
Nearly Minimax Optimal Reinforcement Learning with Linear Function Approximation

Pihe Hu¹ Yu Chen² Longbo Huang¹

Abstract

We study reinforcement learning with linear function approximation where the transition probability and reward functions are linear with respect to a feature mapping $\phi(s, a)$. Specifically, we consider the episodic inhomogeneous linear Markov Decision Process (MDP), and propose a novel computation-efficient algorithm, LSVI-UCB⁺, which achieves an $\tilde{O}(Hd\sqrt{T})$ regret bound where H is the episode length, d is the feature dimension, and T is the number of steps. LSVI-UCB⁺ builds on weighted ridge regression and upper confidence value iteration with a Bernstein-type exploration bonus. Our statistical results are obtained with novel analytical tools, including a new Bernstein self-normalized bound with conservatism on elliptical potentials, and refined analysis of the correction term. To the best of our knowledge, this is the first minimax optimal algorithm for linear MDPs up to logarithmic factors, which closes the \sqrt{Hd} gap between the best known upper bound of $\tilde{O}(\sqrt{H^3d^3T})$ in (Jin et al., 2020) and lower bound of $\Omega(Hd\sqrt{T})$ for linear MDPs.

1. Introduction

Reinforcement Learning (RL) has demonstrated phenomenal empirical success in many areas, including games, robotic control, etc., where improving sample complexity is always an important topic. When the state space and action space are finite, the Markov decision process (MDP) has been proven to achieve nearly minimax optimal sample-complexity with the generative model in (Azar et al., 2013). For harder RL settings, nearly minimax optimal sample-complexities are obtained in (Azar et al., 2017) for finite

horizon episodic MDPs, and (He et al., 2021; Tossou et al., 2019) for infinite horizon MDPs¹. However, MDPs are known to suffer from the curse-of-dimensionality due to large and possibly infinite state and action space.

Function approximation is an essential approach for handling large MDPs, which assumes that the problem structure has a compact representation concerning state or state-action pairs and enables the development of nearly minimax optimal theoretical guarantees for RL problems. Linear function approximation is one of the most fundamental function approximations. It has a significant impact since many problems can be linearly-parameterized structurally or combined with embedding, where linear MDPs and linear mixture MDPs are two of the most popular models. Representative works for these two settings are presented in Table 1, from which we find that the best regret upper bound for linear MDPs and linear mixture MDPs are $\tilde{O}(\sqrt{H^3d^3T})$ by LSVI-UCB in (Jin et al., 2020) and $\tilde{O}(\sqrt{H^2d^2T + H^3d^3T})$ by UCRL-VTR⁺ in (Zhou et al., 2021), respectively. However, the best-known lower bound for these two settings are both $\Omega(Hd\sqrt{T})$ according to (Zhou et al., 2021), which means the following fundamental question remains open:

Does there exist a computation-efficient and minimax optimal algorithm for RL with linear function approximation?

We answer this question affirmatively for linear MDPs by designing the LSVI-UCB⁺ algorithm, which reaches minimax optimal regret up to logarithmic factors. LSVI-UCB⁺ overcomes barriers to nearly minimax optimality in existing works (Jin et al., 2020; Wang et al., 2020b; Zanette et al., 2020b; Wang et al., 2020a) for linear MDPs and their variations, including overly aggressive exploration and extra cost for building a uniform convergence argument by covering net. It constructs a Bernstein-type bonus to perform efficient exploration, which enables a \sqrt{H} factor reduction in regret. Besides, the extra \sqrt{d} dependency from building the uniform convergence argument can be removed by our novel technique of bounding the correction term $(\hat{\mathbb{P}} - \mathbb{P})(\hat{V} - V^*)$. Notably, minimax optimal algorithms (Azar et al., 2017; Zanette & Brunskill, 2019) for tabular MDPs also utilize the

¹Institute for Interdisciplinary Information Sciences, Tsinghua University, Beijing, China ²Department of Mathematical Sciences, Tsinghua University, Beijing, China. Correspondence to: Longbo Huang <longbohuang@tsinghua.edu.cn>.

¹An algorithm is nearly minimax optimal if its sample complexity matches the minimax lower bound up to logarithmic factors.

Table 1. Theoretical results on RL with linear function approximation, where † denotes that rewards are adversarial, and the lower bound holds for both settings.

