluation phase, in order to compute the vector $\phi_{V_{k,h+1}}(s_h^k, a_h^k)$ and the vector $\phi_{V_{k,h+1}^2}(s_h^k, a_h^k)$, POWERS needs to compute the term $\sum_{s' \in \mathcal{S}} \psi(s') \phi_{V_{k,h+1}}(s')$ and $\sum_{s' \in \mathcal{S}} \psi(s') V_{k,h+1}^2(s')$, which need two accesses to the oracle \mathcal{O} . Given the vector $\phi_{V_{k,h+1}}(s_h^k, a_h^k)$ and $\phi_{V_{k,h+1}^2}(s_h^k, a_h^k)$, the covariance matrix can be computed in $O(d^2)$ time, and the estimators $\hat{\theta}_{k+1,h}$ and $\tilde{\theta}_{k+1,h}$ can be computed in $O(d^3)$ time. Therefore, the policy evaluation phase can be computed in $O(d^3 HK)$ time with $O(HK)$ accesses to the oracle \mathcal{O} .

In the policy improvement phase, POWERS will update the policy $\pi_h^k(s|a)$ for each state-action pair (s, a) :

$$\pi_h^{k+1}(a|s) \propto \pi_h^k(a|s) \exp \left\{ \alpha Q_{k,h}(s, a) \right\},$$

which leads to an $O(|\mathcal{S}||\mathcal{A}|K)$ computation complexity. In order to make the computation more efficient for large state space or even continuous state space \mathcal{S} , an alternative approach is to calculate the policy $\pi_h^k(\cdot|s_h^k)$ directly. More specifically, in Line 8, we only need the value of policy π_h^k for state s_h^k , which can be calculated as follows

$$\begin{aligned} \pi_h^k(a|s_h^k) &\propto \exp \left\{ \alpha \sum_{i=1}^{k-1} Q_{i,h}(s_h^k, a) \right\} \\ &\propto \exp \left\{ \alpha \sum_{i=1}^{k-1} \left[r_h^i(s_h^k, a) + \langle \hat{\theta}_{i,h}, \phi_{V_{i,h+1}}(s_h^k, a) \rangle \right. \right. \\ &\quad \left. \left. + \hat{\beta}_i \left\| \hat{\Sigma}_{i,h}^{-1/2} \phi_{V_{i,h+1}}(s_h^k, a) \right\|_2 \right]_{[0, H-h+1]} \right\}. \end{aligned}$$

Therefore, given the covariance matrix $\widehat{\Sigma}_{i,h}$, the estimator $\widehat{\theta}_{i,h}$ and the term $\sum_{s' \in \mathcal{S}} \psi(s') V_{i,h+1}(s')$, the policy $\pi_h^k(\cdot | s_h^k)$ can be computed in $O(d^3 K |\mathcal{A}|)$ time complexity and it will take in total $O(d^3 H K^2 |\mathcal{A}|)$ time complexity to compute all policies $\pi_h^k(\cdot | s_h^k)$ for $k \in [K]$ and $h \in [H]$. By taking the best of both computing approaches, the time complexity for the policy improvement phase of POWERS is $O(\min(d^3 H K^2 |\mathcal{A}|, |\mathcal{S}| |\mathcal{A}| K))$.

Combining the time complexity of the two phases, the total time complexity of POWERS is $O(\min(d^3 H K^2 |\mathcal{A}|, |\mathcal{S}| |\mathcal{A}| K) + d^3 H K)$ with $O(HK)$ accesses to the oracle \mathcal{O} .

5 MAIN RESULTS

In this section, we provide the regret bound for our algorithm POWERS. Here, $T = KH$ is the number of interactions with the MDP.

Theorem 5.1. For any linear mixture MDP M , if we set the parameter $\lambda = 1/B^2$ in POWERS, then with probability at least $1 - 6\delta$, the regret of POWERS is upper bounded as follows:

$$\text{Regret}(M, K) \leq \widetilde{O}(\alpha T H^2 + \alpha^{-1} H \log |\mathcal{A}| + \sqrt{d^2 H^2 T} + \sqrt{d H^3 T} + d^2 H^3 + d^{2.5} H^{2.5}).$$

Remark 5.2. If we set the learning rate $\alpha = O(\sqrt{\log |\mathcal{A}| / (H^2 K)})$ in POWERS, then Theorem 5.1 suggests that with high probability, the regret of POWERS is upper bounded by $\widetilde{O}(\sqrt{d^2 H^2 T} + \sqrt{d H^3 T} + d^2 H^3 + d^{2.5} H^{2.5} + \sqrt{H^3 \log |\mathcal{A}| T})$. When $T \geq d^3 H^3, d \geq H, d \geq \log |\mathcal{A}|$, the regret can be simplified as $\widetilde{O}(dH\sqrt{T})$. Compared with the result of Cai et al. (2020), the POWERS improves the upper bound of regret by a factor of \sqrt{H} . In addition, compared with the UCRL-VTR+ algorithm proposed by Zhou et al. (2021a) for episodic linear mixture MDPs with deterministic rewards, our algorithm POWERS provides a robustness guarantee against adversarial rewards while achieving the same regret guarantee $\widetilde{O}(dH\sqrt{T})$.

The following theorem gives a lower bound of the regret for any algorithms for learning the adversarial linear mixture MDPs.

Theorem 5.3. Suppose $B \geq 2, d \geq 4, H \geq 3, K \geq (d-1)^2 H/2$, then for any algorithm, there exists a time-inhomogenous, episodic B -bounded adversarial linear mixture MDP M , such that the expected regret for this MDP is lower bounded by $\Omega(dH\sqrt{T})$.

Remark 5.4. Theorem 5.3 suggests that when the number of episodes K is large enough, for any algorithm, the regret of learning time-inhomogenous

episodic adversarial linear mixture MDPs is at least $\Omega(dH\sqrt{T})$. Furthermore, the lower bound of regret in Theorem 5.3 matches the upper bound in Theorem 5.1 up to logarithmic factors, which suggests that POWERS is nearly minimax optimal for learning adversarial linear mixture MDPs.

6 PROOF OVERVIEW

In this section, we provide the proof sketch of Theorem 5.1. The complete proof is deferred to the supplementary material. Our proof is based on Lemma 4.1, which says with high probability, the true parameter θ_h is contained in the confidence set $\mathcal{C}_{k,h}$. For simplicity, we denote by \mathcal{E} the event when the result of Lemma 4.1 holds and due to Lemma 4.1, we have $\Pr(\mathcal{E}) \geq 1 - 3\delta$. Recall the definition of regret, we have

$$\text{Regret}(K) = \underbrace{\sum_{k=1}^K (V_{k,1}^*(s_1^k) - V_{k,1}(s_1^k))}_{I_1} + \underbrace{\sum_{k=1}^K (V_{k,1}(s_1^k) - V_{k,1}^{\pi^k}(s_1^k))}_{I_2}.$$

Therefore, the proof can be divided into two main steps.

Step 1: Bounding the term I_1 .

Lemma 6.1. On the event \mathcal{E} , for all $k \in [K]$, we have

$$V_{k,1}^*(s_1^k) - V_{k,1}(s_1^k) \leq \mathbb{E} \left[\sum_{h=1}^H \left\{ \mathbb{E}_{a \sim \pi_h^*(\cdot | s_h)} [Q_{k,h}(s_h, a)] - \mathbb{E}_{a \sim \pi_h^k(\cdot | s_h)} [Q_{k,h}(s_h, a)] \mid s_1 = s_1^k \right\} \right].$$

Here $\mathbb{E}[\cdot | s = s_1^k]$ is the expectation with respect to the randomness of the state-action sequence $\{(s_h, a_h)\}_{h=1}^H$, where $a_h \sim \pi_h^*(\cdot | s_h)$ and $s_{h+1} \sim \mathbb{P}_h(\cdot | s_h, a_h)$.

Lemma 6.1 suggests that the regret can be decomposed as the sum of the advantages at different stages h . Furthermore, since the initial state s_1^k and the optimal policy π^* is fixed across different episode k , the state-action sequence $(s_1, a_1, \dots, s_H, a_H)$ induced by the policy π^* follows the same distribution across different episode k .

In each episode k , the POWERS update the policy π_h^k by the following rule:

$$\pi_h^{k+1}(a|s) \propto \pi_h^k(a|s) \exp \{ \alpha Q_{k,h}(s, a) \}.$$

By the above update rule, we have the following lemma.

Lemma 6.2. On the event \mathcal{E} , for all $k \in [K], h \in [H], s \in \mathcal{S}$, we have

$$\mathbb{E}_{a \sim \pi_h^*(\cdot | s)} [Q_{k,h}(s, a)] - \mathbb{E}_{a \sim \pi_h^k(\cdot | s)} [Q_{k,h}(s, a)]$$

$$\leq \frac{\alpha H^2}{2} + \alpha^{-1} \left(D_{KL}(\pi_h^*(\cdot|s) \|\pi_h^k(\cdot|s)) - D_{KL}(\pi_h^*(\cdot|s) \|\pi_h^{k+1}(\cdot|s)) \right).$$

Substituting the result of Lemma 6.2 into the result of Lemma 4.1, we have the following upper bound for the term I_1 :

$$\begin{aligned} I_1 &\leq \frac{\alpha K H^3}{2} + \sum_{k=1}^K \alpha^{-1} \mathbb{E} \left[\sum_{h=1}^H \left\{ D_{KL}(\pi_h^*(\cdot|s_h) \|\pi_h^k(\cdot|s_h)) - KL(\pi_h^*(\cdot|s_h) \|\pi_h^{k+1}(\cdot|s_h)) \right\} \right] \\ &= \frac{\alpha K H^3}{2} + \alpha^{-1} \mathbb{E} \left[\sum_{h=1}^H \left\{ D_{KL}(\pi_h^*(\cdot|s_h) \|\pi_h^1(\cdot|s_h)) - D_{KL}(\pi_h^*(\cdot|s_h) \|\pi_h^{K+1}(\cdot|s_h)) \right\} \right] \\ &\leq \frac{\alpha K H^3}{2} + \alpha^{-1} \mathbb{E} \left[\sum_{h=1}^H D_{KL}(\pi_h^*(\cdot|s_h) \|\pi_h^1(\cdot|s_h)) \right] \\ &\leq \frac{\alpha K H^3}{2} + \alpha^{-1} H \log |\mathcal{A}|, \end{aligned}$$

where s_1 is the fixed initial state, $a_h \sim \pi_h^*(\cdot|s_h)$, $s_{h+1} \sim \mathbb{P}_h(\cdot|s_h, a_h)$, the second inequality holds due to Kullback–Leibler divergence is non-negative and the third inequality holds due to the fact that initial policy π_h^1 is uniform over the action space \mathcal{A} .

Step 2: Bounding the term I_2 .