Algorithm	Regret
Linear MDP	
OPT-RLSVI (Zanette et al., 2020a)	$\tilde{O}(H^2 d^2 \sqrt{T})$
LSVI-UCB (Jin et al., 2020)	$\tilde{O}(\sqrt{H^3 d^3 T})$
LSVI-UCB ⁺ (this paper)	$\tilde{O}(Hd\sqrt{T})$
Linear Mixture MDP	
OPPO [†] (Cai et al., 2020)	$\tilde{O}(\sqrt{H^3 d^2 T})$
UCRL-VTR (Ayoub et al., 2020)	$\tilde{O}(\sqrt{H^3 d^2 T})$
UCRL-VTR ⁺ (Zhou et al., 2021)	$\tilde{O}(\sqrt{H^2 d^2 T + H^3 d T})$
Lower Bound (Zhou et al., 2021)	$\Omega(Hd\sqrt{T})$

Bernstein-type bonus for exploration with refined consideration of the correction term in analysis. However, results for tabular MDPs cannot be applied directly to our settings due to the need for building the Bernstein inequality for vector-valued martingales in linear settings. It is worth mentioning that the above Bernstein inequality has been studied in (Zhou et al., 2021), where the UCRL-VTR⁺ is nearly minimax optimal when $d \geq H$ for linear mixture MDPs. By contrast, our proposed LSVI-UCB⁺ algorithm achieves nearly minimax optimal regret without requiring this assumption for linear MDPs and can be further generalized to linear mixture MDPs such that the nearly minimax optimal regret can also be obtained without $d \geq H$. This is because our proposed Bernstein inequality (Theorem 7.1) is sharper by considering the conservatism on elliptical potentials. Our contributions are summarized below:

- We develop a novel Bernstein bound of $\tilde{O}(\sigma\sqrt{d} + R)$ for self-normalized martingales, which is sharper than the analog inequalities in (Zhou et al., 2021). By utilizing the conservatism on elliptical potentials, the bound can be further improved to $\tilde{O}(\sigma\sqrt{d})$, which serves as a new analytical tool for RL.
- We propose the LSVI-UCB⁺ algorithm based on a Bernstein-type exploration bonus and weighted ridge regression, with weights determined by value function variances and exploration uncertainty. LSVI-UCB⁺ achieves an $\tilde{O}(Hd\sqrt{T})$ regret, and is minimax optimal up to logarithmic factors in large-sample regime.
- We improve the analytical framework of statistical complexity for linear MDPs by bounding the correction term $(\hat{\mathbb{P}} - \mathbb{P})(\hat{V} - V)$. Combined with the Bernstein self-normalized bound, this new analytical framework can remove the extra dependencies on H and d , which is very different from the traditional Hoeffding bound used in (Jin et al., 2020; Wang et al., 2020b;a).

Notations Scalars are denoted in lower case letters, and vectors/matrices are denoted in boldface letters. Denote $\|\mathbf{x}\|_{\Lambda}^2 = \mathbf{x}^\top \Lambda \mathbf{x}$ for vector \mathbf{x} and positive definite matrix Λ . Denote $\{1, \dots, n\}$ as $[n]$ and the truncated value of x in $[a, b]$ interval as $[x]_{[a,b]}$ for $a \leq b$. Define $a_n = O(b_n)$ if there exists an absolute constant $c > 0$ such that $a_n \leq cb_n$ holds for all $n \geq 1$ and define $a_n = \Omega(b_n)$ for inverse direction. $\tilde{O}(\cdot)$ further suppresses the polylogarithmic factors in $O(\cdot)$.

2. Related Work

Linear Bandits Linear stochastic bandits can be regarded as a special case of linearly-parameterized MDPs with episode length 1. (Dani et al., 2008) proposes an algorithm with $O(d\sqrt{T \log^3 T})$ regret by building confidence ball with Freedman inequality (Freedman, 1975). (Abbasi-Yadkori et al., 2011) improves the regret to $O(d\sqrt{T \log^2 T})$ with a self-normalized tail inequality, derived by the method of mixture (Victor et al., 2009). (Li et al., 2021) further proposes an algorithm with $O(d\sqrt{T \log T} \text{poly}(\log \log T))$ regret by bounding the supremum of self-normalized processes, which matches the lower bound up to a $\text{poly}(\log \log T)$ factor. The self-normalized tail inequalities for linear bandits in these works are all Hoeffding-type, i.e., only consider sub-Gaussian noises. However, for linear RL, Bernstein-type inequalities considering the sub-exponential noise, are necessary for sharper statistical results.