Lemma 6.3. On the event \mathcal{E} , for all $k \in [K]$, $h \in [H]$, we have

$$\begin{aligned} &Q_{k,h}(s_h^k, a_h^k) - Q_{k,h}^{\pi^k}(s_h^k, a_h^k) \\ &\leq [\mathbb{P}_h(V_{k,h+1} - V_{k,h+1}^{\pi^k})](s_h^k, a_h^k) \\ &\quad + 2\hat{\beta}_k \bar{\sigma}_{k,h} \min \left\{ \left\| \hat{\Sigma}_{k,h}^{-1/2} \phi_{V_{k,h+1}}(s_h^k, a_h^k) / \bar{\sigma}_{k,h} \right\|_2, 1 \right\}. \end{aligned}$$

Lemma 6.3 suggests that the difference between the true action-value function $Q_{k,h}^{\pi^k}(s_h^k, a_h^k)$ and the estimated value function $Q_{k,h}(s_h^k, a_h^k)$ can be bounded by the expected difference at the next-stage difference. Furthermore, for the expected difference at the next-stage and the exact difference at the next-stage, we have the following equation

$$\begin{aligned} &[\mathbb{P}_h(V_{k,h+1} - V_{k,h+1}^{\pi^k})](s_h^k, a_h^k) \\ &= Q_{k,h+1}(s_{h+1}^k, a_{h+1}^k) - Q_{k,h+1}^{\pi^k}(s_{h+1}^k, a_{h+1}^k) \\ &\quad + A_{h,k} + B_{h+1,k}, \end{aligned}$$

where $A_{h,k} = [\mathbb{P}_h(V_{k,h+1} - V_{k,h+1}^{\pi^k})](s_h^k, a_h^k) - (V_{k,h+1}(s_{h+1}^k) - V_{k,h+1}^{\pi^k}(s_{h+1}^k))$ is the noise from the state transition and $B_{h,k} = \mathbb{E}_{a \sim \pi_h^k(\cdot|s_h^k)} [Q_{k,h}(s_h^k, a) -$

$Q_{k,h}^{\pi^k}(s_h^k, a)] - (Q_{k,h}(s_h^k, a_h^k) - Q_{k,h}^{\pi^k}(s_h^k, a_h^k))$ is the noise from the stochastic policy. These noises form a martingale difference sequence and we define two high probability events for them:

$$\begin{aligned} \mathcal{E}_1 &= \left\{ \forall h \in [H], \sum_{k=1}^K \sum_{h'=h}^H A_{h',k} + B_{h',k} \leq 4H \sqrt{T \log(H/\delta)} \right\}, \\ \mathcal{E}_2 &= \left\{ \sum_{k=1}^K \sum_{h=1}^H A_{h,k} \leq 2H \sqrt{2T \log(1/\delta)} \right\}. \end{aligned}$$

Then according to the Azuma–Hoeffding inequality, we have $\Pr(\mathcal{E}_1) \geq 1 - \delta$ and $\Pr(\mathcal{E}_2) \geq 1 - \delta$. Furthermore, on the events \mathcal{E}_1 and \mathcal{E}_2 , we can telescope the inequality in Lemma 6.3 over the K episodes, and obtain the following lemma.

Lemma 6.4. On the event $\mathcal{E} \cap \mathcal{E}_1 \cap \mathcal{E}_2$, for all $h \in [H]$, we have

$$\begin{aligned} &\sum_{k=1}^K (V_{k,h}(s_h^k) - V_{k,h}^{\pi^k}(s_h^k)) \\ &\leq 2\hat{\beta}_K \sqrt{\sum_{k=1}^K \sum_{h=1}^H \bar{\sigma}_{k,h}^2} \sqrt{2Hd \log(1 + K/\lambda)} \\ &\quad + 4H \sqrt{T \log(H/\delta)}, \end{aligned}$$

and

$$\begin{aligned} &\sum_{k=1}^K \sum_{h=1}^H [\mathbb{P}_h(V_{k,h+1} - V_{k,h+1}^{\pi^k})](s_h^k, a_h^k) \\ &\leq 2H \hat{\beta}_K \sqrt{\sum_{k=1}^K \sum_{h=1}^H \bar{\sigma}_{k,h}^2} \sqrt{2Hd \log(1 + K/\lambda)} \\ &\quad + 4H^2 \sqrt{T \log(H/\delta)}. \end{aligned}$$

Lemma 6.4 shows that the regret can be upper bounded by the total estimated variance $\sum_{k=1}^K \sum_{h=1}^H \bar{\sigma}_{k,h}^2$. For the total variance $\sum_{k=1}^K \sum_{h=1}^H [\mathbb{V}_h V_{k,h+1}^{\pi^k}](s_h^k, a_h^k)$, we introduce the high probability event \mathcal{E}_3 :

$$\mathcal{E}_3 = \left\{ \sum_{k=1}^K \sum_{h=1}^H [\mathbb{V}_h V_{k,h+1}^{\pi^k}](s_h^k, a_h^k) \leq 3HT + 3H^3 \log(1/\delta) \right\}.$$

Since the stochastic noise in the linear regression only comes from the stochastic transition probability \mathbb{P}_h rather than the adversarial reward, and \mathbb{P}_h is fixed across different episodes, Lemma C.5 in Jin et al. (2018) suggests that $\Pr(\mathcal{E}_3) \geq 1 - \delta$ and on the event $\mathcal{E} \cap \mathcal{E}_1 \cap \mathcal{E}_2 \cap \mathcal{E}_3$, the following lemma gives an upper bound of the total estimated variance.

Lemma 6.5. On the event $\mathcal{E} \cap \mathcal{E}_1 \cap \mathcal{E}_2 \cap \mathcal{E}_3$, we have

$$\begin{aligned} \sum_{k=1}^K \sum_{h=1}^H \bar{\sigma}_{k,h}^2 &\leq 2HT/d + 179HT \\ &+ 165d^3 H^4 \log^2(4K^2 H/\delta) \log^2(1 + KH^4/\lambda) \\ &+ 2062H^5 d^2 \log^2(4K^2 H/\delta) \log^2(1 + K/\lambda). \end{aligned}$$

Combining the results of Lemmas 6.4 and 6.5, the term I_2 can be upper bounded by

$$I_2 = \tilde{O}(\sqrt{dH^3 T} + \sqrt{d^2 H^2 T} + d^2 H^3 + d^{2.5} H^{2.5}).$$

Finally, combining the upper bounds of I_1 and I_2 , we have

$$\begin{aligned} \text{Regret}(K) &= I_1 + I_2 \\ &= \tilde{O}(\sqrt{dH^3 T} + \sqrt{d^2 H^2 T} + d^2 H^3 + d^{2.5} H^{2.5}) \\ &\quad + \alpha H^2 + \alpha^{-1} H \log |\mathcal{A}|. \end{aligned}$$

7 EXPERIMENTS

In this section, we carry out experiments to evaluate the empirical performance of our algorithm POWERS.

In this experiment, we construct the MDP M with dimension $d = 5$ and episode length $H = 20$. In this MDP, the state space \mathcal{S} consists of $H + 2$ different states s_1, \dots, s_{H+2} and the action space $\mathcal{A} = \{-1, 1\}^{d-1}$ consists of 2^{d-1} different actions. For each stage $h \in [H]$ and episode $k \in [K]$, the adversarial reward function r_h^k satisfies that $r_h^k(s_h, \mathbf{a}) = 0$, ($1 \leq h \leq H + 1$) and $r_h^k(s_{H+2}, \mathbf{a}) = 1$. For each stage $h \in [H]$ and corresponding transition probability function \mathbb{P}_h , s_{H+1} and s_{H+2} are absorbing states. For other states $s_h (1 \leq h \leq H)$, the transition probability satisfies that

$$\begin{aligned} \mathbb{P}_h(s_{h+1} | s_h, \mathbf{a}) &= 0.95 - \langle 0.01 \cdot \mathbf{1}_{d-1}, \mathbf{a} \rangle, \\ \mathbb{P}_h(s_{H+2} | s_h, \mathbf{a}) &= 0.05 + \langle 0.01 \cdot \mathbf{1}_{d-1}, \mathbf{a} \rangle, \end{aligned}$$

where each $\mathbf{1}_{d-1}$ is a $(d - 1)$ -dimensional vector of all ones. In the experiment, we set the regularization parameter $\lambda = 1$, learning rate $\alpha = 1$ and use grid search to select the parameters $\hat{\beta}$, $\tilde{\beta}$, $\bar{\beta}$ which obtain the best performance. We compare our algorithm with two baseline benchmarks : **Random** (uniformly and randomly choose actions from \mathcal{A}) and **OPPO** (Cai et al., 2020). The regrets of different algorithms for the first 1000 episodes averaged over 20 runs are plotted in Figure 1.

From Figure 1, it can be seen that both OPPO (Cai et al., 2020) and POWERS obtain a sub-linear cumulative regret. Moreover, we can see that POWERS outperforms the OPPO algorithm. This is consistent with our theoretical result in Theorem 5.1, which suggests that POWERS has a better regret bound than OPPO thanks to the use of Bernstein-type bonus.




Figure 1: Cumulative regret comparison in the first 1000 episodes of different algorithms. Results are averaged over 20 runs.

8 CONCLUSION AND FUTURE WORK

In this work, we considered learning adversarial Markov decision processes under the linear mixture MDP assumption. We proposed a novel algorithm POWERS and proved that with high probability, the regret of POWERS is upper bounded by $\tilde{O}(dH\sqrt{T})$, which matches the lower bound up to logarithmic factors. Currently, our work requires the full information feedback of the reward and it remains an open problem that if there exists a provably efficient algorithm for learning adversarial linear mixture MDPs with bandit-feedback on the reward. We leave it as future work.

Acknowledgments

We thank the anonymous reviewers for their helpful comments. Part of this work was done when JH, DZ and QG participated the Theory of Reinforcement Learning program at the Simons Institute for the Theory of Computing in Fall 2020. JH, DZ and QG are partially supported by the National Science Foundation CAREER Award 1906169, IIS-1904183, and AWS Machine Learning Research Award. The views and conclusions contained in this paper are those of the authors and should not be interpreted as representing any funding agencies.

References

- ABBASI-YADKORI, Y., BARTLETT, P., BHATIA, K., LAZIC, N., SZEPESVARI, C. and WEISZ, G. (2019a). Politex: Regret bounds for policy iteration using expert prediction. In *International Conference on Machine Learning*. PMLR.

- ABBASI-YADKORI, Y., LAZIC, N., SZEPESVARI, C. and WEISZ, G. (2019b). Exploration-enhanced politex. *arXiv preprint arXiv:1908.10479* .
- ABBASI-YADKORI, Y., PÁL, D. and SZEPESVÁRI, C. (2011). Improved algorithms for linear stochastic bandits. In *Advances in Neural Information Processing Systems*.
- AYOUB, A., JIA, Z., SZEPESVARI, C., WANG, M. and YANG, L. (2020). Model-based reinforcement learning with value-targeted regression. In *International Conference on Machine Learning*. PMLR.
- BAGNELL, J. A. and SCHNEIDER, J. (2003). Covariant policy search .
- BAXTER, J. and BARTLETT, P. L. (2000). Direct gradient-based reinforcement learning. In *2000 IEEE International Symposium on Circuits and Systems (ISCAS)*, vol. 3. IEEE.
- CAI, Q., YANG, Z., JIN, C. and WANG, Z. (2020). Provably efficient exploration in policy optimization. In *International Conference on Machine Learning*. PMLR.
- CESA-BIANCHI, N. and LUGOSI, G. (2006). *Prediction, learning, and games*. Cambridge university press.
- DEKEL, O. and HAZAN, E. (2013). Better rates for any adversarial deterministic mdp. In *International Conference on Machine Learning*. PMLR.
- DU, S., KRISHNAMURTHY, A., JIANG, N., AGARWAL, A., DUDIK, M. and LANGFORD, J. (2019). Provably efficient rl with rich observations via latent state decoding. In *International Conference on Machine Learning*. PMLR.
- EFRONI, Y., SHANI, L., ROSENBERG, A. and MANNOR, S. (2020). Optimistic policy optimization with bandit feedback. *arXiv preprint arXiv:2002.08243* .
- EVEN-DAR, E., KAKADE, S. M. and MANSOUR, Y. (2009). Online markov decision processes. *Mathematics of Operations Research* **34** 726–736.
- GERGELY NEU, A. G., SZEPESVÁRI, C. and ANTOS, A. (2010). Online markov decision processes under bandit feedback. In *Proceedings of the Twenty-Fourth Annual Conference on Neural Information Processing Systems*.
- HAO, B., LAZIC, N., ABBASI-YADKORI, Y., JOULANI, P. and SZEPESVARI, C. (2020). Provably efficient adaptive approximate policy iteration. *arXiv preprint arXiv:2002.03069* .
- HE, J., ZHOU, D. and GU, Q. (2021). Logarithmic regret for reinforcement learning with linear function approximation. In *International Conference on Machine Learning*. PMLR.
- JIA, Z., YANG, L., SZEPESVARI, C. and WANG, M. (2020). Model-based reinforcement learning with value-targeted regression. In *Learning for Dynamics and Control*. PMLR.
- JIANG, N., KRISHNAMURTHY, A., AGARWAL, A., LANGFORD, J. and SCHAPIRE, R. E. (2017). Contextual decision processes with low bellman rank are pac-learnable. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*. JMLR. org.
- JIN, C., ALLEN-ZHU, Z., BUBECK, S. and JORDAN, M. I. (2018). Is q-learning provably efficient? In *Advances in Neural Information Processing Systems*.
- JIN, C., JIN, T., LUO, H., SRA, S. and YU, T. (2020a). Learning adversarial markov decision processes with bandit feedback and unknown transition. In *International Conference on Machine Learning*. PMLR.
- JIN, C., YANG, Z., WANG, Z. and JORDAN, M. I. (2020b). Provably efficient reinforcement learning with linear function approximation. In *Conference on Learning Theory*.
- KAKADE, S. and LANGFORD, J. (2002). Approximately optimal approximate reinforcement learning. In *In Proc. 19th International Conference on Machine Learning*. Citeseer.
- KAKADE, S. M. (2001). A natural policy gradient. *Advances in neural information processing systems* **14**.
- KAKADE, S. M. ET AL. (2003). *On the sample complexity of reinforcement learning*. Ph.D. thesis.
- KIRSCHNER, J. and KRAUSE, A. (2018). Information directed sampling and bandits with heteroscedastic noise. In *Conference On Learning Theory*. PMLR.
- MODI, A., JIANG, N., TEWARI, A. and SINGH, S. (2019). Sample complexity of reinforcement learning using linearly combined model ensembles. In *International Conference on Artificial Intelligence and Statistics*.
- NEU, G., GYÖRGY, A. and SZEPESVÁRI, C. (2010). The online loop-free stochastic shortest-path problem. In *COLT*, vol. 2010. Citeseer.
- NEU, G., GYORGY, A. and SZEPESVÁRI, C. (2012). The adversarial stochastic shortest path problem with unknown transition probabilities. In *Artificial Intelligence and Statistics*. PMLR.
- NEU, G., JONSSON, A. and GÓMEZ, V. (2017). A unified view of entropy-regularized markov decision processes. *arXiv preprint arXiv:1705.07798* .
- PUTERMAN, M. L. (1994). Markov decision processes: Discrete stochastic dynamic programming.