RL with Linear Function Approximation Recent works have focused on designing statistically and/or computationally efficient algorithms for RL with linear function approximation. The first sample efficient algorithm is introduced by (Jiang et al., 2017), where low Bellman rank is considered. Subsequent works on this setting include (Dann et al., 2018; Sun et al., 2019). (Yang & Wang, 2019) develops the first statistically and computationally efficient algorithm for linear MDPs with a simulator, where the transition probability and reward functions are linear concerning a feature mapping $\phi(s, a)$. Subsequently, (Jin et al., 2020) considers RL settings for linear MDPs and propose LSVI-UCB algorithm reaching $\tilde{O}(\sqrt{H^3 d^3 T})$ regret. Concurrently, (Zanette et al., 2020a) provides a Thompson sampling based algorithm with regret bound of $\tilde{O}(d^2 H^2 \sqrt{T})$. More works generalize linear MDPs includes (Zanette et al., 2020b) for low inherent bellman error, (Wang et al., 2020b) for linear Q function, and (Wang et al., 2020a) for bounded Eluder dimension.

Another popular linearly-parameterized MDP is the linear mixture MDP, where transition probability is linear to the feature function over (state, action, next state) triples. (Modi et al., 2020) firstly considers the statistical complexity of this setting and (Yang & Wang, 2020) provides an $\tilde{O}(H^2 d \log T \sqrt{T})$ regret with special case of low-dimensional representation of the transition matrix. Sub-

sequently, (Jia et al., 2020; Ayoub et al., 2020) proposes UCRL-VTR algorithm with $\tilde{O}(\sqrt{H^3 d^2 T})$ regret, and (Cai et al., 2020) considers adversarial rewards setting, giving same regret. Notably, the nearly minimax optimal regret for linear mixture MDP is first obtained by UCRL-VTR⁺ in (Zhou et al., 2021) under $d \geq H$ case.

3. Preliminaries

We consider episodic finite horizon MDP $\mathcal{M} = \{\mathcal{S}, \mathcal{A}, H, \{\mathbb{P}_h\}_h, \{r_h\}_h\}$, where \mathcal{S} is the state space, \mathcal{A} is the action space, $H \in \mathbb{Z}^+$ is the length of each episode, $\mathbb{P}_h : \mathcal{S} \times \mathcal{A} \rightarrow \Delta(\mathcal{S})$ and $r_h : \mathcal{S} \times \mathcal{A} \rightarrow [0, 1]$ are time-dependent transition probability and deterministic reward function. We assume that \mathcal{S} is a measurable space with possibly infinite number of elements and \mathcal{A} is a finite set.

For a time-inhomogeneous MDP, the policy is time-dependent, which is denoted as $\pi = \{\pi_1, \dots, \pi_H\}$. Here $\pi_h(s)$ is the action that agent takes at state s at the h -th step. The value function $V_h^\pi : \mathcal{S} \rightarrow \mathbb{R}$ is the expected value of cumulative rewards received under policy π when starting from a state at h -th step, given as

$$V_h^\pi(s) := \mathbb{E} \left[\sum_{h'=h}^H r_{h'}(s_{h'}, \pi_{h'}(s_{h'})) \mid s_h = s, \pi \right],$$

for any $s \in \mathcal{S}, h \in [H]$. The state-action function $Q_h^\pi : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$ gives the expected value of cumulative rewards starting from a state-action pair at h -th step, defined as

$$Q_h^\pi(s, a) = \mathbb{E} \left[\sum_{h'=h}^H r_{h'}(s_{h'}, a_{h'}) \mid s_h = s, a_h = a, \pi \right],$$

for any $(s, a) \in \mathcal{S} \times \mathcal{A}, h \in [H]$. For any function $V : \mathcal{S} \rightarrow \mathbb{R}$, we denote $\mathbb{P}_h V(s, a) = \mathbb{E}_{s' \sim \mathbb{P}_h(\cdot | s, a)} V(s')$ and $[\mathbb{V}_h V](s, a) = \mathbb{P}_h V^2(s, a) - [\mathbb{P}_h V(s, a)]^2$, where V^2 stands for the function whose value at s is $V^2(s)$. The Bellman equation associated with a policy π is

$$\begin{aligned} Q_h^\pi(s, a) &= r_h(s, a) + \mathbb{P}_h V_{h+1}^\pi(s, a) \\ V_h^\pi(s) &= Q_h^\pi(s, \pi_h(s)) \end{aligned}$$

for any $(s, a) \in \mathcal{S} \times \mathcal{A}, h \in [H]$. Since the action space and the episode length are both finite, there always exists an optimal policy π^* such that $V_h^*(s) = \sup_\pi V_h^\pi(s)$ for any $s \in \mathcal{S}, h \in [H]$, with Bellman optimality equation as

$$\begin{aligned} Q_h^*(s, a) &= r_h(s, a) + \mathbb{P}_h V_{h+1}^*(s, a) \\ V_h^*(s) &= \max_{a \in \mathcal{A}} Q_h^*(s, a) \end{aligned}$$

for any $(s, a) \in \mathcal{S} \times \mathcal{A}, h \in [H]$.