- ROSENBERG, A. and MANSOUR, Y. (2019a). Online convex optimization in adversarial markov decision processes. In *International Conference on Machine Learning*. PMLR.
- ROSENBERG, A. and MANSOUR, Y. (2019b). Online stochastic shortest path with bandit feedback and unknown transition function. *Advances in Neural Information Processing Systems* **32** 2212–2221.
- SCHULMAN, J., LEVINE, S., ABBEEL, P., JORDAN, M. and MORITZ, P. (2015). Trust region policy optimization. In *International conference on machine learning*. PMLR.
- SCHULMAN, J., WOLSKI, F., DHARIWAL, P., RADFORD, A. and KLIMOV, O. (2017). Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*.
- SUN, W., JIANG, N., KRISHNAMURTHY, A., AGARWAL, A. and LANGFORD, J. (2019). Model-based rl in contextual decision processes: Pac bounds and exponential improvements over model-free approaches. In *Conference on Learning Theory*. PMLR.
- SUTTON, R. S., MCALLESTER, D. A., SINGH, S. P., MANSOUR, Y. ET AL. (1999). Policy gradient methods for reinforcement learning with function approximation. In *NIPs*, vol. 99. Citeseer.
- TAMAR, A., MANNOR, S. and XU, H. (2014). Scaling up robust mdps using function approximation. In *International conference on machine learning*. PMLR.
- WANG, Y., WANG, R., DU, S. S. and KRISHNAMURTHY, A. (2019). Optimism in reinforcement learning with generalized linear function approximation. *arXiv preprint arXiv:1912.04136*.
- WANG, Y. and ZOU, S. (2021). Online robust reinforcement learning with model uncertainty. *Advances in Neural Information Processing Systems* **34**.
- WILLIAMS, R. J. (1992). Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning* **8** 229–256.
- WU, Y., ZHOU, D. and GU, Q. (2022). Nearly minimax optimal regret for learning infinite-horizon average-reward mdps with linear function approximation. In *International Conference on Artificial Intelligence and Statistics*.
- YANG, L. and WANG, M. (2019a). Sample-optimal parametric q-learning using linearly additive features. In *International Conference on Machine Learning*.
- YANG, L. F. and WANG, M. (2019b). Reinforcement learning in feature space: Matrix bandit, kernels, and regret bound. In *International Conference on Machine Learning*.
- YU, J. Y., MANNOR, S. and SHIMKIN, N. (2009). Markov decision processes with arbitrary reward processes. *Mathematics of Operations Research* **34** 737–757.
- ZANETTE, A., BRANDFONBRENER, D., BRUNSKILL, E., PIROTTA, M. and LAZARIC, A. (2020a). Frequentist regret bounds for randomized least-squares value iteration. In *International Conference on Artificial Intelligence and Statistics*.
- ZANETTE, A., LAZARIC, A., KOCHENDERFER, M. and BRUNSKILL, E. (2020b). Learning near optimal policies with low inherent bellman error. In *International Conference on Machine Learning*.
- ZHANG, K., HU, B. and BASAR, T. (2020). On the stability and convergence of robust adversarial reinforcement learning: A case study on linear quadratic systems. *Advances in Neural Information Processing Systems* **33** 22056–22068.
- ZHOU, D., GU, Q. and SZEPESVARI, C. (2021a). Nearly minimax optimal reinforcement learning for linear mixture markov decision processes. In *Conference on Learning Theory*. PMLR.
- ZHOU, D., HE, J. and GU, Q. (2021b). Provably efficient reinforcement learning for discounted mdps with feature mapping. In *International Conference on Machine Learning*. PMLR.
- ZIMIN, A. and NEU, G. (2013). Online learning in episodic markovian decision processes by relative entropy policy search. In *Neural Information Processing Systems 26*.

Supplementary Material: Near-optimal Policy Optimization Algorithms for Learning Adversarial Linear Mixture MDPs

A PROOF OF THE MAIN RESULTS

A.1 Proof of Theorem 5.1

In this section, we provide the proof of Theorems 5.1 and first propose the following lemmas.

Lemma A.1 (Azuma–Hoeffding inequality, Cesa-Bianchi and Lugosi 2006). Let $\{x_i\}_{i=1}^n$ be a martingale difference sequence with respect to a filtration $\{\mathcal{G}_i\}$ satisfying $|x_i| \leq M$ for some constant M , x_i is \mathcal{G}_{i+1} -measurable, $\mathbb{E}[x_i|\mathcal{G}_i] = 0$. Then for any $0 < \delta < 1$, with probability at least $1 - \delta$, we have

$$\sum_{i=1}^n x_i \leq M\sqrt{2n \log(1/\delta)}.$$

Lemma A.2. On the event \mathcal{E} , for all $k \in [K], h \in [H], s \in \mathcal{S}, a \in \mathcal{A}$, we have

$$Q_{k,h}(s, a) \geq r_h^k(s, a) + [\mathbb{P}_h V_{k,h+1}](s, a).$$

Furthermore, on the event \mathcal{E} , for all $k \in [K], h \in [H], s \in \mathcal{S}, a \in \mathcal{A}$, we have

$$Q_{k,h}(s, a) - Q_{k,h}^{\pi^k}(s, a) \geq 0, V_{k,h}(s) - V_{k,h}^{\pi^k}(s) \geq 0.$$

Proof of Lemma A.2. For the lower bound of $Q_{k,h}(s, a) - Q_{k,h}^{\pi^k}(s, a)$, we have

$$\begin{aligned} & r_h^k(s, a) + \langle \widehat{\boldsymbol{\theta}}_{k,h}, \boldsymbol{\phi}_{V_{k,h+1}}(s, a) \rangle + \widehat{\beta}_k \|\widehat{\boldsymbol{\Sigma}}_{k,h}^{-1/2} \boldsymbol{\phi}_{V_{k,h+1}}(s, a)\|_2 \\ &= r_h^k(s, a) + [\mathbb{P}_h V_{k,h+1}](s, a) + \langle \widehat{\boldsymbol{\theta}}_{k,h} - \boldsymbol{\theta}_h, \boldsymbol{\phi}_{V_{k,h+1}}(s, a) \rangle + \widehat{\beta}_k \|\widehat{\boldsymbol{\Sigma}}_{k,h}^{-1/2} \boldsymbol{\phi}_{V_{k,h+1}}(s, a)\|_2 \\ &\geq r_h^k(s, a) + [\mathbb{P}_h V_{k,h+1}](s, a) + \widehat{\beta}_k \|\widehat{\boldsymbol{\Sigma}}_{k,h}^{-1/2} \boldsymbol{\phi}_{V_{k,h+1}}(s, a)\|_2 - \|\widehat{\boldsymbol{\Sigma}}_{k,h}^{1/2}(\boldsymbol{\theta} - \widehat{\boldsymbol{\theta}}_{k,h})\|_2 \|\widehat{\boldsymbol{\Sigma}}_{k,h}^{-1/2} \boldsymbol{\phi}_{V_{k,h+1}}(s, a)\|_2 \\ &\geq r_h^k(s, a) + [\mathbb{P}_h V_{k,h+1}](s, a), \end{aligned} \tag{A.1}$$

where the first inequality holds due to Cauchy-Schwarz inequality and the second inequality holds due to the definition of event \mathcal{E} . Therefore, we have

$$\begin{aligned} Q_{k,h}(s, a) &= \left[r_h^k(s, a) + \langle \widehat{\boldsymbol{\theta}}_{k,h}, \boldsymbol{\phi}_{V_{k,h+1}}(s, a) \rangle + \widehat{\beta}_k \|\widehat{\boldsymbol{\Sigma}}_{k,h}^{-1/2} \boldsymbol{\phi}_{V_{k,h+1}}(s, a)\|_2 \right]_{[0, H-h+1]} \\ &\geq \min \left\{ r_h^k(s, a) + [\mathbb{P}_h V_{k,h+1}](s, a), H - h + 1 \right\} \\ &\geq r_h^k(s, a) + [\mathbb{P}_h V_{k,h+1}](s, a), \end{aligned} \tag{A.2}$$

where the first inequality holds due to (A.1) and the second inequality holds due to $r_h^k(s, a) + [\mathbb{P}_h V_{k,h+1}](s, a) \leq 1 + (H - h) = H - h + 1$. Now, we prove the second part of Lemma A.2 by induction. The statement holds for stage $h = H + 1$, since

$$V_{k,h}(s) = V_{k,h}^{\pi^k}(s) = 0.$$

When the second part of Lemma A.2 holds for stage $h + 1$, we have

$$Q_{k,h}(s, a) \geq r_h^k(s, a) + [\mathbb{P}_h V_{k,h+1}](s, a)$$

$$\begin{aligned}
 &\geq r_h^k(s, a) + [\mathbb{P}_h V_{k,h+1}^{\pi^k}](s, a) \\
 &= Q^{\pi^k}(s, a),
 \end{aligned} \tag{A.3}$$

where the first inequality holds due to (A.2) and the second inequality holds due to the induction assumption. Furthermore, we have

$$V_{k,h}(s) = \mathbb{E}_{a \sim \pi_h^k(\cdot|s)}[Q_{k,h}(s, a)] \geq \mathbb{E}_{a \sim \pi_h^k(\cdot|s)}[Q_{k,h}^{\pi^k}(s, a)] = V_{k,h}^{\pi^k}(s).$$

Therefore, we finish the induction step and complete the proof of Lemma A.2. \square