The structural assumption we make in this paper is a linear structure in both transition and reward, which has been considered in (Yang & Wang, 2019; Jin et al., 2020; Zanette et al., 2020a). The formal definition is as follows.

Definition 3.1 (Linear MDP). A MDP $\mathcal{M} = \{\mathcal{S}, \mathcal{A}, H, \{\mathbb{P}_h\}_h, \{r_h\}_h\}$ is a linear MDP with a known feature mapping $\phi : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}^d$, if for any $h \in [H]$, there exist $|\mathcal{S}|$ unknown d -dimensional measures $\mu_h = (\mu_h(1), \dots, \mu_h(|\mathcal{S}|)) \in \mathbb{R}^{d \times |\mathcal{S}|}$ and an unknown vector $\theta_h \in \mathbb{R}^d$, such that for any $(s, a) \in \mathcal{S} \times \mathcal{A}$, we have

$$\mathbb{P}_h(\cdot \mid s, a) = \langle \phi(s, a), \mu_h(\cdot) \rangle, \quad r_h(s, a) = \langle \phi(s, a), \theta_h \rangle.$$

We make the following assumptions, similar to existing literature (Jin et al., 2020; Agarwal et al., 2019). Specifically, for any $h \in [H]$, (i) $\sup_{s, a} \|\phi(s, a)\|_2 \leq 1$, (ii) $\|\mu_h v\|_2 \leq \sqrt{d}$ for any vector $v \in \mathbb{R}^{|\mathcal{S}|}$ with $\|v\|_\infty \leq 1$, (iii) $\|\theta_h\|_2 \leq W$, and (iv) $r_h(s, a) \in [0, 1]$ for all $(s, a) \in \mathcal{S} \times \mathcal{A}$.

In this paper, we focus on the setting where the reward function $\{r_h\}_{h \in [H]}$, i.e., $\{\theta_h\}_{h \in [H]}$ is known, but our algorithm can readily be extended to handle unknown rewards.

Learning Protocol In every episode k , the learner first proposes a policy π^k based on all history information up to the end of episode $k - 1$. The learner then executes π^k to generate a single trajectory $\tau^k = \{s_h^k, a_h^k\}_{h=1}^H$ with $a_h^k = \pi_h^k(s_h^k)$ and $s_{h+1}^k \sim \mathbb{P}_h(\cdot | s_h^k, a_h^k)$. The goal of the learner is to learn the optimal policy by interacting with the environment during K episodes. For the k -th episode, the initial state s_1^k is picked by the adversary and the optimal policy will minimize the cumulative regret over K episodes:

$$\text{Regret}(K) = \sum_{k=1}^K [V_1^*(s_1^k) - V_1^{\pi^k}(s_1^k)].$$

4. Strategic Exploration in Linear MDP

Section 4.1 illustrates the standard ways of strategic exploration in linear MDPs in existing works, i.e., optimistic value iteration with parameters estimated by linear ridge regression. Next, we point out in Section 4.2 barriers to minimax optimality in existing algorithms, which also helps explain our algorithm design Section 5.

4.1. Optimistic Learning in Linear MDPs

Optimistic learning evolves in an episodic fashion. In episode k , the agent first estimates unknown parameters of the linear MDP by historical data up to episodes $k - 1$. One standard approach is estimating the parameter $\omega_h^* = \theta_h + \mu_h V_{h+1}^*$ by linear ridge regression, as LSVI-UCB in (Jin et al., 2020) and its variants in (Wang et al., 2020b;a), since the optimal Q function $Q_h^*(s, a) = \langle \omega_h^*, \phi(s, a) \rangle$ according to Proposition 2.3 in (Jin et al., 2020). Subsequently, an optimistic Q function $Q_{k,h}$ in Eq. (2) is constructed with the learned parameter $\omega_{k,h}$ and the exploration bonus $\beta \|\phi(\cdot, \cdot)\|_{\Lambda_{k,h}^{-1}}$. The agent then follows a greedy policy π^k

of $Q_{k,h}$ to interact with the environment and repeat the above procedure in the next episode.

We illustrate two major steps of linear ridge regression and the construction of the optimistic Q function below.

Linear Ridge Regression To estimate the optimal ω_h^* , the following regularized least-squares problem is proposed:

$$\min_{\omega \in \mathbb{R}^d} \sum_{i=1}^{k-1} [V_{k,h+1}(s_{h+1}^i) - \omega^\top \phi(s_h^i, a_h^i)]^2 + \lambda \|\omega\|_2^2 \quad (1)$$

The closed-form solution to Eq. (1) is

$$\omega_{k,h} = \Lambda_{k,h}^{-1} \sum_{i=1}^{k-1} \phi(s_h^i, a_h^i)^\top V_{k,h+1}(s_{h+1}^i)$$

where $\Lambda_{k,h} = \sum_{i=1}^{k-1} \phi(s_h^i, a_h^i) \phi(s_h^i, a_h^i)^\top + \lambda \mathbf{I}$, and $V_{k,h}(\cdot)$ is given in Eq. (3).

Optimistic Estimator $\omega_{k,h}$ is then used to build an optimistic state-action function in Eq. (2) with exploration bonus $\beta \|\phi(\cdot, \cdot)\|_{\Lambda_{k,h}^{-1}}$ to encourage exploration, and optimistic value function is given in Eq. (3) as well. Notice that these two functions are built in a backwards fashion from stage H to 1, such that named as optimistic value iteration.

$$Q_{k,h}(\cdot, \cdot) = \langle \omega_{k,h}, \phi(\cdot, \cdot) \rangle + \beta \|\phi(\cdot, \cdot)\|_{\Lambda_{k,h}^{-1}} \quad (2)$$

$$V_{k,h}(\cdot) = \max_{a \in \mathcal{A}} Q_{k,h}(\cdot, a) \quad (3)$$

In particular, denote the optimistic confidence set $\mathcal{C}_{k,h} := \{\omega : \|\omega - \omega_{k,h}\|_{\Lambda_{k,h}} \leq \beta\}$ such that $Q_{k,h}(\cdot, \cdot) = \max_{\omega \in \mathcal{C}_{k,h}} \langle \omega, \phi(\cdot, \cdot) \rangle$. Notably, confidence set $\mathcal{C}_{k,h}$ is an ellipsoid centered at $\omega_{k,h}$, with shape parameter $\Lambda_{k,h}$ and radius β (usually named as the exploration radius). It can be proved that with high probability, $\omega_h^* \in \mathcal{C}_{k,h}$ by using self-normalized tail inequalities for vector-valued martingales, e.g., Theorem 1 in (Abbasi-Yadkori et al., 2011), used broadly in the analysis of linear bandits or RL with linear function approximation. Consequently, functions in Eq. (2), (3) obtains optimism in high probability.

4.2. Barriers to Minimax Optimality

The above optimistic learning based value iteration is a commonly adopted paradigm of RL with linear function approximation in existing works, e.g., (Jin et al., 2020; Wang et al., 2020b;a). However, the best-known regret upper bound for linear MDPs is $\tilde{O}(\sqrt{H^3 d^3 T})$ by LSVI-UCB algorithm in (Jin et al., 2020), while the best known lower bound is $\Omega(Hd\sqrt{T})$ according to (Zhou et al., 2021). As shown in Section 6, the lower bound is tight. We analyze where the \sqrt{Hd} gap comes from and then propose corresponding solutions, which immediately sheds light on designing the efficient LSVI-UCB⁺ algorithm in the next section.

4.2.1. OVERLY AGGRESSIVE EXPLORATION

The tradeoff between exploitation and exploration is a central task for RL algorithms, implemented by designing exploration bonuses in optimistic learning. The current \sqrt{H} gap stems from the overly aggressive exploration, which means that the current exploration radius $\beta = \tilde{O}(Hd)$ in existing works, e.g., (Jin et al., 2020; Wang et al., 2020b; Ayoub et al., 2020) is too large and leads to insufficient exploitation. The underlying reason remains that a bonus with $\tilde{O}(Hd)$ radius is intrinsically Hoeffding-type since it has the order of the magnitude of the considered martingale difference sequence (MDS). We prove that a Bernstein-type bonus, based on the variance of the MDS, combined with the Law of Total Variance (LTV) (Lattimore & Hutter, 2012), can reduce one \sqrt{H} factor of regrets in linear MDPs. The motivation for this improvement comes from prior works (Azar et al., 2017; Jin et al., 2018; Zanette & Brunskill, 2019) for tabular MDPs, which succeeded in achieving \sqrt{H} regret reduction by introducing a Bernstein-type bonus. For linear mixtures MDPs, UCRL-VTR⁺ in (Zhou et al., 2021) firstly introduces a Bernstein-type bonus and also achieves a \sqrt{H} regret reduction. However, a direct adaption of UCRL-VTR⁺ in linear MDPs will not improve the regret due to the extra cost of building a uniform convergence argument.