Lemma A.3. [Lemma 11, Abbasi-Yadkori et al. 2011] Let $\{\mathbf{x}_t\}_{t=1}^{\infty}$ be a sequence in \mathbb{R}^d , $\mathbf{V}_0 = \lambda \mathbf{I}$ and define $\mathbf{V}_t = \mathbf{V}_0 + \sum_{i=1}^t \mathbf{x}_i \mathbf{x}_i^\top$. If $\|\mathbf{x}_i\|_2 \leq L$ holds for each i , then for each t , we have

$$\sum_{i=1}^t \min\{1, \|\mathbf{x}_i\|_{\mathbf{V}_{i-1}^{-1}}\} \leq 2d \log\left(\frac{d\lambda + tL^2}{d\lambda}\right).$$

Proof of Theorem 5.1. For the regret, we have

$$\begin{aligned}
 \text{Regret}(K) &= \sup_{\pi} \sum_{k=1}^K (V_{k,1}^{\pi}(s_1^k) - V_{k,1}^{\pi^k}(s_1^k)) \\
 &= \sum_{k=1}^K (V_{k,1}^*(s_1^k) - V_{k,1}^{\pi^k}(s_1^k)) \\
 &= \underbrace{\sum_{k=1}^K (V_{k,1}^*(s_1^k) - V_{k,1}(s_1^k))}_{I_1} + \underbrace{\sum_{k=1}^K (V_{k,1}(s_1^k) - V_{k,1}^{\pi^k}(s_1^k))}_{I_2}.
 \end{aligned}$$

For the term I_1 , applying Lemma 6.1, we have

$$\begin{aligned}
 I_1 &= \sum_{k=1}^K (V_{k,1}^*(s_1^k) - V_{k,1}(s_1^k)) \\
 &\leq \sum_{k=1}^K \mathbb{E} \left[\sum_{h=1}^H \left\{ \mathbb{E}_{a \sim \pi_h^*(\cdot|s_h)}[Q_{k,h}(s_h, a)] - \mathbb{E}_{a \sim \pi_h^k(\cdot|s_h)}[Q_{k,h}(s_h, a)] \right\} \middle| s_1 = s_1^k \right] \\
 &\leq \sum_{k=1}^K \mathbb{E} \left[\sum_{h=1}^H \left\{ \frac{\alpha H^2}{2} + \alpha^{-1} \left(D_{KL}(\pi_h^*(\cdot|s_h) \|\pi_h^k(\cdot|s_h)) - D_{KL}(\pi_h^*(\cdot|s_h) \|\pi_h^{k+1}(\cdot|s_h)) \right) \right\} \middle| s_1 = s_1^k \right] \\
 &= \frac{\alpha K H^3}{2} + \sum_{k=1}^K \alpha^{-1} \mathbb{E} \left[\sum_{h=1}^H \left\{ D_{KL}(\pi_h^*(\cdot|s_h) \|\pi_h^k(\cdot|s_h)) - D_{KL}(\pi_h^*(\cdot|s_h) \|\pi_h^{k+1}(\cdot|s_h)) \right\} \middle| s_1 = s_1^k \right], \tag{A.4}
 \end{aligned}$$

where $a_h \sim \pi_h^*(\cdot|s_h)$, $s_{h+1} \sim \mathbb{P}_h(\cdot|s_h, a_h)$, the first inequality holds due to Lemma 6.1 and the second inequality holds due to Lemma 6.2. For Kullback–Leibler divergence $D_{KL}(\pi_h^*(\cdot|s_h) \|\pi_h^1(\cdot|s_h))$, we have $0 \leq D_{KL}(\pi_h^*(\cdot|s_h) \|\pi_h^1(\cdot|s_h))$ and

$$\begin{aligned}
 D_{KL}(\pi_h^*(\cdot|s_h) \|\pi_h^1(\cdot|s_h)) &= \sum_{a \in \mathcal{A}} \pi_h^*(a|s_h) \log \left(\frac{\pi_h^*(a|s_h)}{\pi_h^1(a|s_h)} \right) \\
 &= \sum_{a \in \mathcal{A}} \pi_h^*(a|s_h) \log (\pi_h^*(a|s_h) \times |\mathcal{A}|) \\
 &= \log |\mathcal{A}| + \sum_{a \in \mathcal{A}} \pi_h^*(a|s_h) \log (\pi_h^*(a|s_h)) \\
 &\leq \log |\mathcal{A}|,
 \end{aligned} \tag{A.5}$$

where the first equation holds due to $\pi_h^1(a|s_h) = 1/|\mathcal{A}|$ and the inequality holds on due to $0 \leq \pi_h^*(a|s_h) \leq 1$. Substituting (A.5) into (A.4), we have

$$\begin{aligned}
 I_1 &\leq \frac{\alpha KH^3}{2} + \sum_{k=1}^K \alpha^{-1} \mathbb{E} \left[\sum_{h=1}^H \left\{ D_{KL}(\pi_h^*(\cdot|s_h) \|\pi_h^k(\cdot|s_h)) - D_{KL}(\pi_h^*(\cdot|s_h) \|\pi_h^{k+1}(\cdot|s_h)) \right\} \right] \\
 &= \frac{\alpha KH^3}{2} + \alpha^{-1} \mathbb{E} \left[\sum_{h=1}^H \left\{ D_{KL}(\pi_h^*(\cdot|s_h) \|\pi_h^1(\cdot|s_h)) - D_{KL}(\pi_h^*(\cdot|s_h) \|\pi_h^{K+1}(\cdot|s_h)) \right\} \right] \\
 &\leq \frac{\alpha KH^3}{2} + \alpha^{-1} \mathbb{E} \left[\sum_{h=1}^H \left\{ D_{KL}(\pi_h^*(\cdot|s_h) \|\pi_h^1(\cdot|s_h)) \right\} \right] \\
 &\leq \frac{\alpha KH^3}{2} + \alpha^{-1} H \log |\mathcal{A}|, \tag{A.6}
 \end{aligned}$$

where s_1 is the fixed initial state, $a_h \sim \pi_h^*(\cdot|s_h)$, $s_{h+1} \sim \mathbb{P}_h(\cdot|s_h, a_h)$, the first inequality holds due to (A.4) and the second inequality holds due to Kullback–Leibler divergence is non-negative and the third inequality holds due to (A.5). For the term I_2 , we have

$$\begin{aligned}
 I_2 &= \sum_{k=1}^K (V_{k,1}(s_1^k) - V_{k,1}^{\pi^k}(s_1^k)) \\
 &\leq 2\hat{\beta}_K \sqrt{\sum_{k=1}^K \sum_{h=1}^H \bar{\sigma}_{k,h}^2 \sqrt{2Hd \log(1+K/\lambda)} + 4H\sqrt{T \log(H/\delta)}} \\
 &\leq 56\sqrt{dH^3 T} \log(4K^2 H/\delta) \log(1+K/\lambda) + 492\sqrt{d^2 H^2 T} \log(4K^2 H/\delta) \log(1+K/\lambda) \\
 &\quad + 1670d^2 H^3 \log^2(4K^2 H/\delta) \log^2(1+K/\lambda) + 473d^{2.5} H^{2.5} \log^2(4K^2 H/\delta) \log^2(1+KH^4/\lambda), \tag{A.7}
 \end{aligned}$$

where the first inequality holds due to Lemma 6.4 and the second inequality holds due to Lemma 6.5 with the fact that $\sqrt{a+b+c+d} \leq \sqrt{a} + \sqrt{b} + \sqrt{c} + \sqrt{d}$. Substituting (A.6) and (A.7) into (A.5), we finish the proof of Theorem 5.1. \square

A.2 Proof of Theorem 5.3

In this section, we provide the proof of the lower bounds of the regret and the lower bound is based on previous work (Zhou et al., 2021b,a).

Proof of Theorem 5.3. To prove the lower bound, we construct a series of hard-to-learn adversarial MDPs introduced by Zhou et al. (2021b,a). To be more specific, the state space \mathcal{S} consists of states s_1, \dots, s_{H+2} , where s_{H+1} and s_{H+2} are absorbing states. The action space $\mathcal{A} = \{-1, 1\}^{d-1}$ consists of 2^{d-1} different actions. The adversarial reward function r_h^k satisfies that $r_h^k(s_h, \mathbf{a}) = 0$ ($1 \leq h \leq H+1$) and $r_h^k(s_{H+2}, \mathbf{a}) = 1$. For the transition probability function \mathbb{P}_h , s_{H+1} and s_{H+2} are absorbing states, which will always stay at the same state, and for other state s_h ($1 \leq h \leq H$), we have

$$\begin{aligned}
 \mathbb{P}_h(s_{h+1}|s_h, \mathbf{a}) &= 1 - \delta - \langle \boldsymbol{\mu}_h, \mathbf{a} \rangle, \\
 \mathbb{P}_h(s_{H+2}|s_h, \mathbf{a}) &= \delta + \langle \boldsymbol{\mu}_h, \mathbf{a} \rangle,
 \end{aligned}$$

where each $\boldsymbol{\mu}_h \in \{-\Delta, \Delta\}^d$ with $\Delta = \sqrt{\delta/K}/(4\sqrt{2})$ and $\delta = 1/H$. Furthermore, these hard-to-learn adversarial MDPs can be represented as linear mixture MDPs with the following feature mapping $\phi: \mathcal{S} \times \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}^d$ and vector $\boldsymbol{\theta}_h$:

$$\begin{aligned}
 \phi(s_{h+1}|s_h, \mathbf{a}) &= (\alpha(1-\delta), -\beta\mathbf{a}), h \in [H], \\
 \phi(s_{H+2}|s_h, \mathbf{a}) &= (\alpha\delta, \beta\mathbf{a}), h \in [H], \\
 \phi(s_{h+1}|s_h, \mathbf{a}) &= (\alpha, \mathbf{0}), h \in [H], \\
 \phi(s_{h+1}|s_h, \mathbf{a}) &= (0, \mathbf{0}), h \in [H],
 \end{aligned}$$

$$\boldsymbol{\theta}_h = (1/\alpha, \boldsymbol{\mu}_h/\beta), h \in [H],$$

where $\mathbf{0} = 0^{d-1}$ is a $(d-1)$ -dimensional vector of all zeros, $\alpha = \sqrt{1/(1 + (d-1)\Delta)}$ and $\beta = \sqrt{\Delta/(1 + (d-1)\Delta)}$. Since $B \geq 2$ and $K \geq (d-1)^2 H/2$, we have $\|\boldsymbol{\theta}_h\|_2 = (1 + \Delta(d-1))^2 \leq B$ and these hard-to-learn MDPs are B -bounded linear mixture MDPs.