4.2.2. EXTRA UNIFORM CONVERGENCE COST

Introducing a ε -covering net is a common approach to build a uniform convergence argument over a function class. Many algorithms for RL with linear function approximation achieve polynomial sample complexity with this approach. However, this brings extra dependency on d in the regret, as presented in prior analysis, e.g., LSVI-UCB in (Jin et al., 2020) and its variants in (Wang et al., 2020b;a). Specifically, when bounding the deviation term $(\hat{\mathbb{P}} - \mathbb{P})\hat{V}$, the self-normalized tail inequality cannot be applied directly since \hat{V} is not well-measurable. Prior works fix a value function $V(\cdot) \in \hat{\mathcal{V}}$, where $\hat{\mathcal{V}}$ is the function class contains all possible \hat{V} , and build a uniform convergence argument by taking uniform bound over all functions in the ε -covering net $\hat{\mathcal{N}}_\varepsilon$ of $\hat{\mathcal{V}}$. In this way, a self-normalized bound concerning \hat{V} can be established (refer to proof of Lemma C.8 in Appendix for details). However, the covering number of $\hat{\mathcal{N}}_\varepsilon$ highly depends on the feature space dimension, resulting in extra dependency on d in the regret. We propose a novel technique of bounding the deviation term by dominant term $(\hat{\mathbb{P}} - \mathbb{P})V^*$ and the correction term $(\hat{\mathbb{P}} - \mathbb{P})(\hat{V} - V^*)$ separately, to remove the extra dependency on d . Note that bounding the correction term is also required for RL algorithms in tabular MDPs to achieve minimax optimality. However, adopting this idea to linear MDPs is nontrivial since we need to build the self-normalized bound for vector-valued martingale other than the well-studied scalar bound in tabular MDPs.

5. Optimal Exploration for linear MDPs

In this section, we present the proposed LSVI-UCB⁺ algorithm (Algorithm 1), where the optimistic value iteration is performed in lines 6-14, and the learned policy is executed in line 17. The remaining parts of Algorithm 1 are responsible for estimating parameter μ_h by linear weighted ridge regression. Specifically, the estimated variance is given in line 15, whose lower bound is controlled in lines 18-24, and the solution to the regression is given in line 25.

Algorithm 1 LSVI-UCB⁺ for Linear MDPs

Require: Regularization parameter $\lambda, \hat{\beta}, \check{\beta}$.

- 1: **for** step $h = H, \dots, 1$ **do**
- 2: $\hat{\Lambda}_{1,h}, \check{\Lambda}_{1,h} \leftarrow \lambda \mathbf{I}; \hat{\mu}_{1,h} \leftarrow \mathbf{0}$
- 3: **end for**
- 4: **for** episode $k = 1, \dots, K$ **do**
- 5: $\hat{V}_{k,H+1}(\cdot), \check{V}_{k,H+1}(\cdot) \leftarrow 0$
- 6: **for** step $h = H, \dots, 1$ // **Optimistic value iteration do**
- 7: $\hat{Q}_{k,h}(\cdot, \cdot) = r_h(\cdot, \cdot) + \langle \hat{\mu}_{k,h}, \hat{V}_{k,h+1}, \phi(\cdot, \cdot) \rangle + 3\hat{\beta} \|\phi(\cdot, \cdot)\|_{\hat{\Lambda}_{k,h}^{-1}}$ // **Over-optimistic Q function**
- 8: $\hat{Q}_{k,h}(\cdot, \cdot) = r_h(\cdot, \cdot) + \langle \hat{\mu}_{k,h}, \hat{V}_{k,h+1}, \phi(\cdot, \cdot) \rangle + \hat{\beta} \|\phi(\cdot, \cdot)\|_{\hat{\Lambda}_{k,h}^{-1}}$ // **Optimistic Q function**
- 9: $\check{Q}_{k,h}(\cdot, \cdot) = r_h(\cdot, \cdot) + \langle \hat{\mu}_{k,h}, \check{V}_{k,h+1}, \phi(\cdot, \cdot) \rangle - \check{\beta} \|\phi(\cdot, \cdot)\|_{\hat{\Lambda}_{k,h}^{-1}}$ // **Pessimistic Q function**
- 10: $\hat{V}_{k,h}(\cdot) \leftarrow \min \{ \max_{a \in \mathcal{A}} \hat{Q}_{k,h}(\cdot, a), H \}$
- 11: $\check{V}_{k,h}(\cdot) \leftarrow \min \{ \max_{a \in \mathcal{A}} \check{Q}_{k,h}(\cdot, a), H \}$
- 12: $\check{V}_{k,h}(\cdot) \leftarrow \max \{ \max_{a \in \mathcal{A}} \check{Q}_{k,h}(\cdot, a), 0 \}$
- 13: $\pi_h^k(\cdot) \leftarrow \arg \max_{a \in \mathcal{A}} \hat{Q}_{k,h}(\cdot, a)$
- 14: **end for**
- 15: Receive the initial state s_1^k .
- 16: **for** step $h = 1, \dots, H$ **do**
- 17: $a_h^k \leftarrow \pi_h^k(s_h^k)$, and observe $s_{h+1}^k \sim \mathbb{P}_h(\cdot | s_h^k, a_h^k)$.
- 18: $\hat{\sigma}_{k,h} \leftarrow \sqrt{\max\{H, dE_{k,h}, [\hat{V}_{k,h} \hat{V}_{k,h+1}](s_h^k, a_h^k) + U_{k,h}\}}$ according to Eq. (9) and Lemma 7.5.
- 19: $\hat{\Lambda}_{k+1,h} \leftarrow \hat{\Lambda}_{k,h} + \hat{\sigma}_{k,h}^{-2} \phi(s_h^k, a_h^k) \phi(s_h^k, a_h^k)^\top$
- 20: **if** $\|\hat{\sigma}_{k,h}^{-1} \phi(s_h^k, a_h^k)\|_{\hat{\Lambda}_{k,h}^{-1}} \leq 1/(Hd^3)$ **then**
- 21: $\varsigma_{k,h} \leftarrow \sqrt{H}$
- 22: **else**
- 23: $\varsigma_{k,h} \leftarrow H\sqrt{d^3}$ // **Enlarge ς_k if necessary**
- 24: **end if**
- 25: $\hat{\sigma}_{k,h} \leftarrow \sqrt{\max\{\varsigma_{k,h}^2, dE_{k,h}, [\hat{V}_{k,h} \hat{V}_{k,h+1}](s_h^k, a_h^k) + U_{k,h}\}}$
- 26: $\hat{\Lambda}_{k+1,h} \leftarrow \hat{\Lambda}_{k,h} + \hat{\sigma}_{k,h}^{-2} \phi(s_h^k, a_h^k) \phi(s_h^k, a_h^k)^\top$
- 27: $\hat{\mu}_{k+1,h} \leftarrow \hat{\Lambda}_{k+1,h}^{-1} \sum_{i=1}^k \hat{\sigma}_{i,h}^{-2} \phi(s_h^i, a_h^i) \delta(s_{h+1}^i)^\top$ // **Solution to weighted ridge regression**
- 28: **end for**
- 29: **end for**

LSVI-UCB⁺ is an optimistic algorithm similar to existing works (Yang & Wang, 2019; Jin et al., 2020; Ayoub et al., 2020), but upgrading the Hoeffding-type bonus to a carefully designed Bernstein-type one. The exploration radius in LSVI-UCB⁺ is proportional to the standard deviation of the optimal value function conditioned on some state-action pair, which accounts for two key novelties of LSVI-UCB⁺:

(i) We replace the linear ridge regression in prior works (Yang & Wang, 2019; Jin et al., 2020; Ayoub et al., 2020) with a carefully designed weighted version such that LTV can be applied. Note that the linear weighted ridge regression estimator was originally built for linear bandits with heteroscedastic noises, e.g., (Lattimore et al., 2015; Kirschner & Krause, 2018). Besides, the regression is performed to estimate μ_h , i.e., transition matrix, instead of estimating indirect variables, e.g., ω_h^* of LSVI-UCB in (Jin et al., 2020).

(ii) A variance estimator, based on the estimated parameter $\hat{\mu}_{k,h}$, is built for the optimal value function to determine the weights in regression. UCRL-VTR⁺ in (Zhou et al., 2021) also introduces weighted ridge regression for linear mixture MDPs, and weights are determined by variances of the constructed optimistic value function. However, the weights in LSVI-UCB⁺ are very different from those in UCRL-VTR⁺, since our variances are estimated with respect to the optimal value function, not the constructed value function.