Since the adversarial reward function is fixed across different episode k , the value function $V_{k,h}^\pi$ is also fixed for each policy π and the optimal policy π^* is pick the action $\mathbf{a}^* = \boldsymbol{\mu}_{h'}/\Delta$ at state $s_{h'} (1 \leq h' \leq H)$. Therefore, the adversarial MDP will degenerate to non-adversarial MDP and the adversarial regret is the same as the non-adversarial regret. For the lower bound of the non-adversarial regret, Theorem 5.6 (Zhou et al., 2021a) shows that for any algorithm, if $B \geq 2, H \geq 3, d \geq 4$ and $K \geq (d-1)^2 H/2$, then there exist a parameter $\boldsymbol{\mu}^* = \{\boldsymbol{\mu}_1^*, \dots, \boldsymbol{\mu}_H^*\}$ such that the expected regret is lower bounded by

$$\text{Regret}(K) = \Omega(dH\sqrt{T}).$$

Therefore, we finish the proof of Theorem 5.3. \square

B PROOF OF LEMMAS IN SECTIONS 4 AND 6

B.1 Proof of Lemma 4.1

We need the following Lemmas:

Lemma B.1 (Bernstein inequality for vector-valued martingales, Zhou et al. 2021a). Let $\{\mathcal{F}_t\}_{t=1}^\infty$ be a filtration and $(\mathbf{x}_t, \eta_t)_{t \geq 1}$ be a stochastic process so that $\mathbf{x}_t \in \mathbb{R}^d$ is \mathcal{F}_t -measurable and $\eta_t \in \mathbb{R}$ is \mathcal{F}_{t+1} -measurable. For constant $R, L, \sigma, \lambda > 0, \boldsymbol{\mu}^* \in \mathbb{R}^d$, let $y_t = \langle \mathbf{x}_t, \boldsymbol{\mu}^* \rangle + \eta_t$ and suppose that

$$|\eta_t| \leq R, \mathbb{E}[\eta_t | \mathcal{F}_t] = 0, \mathbb{E}[\eta_t^2 | \mathcal{F}_t] \leq \sigma^2, \|\mathbf{x}_t\|_2 \leq L.$$

Then, for any $0 \leq \delta \leq 1$, with probability at least $1 - \delta$, we have

$$\forall t > 0, \left\| \sum_{i=1}^t \mathbf{x}_i \eta_i \right\|_{\boldsymbol{\Sigma}_t^{-1}} \leq \beta_t, \|\boldsymbol{\mu}_t - \boldsymbol{\mu}^*\|_{\boldsymbol{\Sigma}_t} \leq \beta_t + \sqrt{\lambda} \|\boldsymbol{\mu}^*\|_2,$$

where $\boldsymbol{\Sigma}_t = \lambda \mathbf{I} + \sum_{i=1}^t \mathbf{x}_i \mathbf{x}_i^\top$, $\mathbf{b}_t = \sum_{i=1}^t \mathbf{x}_i y_i$, $\boldsymbol{\mu}_t = \boldsymbol{\Sigma}_t^{-1} \mathbf{b}_t$ and

$$\beta_t = 8\sigma \sqrt{d \log(1 + tL^2/(d\lambda)) \log(4t^2/\delta)} + 4R \log(4t^2/\delta).$$

Proof of Lemma 4.1. For each $h \in [H]$, by the definition of $[\bar{\mathbb{V}}_h V_{k,h+1}](s_h^k, a_h^k)$ in (4.2) and $[\mathbb{V}_h V_{k,h+1}](s_h^k, a_h^k)$ in (3.1), we have

$$\begin{aligned} & [\bar{\mathbb{V}}_h V_{k,h+1}](s_h^k, a_h^k) - [\mathbb{V}_h V_{k,h+1}](s_h^k, a_h^k) \\ &= \left[\langle \boldsymbol{\phi}_{V_{k,h+1}^2}(s_h^k, a_h^k), \tilde{\boldsymbol{\theta}}_{k,h} \rangle, H^2 \right]_{[0, H^2]} - \left[\langle \boldsymbol{\phi}_{V_{k,h+1}}(s_h^k, a_h^k), \hat{\boldsymbol{\theta}}_{k,h} \rangle_{[0, H]} \right]^2 \\ & \quad - \left\{ [\mathbb{P}_h V_{k,h+1}^2](s_h^k, a_h^k) - ([\mathbb{P}_h V_{k,h+1}](s_h^k, a_h^k))^2 \right\} \\ &= \underbrace{\left[\langle \boldsymbol{\phi}_{V_{k,h+1}^2}(s_h^k, a_h^k), \tilde{\boldsymbol{\theta}}_{k,h} \rangle \right]_{[0, H^2]} - [\mathbb{P}_h V_{k,h+1}^2](s_h^k, a_h^k)}_{I_1} \\ & \quad + \underbrace{\left([\mathbb{P}_h V_{k,h+1}](s_h^k, a_h^k) \right)^2 - \left[\langle \boldsymbol{\phi}_{V_{k,h+1}}(s_h^k, a_h^k), \hat{\boldsymbol{\theta}}_{k,h} \rangle_{[0, H]} \right]^2}_{I_2}. \end{aligned} \tag{B.1}$$

For the term I_1 , we have

$$|I_1| = \left| \left[\langle \boldsymbol{\phi}_{V_{k,h+1}^2}(s_h^k, a_h^k), \tilde{\boldsymbol{\theta}}_{k,h} \rangle \right]_{[0, H^2]} - [\mathbb{P}_h V_{k,h+1}^2](s_h^k, a_h^k) \right|$$

$$\begin{aligned}
 &\leq \left| \langle \phi_{V_{k,h+1}^2}(s_h^k, a_h^k), \tilde{\theta}_{k,h} \rangle - [\mathbb{P}_h V_{k,h+1}^2](s_h^k, a_h^k) \right| \\
 &= \left| \langle \phi_{V_{k,h+1}^2}(s_h^k, a_h^k), \tilde{\theta}_{k,h} \rangle - \langle \phi_{V_{k,h+1}^2}(s_h^k, a_h^k), \theta_h \rangle \right| \\
 &= \left| \langle \phi_{V_{k,h+1}^2}(s_h^k, a_h^k), \tilde{\theta}_{k,h} - \theta_h \rangle \right| \\
 &\leq \|\phi_{V_{k,h+1}^2}(s_h^k, a_h^k)\|_{\tilde{\Sigma}_{k,h}^{-1}} \|\tilde{\theta}_{k,h} - \theta_h\|_{\tilde{\Sigma}_{k,h}}, \tag{B.2}
 \end{aligned}$$

where the first inequality holds due to $0 \leq [\mathbb{P}_h V_{k,h+1}^2](s_h^k, a_h^k) \leq H^2$ and the second inequality holds due to Cauchy-Schwarz inequality. For the term $\|\tilde{\theta}_{k,h} - \theta_h\|_{\tilde{\Sigma}_{k,h}}$, we apply Lemma 4.1 with $\mathbf{x}_t = \phi_{V_{t,h+1}^2}(s_h^t, a_h^t)$, $\eta_t = V_{t,h+1}^2(s_{h+1}^t) - [\mathbb{P}_h V_{t,h+1}^2](s_h^t, a_h^t)$. For \mathbf{x}_t, η_t , we have the following property

$$\begin{aligned}
 \|\mathbf{x}_t\|_2 &= \|\phi_{V_{t,h+1}^2}(s_h^t, a_h^t)\|_2 \leq \max_{s'} V_{t,h+1}^2(s') \leq H^2, \\
 \mathbb{E}[\eta_t | \mathcal{F}_t] &= 0, |\eta_t| = |V_{t,h+1}^2(s_{h+1}^t) - [\mathbb{P}_h V_{t,h+1}^2](s_h^t, a_h^t)| \leq H^2, \\
 \mathbb{E}[\eta_t^2 | \mathcal{F}_t] &\leq H^4.
 \end{aligned}$$

Therefore, with probability at least $1 - \delta/H$, for all $k \in [K]$, we have

$$\|\tilde{\theta}_{k,h} - \theta_h\|_{\tilde{\Sigma}_{k,h}} \leq 8H^2 \sqrt{d \log(1 + kH^4/(d\lambda)) \log(4k^2H/\delta)} + 4H^2 \log(4k^2H/\delta) + \sqrt{\lambda}B. \tag{B.3}$$

Substituting (B.3) into (B.2), we have

$$\begin{aligned}
 |I_1| &\leq \|\phi_{V_{k,h+1}^2}(s_h^k, a_h^k)\|_{\tilde{\Sigma}_{k,h}^{-1}} \left(8H^2 \sqrt{d \log(1 + kH^4/(d\lambda)) \log(4k^2H/\delta)} + 4H^2 \log(4k^2H/\delta) + \sqrt{\lambda}B \right) \\
 &= \tilde{\beta}_k \|\tilde{\Sigma}_{k,h}^{-1/2} \phi_{V_{k,h+1}^2}(s_h^k, a_h^k)\|_2.
 \end{aligned}$$

Since both two terms of I_1 belong to the interval $[0, H^2]$, we have

$$|I_1| \leq \min \left\{ \tilde{\beta}_k \|\tilde{\Sigma}_{k,h}^{-1/2} \phi_{V_{k,h+1}^2}(s_h^k, a_h^k)\|_2, H^2 \right\}. \tag{B.4}$$

For the term I_2 , we have

$$\begin{aligned}
 |I_2| &= \left| ([\mathbb{P}_h V_{k,h+1}](s_h^k, a_h^k))^2 - \left[\langle \phi_{V_{k,h+1}}(s_h^k, a_h^k), \hat{\theta}_{k,h} \rangle_{[0,H]} \right]^2 \right| \\
 &= \left| \left[\langle \phi_{V_{k,h+1}}(s_h^k, a_h^k), \hat{\theta}_{k,h} \rangle_{[0,H]} - [\mathbb{P}_h V_{k,h+1}](s_h^k, a_h^k) \right] \right. \\
 &\quad \times \left. \left[\langle \phi_{V_{k,h+1}}(s_h^k, a_h^k), \hat{\theta}_{k,h} \rangle_{[0,H]} + [\mathbb{P}_h V_{k,h+1}](s_h^k, a_h^k) \right] \right| \\
 &\leq 2H \left| \left[\langle \phi_{V_{k,h+1}}(s_h^k, a_h^k), \hat{\theta}_{k,h} \rangle_{[0,H]} - [\mathbb{P}_h V_{k,h+1}](s_h^k, a_h^k) \right] \right| \\
 &\leq 2H \left| \langle \phi_{V_{k,h+1}}(s_h^k, a_h^k), \hat{\theta}_{k,h} \rangle - \langle \phi_{V_{k,h+1}}(s_h^k, a_h^k), \theta_h \rangle \right| \\
 &= 2H \left| \langle \phi_{V_{k,h+1}}(s_h^k, a_h^k), \hat{\theta}_{k,h} - \theta_h \rangle \right| \\
 &\leq 2H \|\phi_{V_{k,h+1}}(s_h^k, a_h^k)\|_{\tilde{\Sigma}_{k,h}^{-1}} \|\hat{\theta}_{k,h} - \theta_h\|_{\tilde{\Sigma}_{k,h}}, \tag{B.5}
 \end{aligned}$$

where the first inequality and second inequality holds due to $0 \leq [\mathbb{P}_h V_{k,h+1}](s_h^k, a_h^k) \leq H$ and the third inequality holds due to Cauchy-Schwarz inequality. For the term $\|\hat{\theta}_{k,h} - \theta_h\|_{\tilde{\Sigma}_{k,h}}$, we apply Lemma 4.1 with $\mathbf{x}_t = \bar{\sigma}_{k,h}^{-1} \phi_{V_{t,h+1}}(s_h^t, a_h^t)$, $\eta_t = \bar{\sigma}_{k,h}^{-1} V_{t,h+1}(s_{h+1}^t) - \bar{\sigma}_{k,h}^{-1} [\mathbb{P}_h V_{t,h+1}](s_h^t, a_h^t)$. For \mathbf{x}_t, η_t , we have following property

$$\begin{aligned}
 \|\mathbf{x}_t\|_2 &= \|\bar{\sigma}_{k,h}^{-1} \phi_{V_{t,h+1}}(s_h^t, a_h^t)\|_2 \leq \bar{\sigma}_{k,h}^{-1} \max_{s'} |V_{t,h+1}(s')| \leq \sqrt{d}, \\
 \mathbb{E}[\eta_t | \mathcal{F}_t] &= 0, |\eta_t| = |\bar{\sigma}_{k,h}^{-1} V_{t,h+1}(s_{h+1}^t) - \bar{\sigma}_{k,h}^{-1} [\mathbb{P}_h V_{t,h+1}](s_h^t, a_h^t)| \leq \sqrt{d},
 \end{aligned}$$

$$\mathbb{E}[\eta_t^2 | \mathcal{F}_t] \leq \sup \eta_t^2 \leq d.$$

Therefore, with probability at least $1 - \delta/H$, for all $k \in [K]$, we have

$$\|\widehat{\boldsymbol{\theta}}_{k,h} - \boldsymbol{\theta}_h\|_{\widehat{\boldsymbol{\Sigma}}_{k,h}} \leq 8d\sqrt{\log(1 + kH^4/(d\lambda)) \log(4k^2H/\delta)} + 4\sqrt{d} \log(4k^2H/\delta) + \sqrt{\lambda}B. \quad (\text{B.6})$$

Substituting (B.6) into (B.5), we have

$$\begin{aligned} |I_2| &\leq 2H \|\boldsymbol{\phi}_{V_{k,h+1}^2}(s_h^k, a_h^k)\|_{\widehat{\boldsymbol{\Sigma}}_{k,h}^{-1}} \left(8d\sqrt{\log(1 + kH^4/(d\lambda)) \log(4k^2H/\delta)} + 4\sqrt{d} \log(4k^2H/\delta) + \sqrt{\lambda}B \right) \\ &= \widehat{\beta}_k \|\widehat{\boldsymbol{\Sigma}}_{k,h}^{-1/2} \boldsymbol{\phi}_{V_{k,h+1}^2}(s_h^k, a_h^k)\|_2. \end{aligned}$$

Since both two terms of I_2 belong to the interval $[0, H^2]$, we have

$$|I_2| \leq \min \left\{ 2H\widehat{\beta}_k \|\widehat{\boldsymbol{\Sigma}}_{k,h}^{-1/2} \boldsymbol{\phi}_{V_{k,h+1}^2}(s_h^k, a_h^k)\|_2, H^2 \right\}. \quad (\text{B.7})$$

Substituting (B.4) and (B.7) into (B.1), with probability at least $1 - 2\delta/H$, we have

$$|[\bar{\mathbb{V}}_h V_{k,h+1}](s_h^k, a_h^k) - [\mathbb{V}_h V_{k,h+1}](s_h^k, a_h^k)| = |I_1 + I_2| \leq |I_1| + |I_2| \leq E_{k,h}, \quad (\text{B.8})$$

where

$$E_{k,h} = \min \left\{ \widetilde{\beta}_k \|\widetilde{\boldsymbol{\Sigma}}_{k,h}^{-1/2} \boldsymbol{\phi}_{V_{k,h+1}^2}(s_h^k, a_h^k)\|_2, H^2 \right\} + \min \left\{ 2H\bar{\beta}_k \|\widehat{\boldsymbol{\Sigma}}_{k,h}^{-1/2} \boldsymbol{\phi}_{V_{k,h+1}^2}(s_h^k, a_h^k)\|_2, H^2 \right\}.$$

We apply Lemma 4.1 again with $\mathbf{x}_t = \bar{\sigma}_{k,h}^{-1} \boldsymbol{\phi}_{V_{t,h+1}}(s_h^t, a_h^t)$, $\eta_t = \bar{\sigma}_{k,h}^{-1} V_{t,h+1}(s_{h+1}^t) - \bar{\sigma}_{k,h}^{-1} [\mathbb{P}_h V_{t,h+1}](s_h^t, a_h^t)$. For \mathbf{x}_t, η_t , we have following property

$$\begin{aligned} \|\mathbf{x}_t\|_2 &= \|\bar{\sigma}_{k,h}^{-1} \boldsymbol{\phi}_{V_{t,h+1}}(s_h^t, a_h^t)\|_2 \leq \bar{\sigma}_{k,h}^{-1} \max_{s'} V_{t,h+1}(s') \leq \sqrt{d} \\ \mathbb{E}[\eta_t | \mathcal{F}_t] &= 0, |\eta_t| = |\bar{\sigma}_{k,h}^{-1} V_{t,h+1}(s_{h+1}^t) - \bar{\sigma}_{k,h}^{-1} [\mathbb{P}_h V_{t,h+1}](s_h^t, a_h^t)| \leq \sqrt{d}. \end{aligned}$$

With probability at least $1 - 2\delta/H$, for all $t \in [K]$, we have

$$\mathbb{E}[\eta_t^2 | \mathcal{F}_t] = \bar{\sigma}_{t,h}^{-1} [\mathbb{V}_h V_{t,h+1}](s_h^t, a_h^t), \leq \bar{\sigma}_{t,h}^{-1} ([\bar{\mathbb{V}}_h V_{t,h+1}](s_h^t, a_h^t) + E_{t,h}) \leq 1, \quad (\text{B.9})$$

where the first inequality holds due to (B.8) and the second inequality holds due to the definition of $\bar{\sigma}_{t,h}^{-1}$. Therefore, with probability at least $1 - 3\delta/H$, for all $k \in [K]$, we have

$$\|\widehat{\boldsymbol{\theta}}_{k,h} - \boldsymbol{\theta}_h\|_{\widehat{\boldsymbol{\Sigma}}_{k,h}} \leq 8\sqrt{d} \log(1 + kH^4/(d\lambda)) \log(4k^2H/\delta) + 4\sqrt{d} \log(4k^2H/\delta) + \sqrt{\lambda}B = \widehat{\beta}_k.$$

Taking union bound for all $h \in [H]$, we finish the proof. \square

B.2 Proof of Lemma 6.1

Proof of Lemma 6.1. For each $h \in [H]$ and $s \in \mathcal{S}$, we have

$$\begin{aligned} V_{k,h}^*(s) - V_{k,h}(s) &= \mathbb{E}_{a \sim \pi_h^*(\cdot|s)} [Q_{k,h}^*(s, a)] - \mathbb{E}_{a \sim \pi_h^k(\cdot|s)} [Q_{k,h}(s, a)] \\ &= \underbrace{\mathbb{E}_{a \sim \pi_h^*(\cdot|s)} [Q_{k,h}^*(s, a)] - \mathbb{E}_{a \sim \pi_h^*(\cdot|s)} [Q_{k,h}(s, a)]}_I \\ &\quad + \mathbb{E}_{a \sim \pi_h^*(\cdot|s)} [Q_{k,h}(s, a)] - \mathbb{E}_{a \sim \pi_h^k(\cdot|s)} [Q_{k,h}(s, a)]. \end{aligned} \quad (\text{B.10})$$

For the term I , we have

$$\begin{aligned} I &= \mathbb{E}_{a \sim \pi_h^*(\cdot|s)} [Q_{k,h}^*(s, a) - Q_{k,h}(s, a)] \\ &\leq \mathbb{E}_{a \sim \pi_h^*(\cdot|s)} \left[[\mathbb{P}_h(V_{k,h+1}^* - V_{k,h+1})](s, a) \right] \\ &= \mathbb{E}_{a \sim \pi_h^*(\cdot|s), s' \sim \mathbb{P}_h(\cdot|s, a)} [V_{k,h+1}^*(s') - V_{k,h+1}(s')], \end{aligned} \quad (\text{B.11})$$

where the inequality holds due to Lemma A.2. Recursively using (B.11) with all $h \in [H]$, for all $k \in [K]$, we have

$$V_{k,1}^*(s_1^k) - V_{k,1}(s_1^k) \leq \mathbb{E} \left[\sum_{h=1}^H \left\{ \mathbb{E}_{a \sim \pi_h^*(\cdot|s_h)} [Q_{k,h}(s_h, a)] - \mathbb{E}_{a \sim \pi_h^k(\cdot|s_h)} [Q_{k,h}(s_h, a)] \right\} \middle| s_1 = s_1^k \right],$$

where $a_h \sim \pi_h^*(\cdot|s_h)$, $s_{h+1} \sim \mathbb{P}_h(\cdot|s_h, a_h)$. Therefore, we finish the proof. \square

B.3 Proof of Lemma 6.2

Proof of Lemma 6.2. By the update rule of the policy π_h^k , for all $k \in [K], h \in [H], s \in \mathcal{S}$, we have

$$\exp(\alpha Q_{k,h}(s,a)) = \frac{\pi_h^k(a|s) \exp\{\alpha Q_{h,k}(a|s)\}}{\pi_h^k(a|s)} = \frac{\rho \pi_h^{k+1}(a|s)}{\pi_h^k(a|s)}, \quad (\text{B.12})$$

where $\rho = \sum_{a \in \mathcal{A}} \pi_h^k(a|s) \exp\{\alpha Q_{h,k}(a|s)\}$ is fixed for all action a . Thus, we have

$$\begin{aligned} & \sum_{a \in \mathcal{A}} \alpha Q_{k,h}(s,a) (\pi_h^*(a|s) - \pi_h^{k+1}(a|s)) \\ &= \sum_{a \in \mathcal{A}} (\log \rho + \log \pi_h^{k+1}(a|s) - \log \pi_h^k(a|s)) (\pi_h^*(a|s) - \pi_h^{k+1}(a|s)) \\ &= \sum_{a \in \mathcal{A}} \pi_h^*(a|s) (\log \pi_h^*(a|s) - \log \pi_h^k(a|s)) \\ &= \sum_{a \in \mathcal{A}} \pi_h^*(a|s) (\log \pi_h^*(a|s) - \log \pi_h^{k+1}(a|s)) - \sum_{a \in \mathcal{A}} \pi_h^*(a|s) (\log \pi_h^{k+1}(a|s) - \log \pi_h^k(a|s)) \\ &\quad - \sum_{a \in \mathcal{A}} \pi_h^{k+1}(a|s) (\log \pi_h^{k+1}(a|s) - \log \pi_h^k(a|s)) \\ &= D_{KL}(\pi_h^*(\cdot|s) \| \pi_h^{k+1}(\cdot|s)) - D_{KL}(\pi_h^*(\cdot|s) \| \pi_h^k(\cdot|s)) - D_{KL}(\pi_h^{k+1}(\cdot|s) \| \pi_h^k(\cdot|s)), \end{aligned} \quad (\text{B.13})$$

where the first equation holds due to (B.12) and the second equation holds due to $\sum_{a \in \mathcal{A}} (\pi_h^*(a|s) - \pi_h^{k+1}(a|s)) = 0$. Therefore, we have

$$\begin{aligned} & \mathbb{E}_{a \sim \pi_h^*(\cdot|s_h)} [Q_{k,h}(s,a)] - \mathbb{E}_{a \sim \pi_h^k(\cdot|s)} [Q_{k,h}(s,a)] \\ &= \sum_{a \in \mathcal{A}} Q_{k,h}(s,a) (\pi_h^*(a|s) - \pi_h^k(a|s)) \\ &= \sum_{a \in \mathcal{A}} Q_{k,h}(s,a) (\pi_h^*(a|s) - \pi_h^{k+1}(a|s)) + \sum_{a \in \mathcal{A}} Q_{k,h}(s,a) (\pi_h^{k+1}(a|s) - \pi_h^k(a|s)) \\ &\leq \sum_{a \in \mathcal{A}} Q_{k,h}(s,a) (\pi_h^*(a|s) - \pi_h^{k+1}(a|s)) + H \|\pi_h^{k+1}(\cdot|s) - \pi_h^k(\cdot|s)\|_1 \\ &= \alpha^{-1} \left(D_{KL}(\pi_h^*(\cdot|s) \| \pi_h^{k+1}(\cdot|s)) - D_{KL}(\pi_h^*(\cdot|s) \| \pi_h^k(\cdot|s)) - D_{KL}(\pi_h^{k+1}(\cdot|s) \| \pi_h^k(\cdot|s)) \right) \\ &\quad + H \|\pi_h^{k+1}(\cdot|s) - \pi_h^k(\cdot|s)\|_1 \\ &\leq \alpha^{-1} \left(D_{KL}(\pi_h^*(\cdot|s) \| \pi_h^{k+1}(\cdot|s)) - D_{KL}(\pi_h^*(\cdot|s) \| \pi_h^k(\cdot|s)) \right) \\ &\quad + H \|\pi_h^{k+1}(\cdot|s) - \pi_h^k(\cdot|s)\|_1 - \frac{\|\pi_h^{k+1}(\cdot|s) - \pi_h^k(\cdot|s)\|_1^2}{2\alpha} \\ &\leq \frac{\alpha H^2}{2} + \alpha^{-1} \left(D_{KL}(\pi_h^*(\cdot|s_h) \| \pi_h^k(\cdot|s_h)) - D_{KL}(\pi_h^*(\cdot|s_h) \| \pi_h^{k+1}(\cdot|s_h)) \right), \end{aligned} \quad (\text{B.14})$$

where the first inequality holds due to the fact that $0 \leq Q_{k,h}^{\pi_h^k}(s,a) \leq Q_{k,h}(s,a) \leq H$, the second inequality holds due to Pinsker's inequality and the last inequality holds due to the fact that $ax - bx^2 \leq a^2/4b$. Therefore, we finish the proof. \square

B.4 Proof of Lemma 6.3

Proof of Lemma 6.3.

$$\begin{aligned} & Q_{k,h}(s_h^k, a_h^k) - Q_{k,h}^{\pi_h^k}(s_h^k, a_h^k) \\ &= \left[r_h^k(s_h^k, a_h^k) + \langle \widehat{\theta}_{k,h}, \phi_{V_{k,h+1}}(s_h^k, a_h^k) \rangle + \widehat{\beta}_k \left\| \widehat{\Sigma}_{k,h}^{-1/2} \phi_{V_{k,h+1}}(s_h^k, a_h^k) \right\|_2 \right]_{[0, H-h+1]} \end{aligned}$$

$$\begin{aligned}
 & -r_h^k(s_h^k, a_h^k) - [\mathbb{P}_h V_{k,h}^{\pi^k}](s_h^k, a_h^k) \\
 \leq & \left| \langle \widehat{\theta}_{k,h}, \phi_{V_{k,h+1}}(s_h^k, a_h^k) \rangle + \widehat{\beta}_k \|\widehat{\Sigma}_{k,h}^{-1/2} \phi_{V_{k,h+1}}(s_h^k, a_h^k)\|_2 \right| - [\mathbb{P}_h V_{k,h}^{\pi^k}](s_h^k, a_h^k) \\
 \leq & [\mathbb{P}_h V_{k,h+1}](s_h^k, a_h^k) + \left| \langle \widehat{\theta}_{k,h} - \theta_h, \phi_{V_{k,h+1}}(s_h^k, a_h^k) \rangle \right| \\
 & + \widehat{\beta}_k \|\widehat{\Sigma}_{k,h}^{-1/2} \phi_{V_{k,h+1}}(s_h^k, a_h^k)\|_2 - [\mathbb{P}_h V_{k,h+1}^{\pi^k}](s_h^k, a_h^k) \\
 \leq & [\mathbb{P}_h(V_{k,h+1} - V_{k,h+1}^{\pi^k})](s_h^k, a_h^k) + 2\widehat{\beta}_k \|\widehat{\Sigma}_{k,h}^{-1/2} \phi_{V_{k,h+1}}(s_h^k, a_h^k)\|_2 \\
 = & [\mathbb{P}_h(V_{k,h+1} - V_{k,h+1}^{\pi^k})](s_h^k, a_h^k) + 2\widehat{\beta}_k \bar{\sigma}_{k,h} \|\widehat{\Sigma}_{k,h}^{-1/2} \phi_{V_{k,h+1}}(s_h^k, a_h^k) / \bar{\sigma}_{k,h}\|_2, \tag{B.15}
 \end{aligned}$$

where the first inequality holds due to the fact that $x_{[0,z]} - y \leq |x - y|$ when $y \geq 0$, the second inequality holds due to the fact that $|x + y + z| \leq |x| + |y| + |z|$ and the third inequality holds due to event \mathcal{E} . Furthermore, we have

$$Q_{k,h}(s_h^k, a_h^k) - Q_{k,h}^{\pi^k}(s_h^k, a_h^k) \leq H - Q_{k,h}^{\pi^k} \leq H \leq 2\widehat{\beta}_k \bar{\sigma}_{k,h}, \tag{B.16}$$

where the first inequality holds due to $Q_{k,h}(s_h^k, a_h^k) \leq H$, the second inequality holds due to $Q_{k,h}^{\pi^k}(s_h^k, a_h^k) \geq 0$ and the last inequality holds due to $2\widehat{\beta}_k \bar{\sigma}_{k,h} \geq \sqrt{d}H/\sqrt{d} = H$. Combined (B.15) and (B.16), we have

$$\begin{aligned}
 Q_{k,h}(s_h^k, a_h^k) - Q_{k,h}^{\pi^k}(s_h^k, a_h^k) \leq & [\mathbb{P}_h(V_{k,h+1} - V_{k,h+1}^{\pi^k})](s_h^k, a_h^k) \\
 & + 2\widehat{\beta}_k \bar{\sigma}_{k,h} \min \left\{ \|\widehat{\Sigma}_{k,h}^{-1/2} \phi_{V_{k,h+1}}(s_h^k, a_h^k) / \bar{\sigma}_{k,h}\|_2, 1 \right\}.
 \end{aligned}$$

Therefore, we finish the proof. \square

B.5 Proof of Lemma 6.4

Proof of Lemma 6.4.

$$\begin{aligned}
 & V_{k,h}(s_h^k) - V_{k,h}^{\pi^k}(s_h^k) \\
 = & \mathbb{E}_{a \sim \pi_h^k(\cdot | s_h^k)} [Q_{k,h}(s_h^k, a) - Q_{k,h}^{\pi^k}(s_h^k, a)] \\
 = & \mathbb{E}_{a \sim \pi_h^k(\cdot | s_h^k)} [Q_{k,h}(s_h^k, a) - Q_{k,h}^{\pi^k}(s_h^k, a)] - (Q_{k,h}(s_h^k, a_h^k) - Q_{k,h}^{\pi^k}(s_h^k, a_h^k) + Q_{k,h}(s_h^k, a_h^k) - Q_{k,h}^{\pi^k}(s_h^k, a_h^k)) \\
 \leq & \mathbb{E}_{a \sim \pi_h^k(\cdot | s_h^k)} [Q_{k,h}(s_h^k, a) - Q_{k,h}^{\pi^k}(s_h^k, a)] - (Q_{k,h}(s_h^k, a_h^k) - Q_{k,h}^{\pi^k}(s_h^k, a_h^k) + [\mathbb{P}_h(V_{k,h+1} - V_{k,h+1}^{\pi^k})](s_h^k, a_h^k) \\
 & + 2\widehat{\beta}_k \bar{\sigma}_{k,h} \min \left\{ \|\widehat{\Sigma}_{k,h}^{-1/2} \phi_{V_{k,h+1}}(s_h^k, a_h^k) / \bar{\sigma}_{k,h}\|_2, 1 \right\}) \\
 = & \mathbb{E}_{a \sim \pi_h^k(\cdot | s_h^k)} [Q_{k,h}(s_h^k, a) - Q_{k,h}^{\pi^k}(s_h^k, a)] - (Q_{k,h}(s_h^k, a_h^k) - Q_{k,h}^{\pi^k}(s_h^k, a_h^k) \\
 & + V_{k,h+1}(s_{h+1}^k) - V_{k,h+1}^{\pi^k}(s_{h+1}^k) + [\mathbb{P}_h(V_{k,h+1} - V_{k,h+1}^{\pi^k})](s_h^k, a_h^k) - (V_{k,h+1}(s_{h+1}^k) - V_{k,h+1}^{\pi^k}(s_{h+1}^k)) \\
 & + 2\widehat{\beta}_k \bar{\sigma}_{k,h} \min \left\{ \|\widehat{\Sigma}_{k,h}^{-1/2} \phi_{V_{k,h+1}}(s_h^k, a_h^k) / \bar{\sigma}_{k,h}\|_2, 1 \right\}), \tag{B.17}
 \end{aligned}$$

where the inequality holds due to Lemma 6.3. Furthermore, on the event \mathcal{E} and \mathcal{E}_1 , for all $h \in [H]$, we have

$$\begin{aligned}
 & \sum_{k=1}^K (V_{k,h}(s_h^k) - V_{k,h}^{\pi^k}(s_h^k)) \\
 \leq & \sum_{k=1}^K \sum_{h'=h}^H 2\widehat{\beta}_k \bar{\sigma}_{k,h} \min \left\{ \|\widehat{\Sigma}_{k,h}^{-1/2} \phi_{V_{k,h+1}}(s_h^k, a_h^k) / \bar{\sigma}_{k,h}\|_2, 1 \right\} \\
 & + \sum_{k=1}^K \sum_{h'=h}^H \left(\mathbb{E}_{a \sim \pi_h^k(\cdot | s_h^k)} [Q_{k,h}(s_h^k, a) - Q_{k,h}^{\pi^k}(s_h^k, a)] - (Q_{k,h}(s_h^k, a_h^k) - Q_{k,h}^{\pi^k}(s_h^k, a_h^k)) \right) \\
 & + \sum_{k=1}^K \sum_{h'=h}^H \left([\mathbb{P}_h(V_{k,h+1} - V_{k,h+1}^{\pi^k})](s_h^k, a_h^k) - (V_{k,h+1}(s_{h+1}^k) - V_{k,h+1}^{\pi^k}(s_{h+1}^k)) \right)
 \end{aligned}$$

$$\begin{aligned}
 &\leq \sum_{k=1}^K \sum_{h'=h}^H 2\widehat{\beta}_k \bar{\sigma}_{k,h} \min \left\{ \left\| \widehat{\Sigma}_{k,h}^{-1/2} \phi_{V_{k,h+1}}(s_h^k, a_h^k) / \bar{\sigma}_{k,h} \right\|_2, 1 \right\} + 4H\sqrt{T \log(H/\delta)} \\
 &\leq 2\widehat{\beta}_K \sum_{k=1}^K \sum_{h'=h}^H \bar{\sigma}_{k,h} \min \left\{ \left\| \widehat{\Sigma}_{k,h}^{-1/2} \phi_{V_{k,h+1}}(s_h^k, a_h^k) / \bar{\sigma}_{k,h} \right\|_2, 1 \right\} + 4H\sqrt{T \log(H/\delta)} \\
 &\leq 2\widehat{\beta}_K \sqrt{\sum_{k=1}^K \sum_{h=1}^H \bar{\sigma}_{k,h}^2} \sqrt{\sum_{k=1}^K \sum_{h=1}^H \min \left\{ \left\| \widehat{\Sigma}_{k,h}^{-1/2} \phi_{V_{k,h+1}}(s_h^k, a_h^k) / \bar{\sigma}_{k,h} \right\|_2, 1 \right\}} + 4H\sqrt{T \log(H/\delta)} \\
 &\leq 2\widehat{\beta}_K \sqrt{\sum_{k=1}^K \sum_{h=1}^H \bar{\sigma}_{k,h}^2} \sqrt{2Hd \log(1 + K/\lambda)} + 4H\sqrt{T \log(H/\delta)}, \tag{B.18}
 \end{aligned}$$

where the first inequality holds by taking the summation of (B.17) for $k \in [K]$ and $h \leq h' \leq H$, the second inequality holds due to the definition of event \mathcal{E}_1 , the third inequality holds due to $\widehat{\beta}_k \leq \widehat{\beta}_K$, the fourth inequality holds due to Cauchy-Schwarz inequality and the last inequality holds due to Lemma A.3. Furthermore, taking the summation of (B.18), we have

$$\begin{aligned}
 &\sum_{k=1}^K \sum_{h=1}^H [\mathbb{P}_h(V_{k,h+1} - V_{k,h+1}^{\pi^k})](s_h^k, a_h^k) \\
 &= \sum_{k=1}^K \sum_{h=1}^H (V_{k,h+1}(s_{h+1}^k) - V_{k,h+1}^{\pi^k}(s_{h+1}^k)) \\
 &\quad + \sum_{k=1}^K \sum_{h=1}^H \left([\mathbb{P}_h(V_{k,h+1} - V_{k,h+1}^{\pi^k})](s_h^k, a_h^k) - (V_{k,h+1}(s_{h+1}^k) - V_{k,h+1}^{\pi^k}(s_{h+1}^k)) \right) \\
 &\leq \sum_{k=1}^K \sum_{h=1}^H (V_{k,h+1}(s_{h+1}^k) - V_{k,h+1}^{\pi^k}(s_{h+1}^k)) + 2H\sqrt{2T \log(1/\delta)} \\
 &\leq 2H\widehat{\beta}_K \sqrt{\sum_{k=1}^K \sum_{h=1}^H \bar{\sigma}_{k,h}^2} \sqrt{2Hd \log(1 + K/\lambda)} + 4H^2\sqrt{T \log(H/\delta)},
 \end{aligned}$$

where the first inequality holds due to the definition of event \mathcal{E}_2 and the last inequality holds due (B.18). Therefore, we finish the proof. \square

B.6 Proof of Lemma 6.5

Proof of Lemma 6.5. On the event \mathcal{E} , by Lemma 4.1, for all $k \in [K], h \in [H]$, we have

$$[\bar{\mathbb{V}}_{k,h} V_{k,h+1}](s_h^k, a_h^k) + E_{k,h} \geq [\mathbb{V}_h V_{k,h+1}](s_h^k, a_h^k) \geq 0.$$

Therefore, we have

$$\begin{aligned}
 \sum_{k=1}^K \sum_{h=1}^H \bar{\sigma}_{k,h}^2 &= \sum_{k=1}^K \sum_{h=1}^H \max \{ H^2/d, [\bar{\mathbb{V}}_{k,h} V_{k,h+1}](s_h^k, a_h^k) + E_{k,h} \} \\
 &\leq \sum_{k=1}^K \sum_{h=1}^H \frac{H^2}{d} + \sum_{k=1}^K \sum_{h=1}^H [\bar{\mathbb{V}}_{k,h} V_{k,h+1}](s_h^k, a_h^k) + \sum_{k=1}^K \sum_{h=1}^H E_{k,h} \\
 &= \frac{H^2 T}{d} + \underbrace{\sum_{k=1}^K \sum_{h=1}^H ([\bar{\mathbb{V}}_{k,h} V_{k,h+1}](s_h^k, a_h^k) - [\mathbb{V}_h V_{k,h+1}](s_h^k, a_h^k))}_{I_1} + \underbrace{\sum_{k=1}^K \sum_{h=1}^H E_{k,h}}_{I_2}
 \end{aligned}$$

$$+ \underbrace{\sum_{k=1}^K \sum_{h=1}^H ([\mathbb{V}_h V_{k,h+1}](s_h^k, a_h^k) - [\mathbb{V}_h V_{k,h+1}^{\pi^k}](s_h^k, a_h^k))}_{I_3} + \underbrace{\sum_{k=1}^K \sum_{h=1}^H [\mathbb{V}_h V_{k,h+1}^{\pi^k}](s_h^k, a_h^k)}_{I_4}, \quad (\text{B.19})$$

where the inequality holds due to the fact that $\max\{a, b\} \leq a + b$, when $a, b \geq 0$. For the term I_1 , we have

$$I_1 = \sum_{k=1}^K \sum_{h=1}^H ([\bar{\mathbb{V}}_{k,h} V_{k,h+1}](s_h^k, a_h^k) - [\mathbb{V}_h V_{k,h+1}](s_h^k, a_h^k)) \leq \sum_{k=1}^K \sum_{h=1}^H E_{k,h} = I_2, \quad (\text{B.20})$$

where the inequality holds due to the definition of event \mathcal{E} . For the term I_2 , we have

$$\begin{aligned} I_2 &= \sum_{k=1}^K \sum_{h=1}^H E_{k,h} \\ &= \sum_{k=1}^K \sum_{h=1}^H \min \left\{ \tilde{\beta}_k \|\tilde{\Sigma}_{k,h}^{-1/2} \phi_{V_{k,h+1}^2}(s_h^k, a_h^k)\|_2, H^2 \right\} + \sum_{k=1}^K \sum_{h=1}^H \min \left\{ 2H\bar{\beta}_k \|\widehat{\Sigma}_{k,h}^{-1/2} \phi_{V_{k,h+1}}(s_h^k, a_h^k)\|_2, H^2 \right\} \\ &\leq \tilde{\beta}_K \sum_{k=1}^K \sum_{h=1}^H \min \left\{ \|\tilde{\Sigma}_{k,h}^{-1/2} \phi_{V_{k,h+1}^2}(s_h^k, a_h^k)\|_2, 1 \right\} \\ &\quad + 2H\bar{\beta}_K \sum_{k=1}^K \sum_{h=1}^H \bar{\sigma}_{k,h} \min \left\{ \|\widehat{\Sigma}_{k,h}^{-1/2} \phi_{V_{k,h+1}}(s_h^k, a_h^k)\|_2, 1 \right\} \\ &\leq \tilde{\beta}_K \sqrt{T} \sqrt{\sum_{k=1}^K \sum_{h=1}^H \min \left\{ \|\tilde{\Sigma}_{k,h}^{-1/2} \phi_{V_{k,h+1}^2}(s_h^k, a_h^k)\|_2^2, 1 \right\}} \\ &\quad + 2\sqrt{3}H^2\bar{\beta}_K \sqrt{T} \sqrt{\sum_{k=1}^K \sum_{h=1}^H \min \left\{ \|\widehat{\Sigma}_{k,h}^{-1/2} \phi_{V_{k,h+1}}(s_h^k, a_h^k)\|_2^2, 1 \right\}} \\ &\leq \tilde{\beta}_K \sqrt{T} \sqrt{2dH \log(1 + kH^4/(d\lambda))} + 2H^2\bar{\beta}_K \sqrt{3T} \sqrt{2dH \log(1 + K/\lambda)}, \end{aligned} \quad (\text{B.21})$$

where the first inequality holds due to $\tilde{\beta}_K \geq \tilde{\beta}_k \geq H^2$ and $\bar{\beta}_K \bar{\sigma}_{k,h} \geq \bar{\beta}_k \bar{\sigma}_{k,h} \geq H$, the second inequality holds due to Cauchy-Schwarz inequality with the fact that $\bar{\sigma}_{k,h} \leq \max\{H^2/d, H^2 + 2H^2\} \leq 3H^2$ and the last inequality holds due to Lemma A.3.

On the event $\mathcal{E} \cap \mathcal{E}_1 \cap \mathcal{E}_2$, for the term I_3 , we have

$$\begin{aligned} I_3 &= \sum_{k=1}^K \sum_{h=1}^H ([\mathbb{V}_h V_{k,h+1}](s_h^k, a_h^k) - [\mathbb{V}_h V_{k,h+1}^{\pi^k}](s_h^k, a_h^k)) \\ &= \sum_{k=1}^K \sum_{h=1}^H \left([\mathbb{P}_h V_{k,h+1}^2](s_h^k, a_h^k) - ([\mathbb{P}_h V_{k,h+1}](s_h^k, a_h^k))^2 - [\mathbb{P}_h (V_{k,h+1}^{\pi^k})^2](s_h^k, a_h^k) + ([\mathbb{P}_h V_{k,h+1}^{\pi^k}](s_h^k, a_h^k))^2 \right) \\ &\leq \sum_{k=1}^K \sum_{h=1}^H \left([\mathbb{P}_h V_{k,h+1}^2](s_h^k, a_h^k) - [\mathbb{P}_h (V_{k,h+1}^{\pi^k})^2](s_h^k, a_h^k) \right) \\ &\leq 2H \sum_{k=1}^K \sum_{h=1}^H \left([\mathbb{P}_h V_{k,h+1}](s_h^k, a_h^k) - [\mathbb{P}_h V_{k,h+1}^{\pi^k}](s_h^k, a_h^k) \right) \\ &\leq 4H^2 \widehat{\beta}_K \sqrt{\sum_{k=1}^K \sum_{h=1}^H \bar{\sigma}_{k,h}^2 \sqrt{2Hd \log(1 + K/\lambda)} + 8H^3 \sqrt{T \log(H/\delta)}}, \end{aligned} \quad (\text{B.22})$$

where the first inequality holds due to the fact that $V_{k,h+1}^{\pi^k}(s') \leq V_{k,h+1}(s')$, the second inequality holds due to $0 \leq V_{k,h+1}(s'), V_{k,h+1}^{\pi^k}(s') \leq H$ and the last inequality holds due to Lemma 6.4.

On the event \mathcal{E}_3 , for the term I_4 , we have

$$I_4 = \sum_{k=1}^K \sum_{h=1}^H [\mathbb{V}_h V_{k,h+1}^{\pi^k}](s_h^k, a_h^k) \leq 3(HT + H^3 \log(1/\delta)). \quad (\text{B.23})$$

Substituting (B.20), (B.21), (B.22) and (B.23) into (B.19), we have

$$\begin{aligned} \sum_{k=1}^K \sum_{h=1}^H \bar{\sigma}_{k,h}^{-2} &\leq H^2 T/d + 3(HT + H^3 \log(1/\delta)) \\ &\quad + 2\tilde{\beta}_K \sqrt{T} \sqrt{2dH \log(1 + KH^4/(d\lambda))} + 4H^2 \bar{\beta}_K \sqrt{3T} \sqrt{2dH \log(1 + K/\lambda)} \\ &\quad + 4H^2 \widehat{\beta}_K \sqrt{\sum_{k=1}^K \sum_{h=1}^H \bar{\sigma}_{k,h}^2} \sqrt{2Hd \log(1 + K/\lambda) + 8H^3 \sqrt{T \log(H/\delta)}}, \end{aligned}$$

where

$$\begin{aligned} \bar{\beta}_K &= 8d \sqrt{\log(1 + K/\lambda) \log(4K^2 H/\delta)} + 4\sqrt{d} \log(4K^2 H/\delta) + \sqrt{\lambda} B, \\ \widehat{\beta}_K &= 8\sqrt{d} \log(1 + K/\lambda) \log(4K^2 H/\delta) + 4\sqrt{d} \log(4K^2 H/\delta) + \sqrt{\lambda} B, \\ \tilde{\beta}_K &= 8H^2 \sqrt{d \log(1 + KH^4/(d\lambda)) \log(4K^2 H/\delta)} + 4H^2 \log(4K^2 H/\delta) + \sqrt{\lambda} B. \end{aligned}$$

Therefore, by the fact that $x \leq a\sqrt{x} + b$ implies $x \leq a^2 + 2b$, we have

$$\begin{aligned} \sum_{k=1}^K \sum_{h=1}^H \bar{\sigma}_{k,h}^{-2} &\leq 2H^2 T/d + 6(HT + H^3 \log(1/\delta)) \\ &\quad + 4\tilde{\beta}_K \sqrt{T} \sqrt{2dH \log(1 + KH^4/(d\lambda))} + 8H^2 \bar{\beta}_K \sqrt{3T} \sqrt{2dH \log(1 + K/\lambda)} \\ &\quad + 32H^5 d (\widehat{\beta}_K)^2 \log(1 + K/\lambda) + 16H^3 \sqrt{T \log(H/\delta)} \\ &\leq 2H^2 T/d + 6(HT + H^3 \log(1/\delta)) + 330\sqrt{d^3 H^5 T} \log(4K^2 H/\delta) \log(1 + KH^4/\lambda) \\ &\quad + 2048H^5 d^2 \log^2(4K^2 H/\delta) \log^2(1 + K/\lambda) + 16H^3 \sqrt{T \log(H/\delta)} \\ &\leq 2HT/d + 179HT + 165d^3 H^4 \log^2(4K^2 H/\delta) \log^2(1 + KH^4/\lambda) \\ &\quad + 2062H^5 d^2 \log^2(4K^2 H/\delta) \log^2(1 + K/\lambda), \end{aligned}$$

where the second inequality holds due to the definition of parameter $\bar{\beta}_K, \widehat{\beta}_K, \tilde{\beta}_K$ with the fact that $\lambda = 1/B^2 \leq 1$ and the third inequality holds due to Young's inequality. Therefore, we finish the proof. \square