5.1. Linear Weighted Ridge Regression

Denote $\delta(s) \in \mathbb{R}^{|\mathcal{S}|}$ as a one-hot vector that is zero everywhere except that the entry corresponding to state s is one, and define $\epsilon_h^k := \mathbb{P}_h(\cdot | s_h^k, a_h^k) - \delta(s_{h+1}^k)$. Since $\mathbb{E}[\epsilon_h^k | \mathcal{F}_{k,h}] = 0$, $\delta(s_{h+1}^k)$ is an unbiased estimate of $\mathbb{P}_h(\cdot | s_h^k, a_h^k) = \mu_h^\top \phi(s_h^k, a_h^k)$. Thus, μ_h can be learned via regression from $\phi(s_h^k, a_h^k)$ to $\delta(s_{h+1}^k)$. In addition, samples are normalized by the estimated standard deviation $\hat{\sigma}_{k,h}$. Thus, the estimated parameter $\hat{\mu}_{k,h}$ in line 27 of Algorithm 1 is the solution to the following weighted ridge regression problem:

$$\min_{\mu \in \mathbb{R}^{d \times |\mathcal{S}|}} \sum_{i=1}^{k-1} \left\| \left[\mu_h^\top \phi(s_h^i, a_h^i) - \delta(s_{h+1}^i) \right] \hat{\sigma}_{i,h}^{-1} \right\|_2^2 + \lambda \|\mu\|_F^2,$$

where $\|\cdot\|_F$ denotes Frobenius norm. The solution is

$$\hat{\mu}_{k,h} = \hat{\Lambda}_{k,h}^{-1} \sum_{i=1}^{k-1} \hat{\sigma}_{i,h}^{-2} \phi(s_h^i, a_h^i) \delta(s_{h+1}^i)^\top, \quad (4)$$

where $\hat{\Lambda}_{k,h} = \sum_{i=1}^{k-1} \hat{\sigma}_{i,h}^{-2} \phi(s_h^i, a_h^i) \phi(s_h^i, a_h^i)^\top + \lambda \mathbf{I}$. Thus, the estimated transition probability is denoted as

$$\hat{\mathbb{P}}_{k,h}(\cdot | s, a) = \hat{\mu}_{k,h}^\top \phi(s, a)$$

for any $(s, a) \in \mathcal{S} \times \mathcal{A}$. After estimating the transition matrix, line 8 in Algorithm 1 constructs an optimistic state-

action function, which is equivalent to

$$\hat{Q}_{k,h}(\cdot, \cdot) = \max_{\mu \in \hat{\mathcal{C}}_{k,h}} r_h(\cdot, \cdot) + \langle \mu, \hat{V}_{k,h+1}, \phi(\cdot, \cdot) \rangle, \quad (5)$$

where the optimistic confidence set is given by

$$\hat{\mathcal{C}}_{k,h} := \{ \mu : \|(\mu - \hat{\mu}_{k,h}) \hat{V}_{k,h+1}\|_{\hat{\Lambda}_{k,h}} \leq \hat{\beta} \}, \quad (6)$$

and $\hat{\beta}$ is the exploration radius. In addition, the pessimistic state-action function in line 9 is equivalent to

$$\check{Q}_{k,h}(\cdot, \cdot) = \min_{\mu \in \check{\mathcal{C}}_{k,h}} r_h(\cdot, \cdot) + \langle \mu, \check{V}_{k,h+1}, \phi(\cdot, \cdot) \rangle, \quad (7)$$

where the pessimistic confidence set is

$$\check{\mathcal{C}}_{k,h} := \{ \mu : \|(\mu - \hat{\mu}_{k,h}) \check{V}_{k,h+1}\|_{\check{\Lambda}_{k,h}} \leq \check{\beta} \}. \quad (8)$$

Monotonicity Apart from optimistic and pessimistic Q functions, we also construct a over-optimistic Q function $\hat{Q}_{k,h}(\cdot, \cdot)$. Subsequently, corresponding value functions can be proved to have

$$\check{V}_{k,h}(s) \leq V_h^*(s) \leq \hat{V}_{k,h}(s) \leq \min_{i \in [k]} \{\hat{V}_{i,h}(s)\}, \quad \forall s \in \mathcal{S}$$

in high probability. In particular, the pessimistic value function is required for estimating the variance upper bound later, while the over-optimistic value function is also necessary for bounding the correction term.

5.2. Variance Estimation

After estimating the transition matrix in Eq. (4), LSVI-UCB⁺ estimates the variance of the optimal value function $[\mathbb{V}_h V_{h+1}^*](s_h^k, a_h^k)$ and the variance of sub-optimality gap $[\mathbb{V}_h (\hat{V}_{k,h+1} - V_{h+1}^*)](s_h^k, a_h^k)$. This is a major difference with prior UCRL-VTR⁺ algorithm in (Zhou et al., 2021) for linear mixture MDPs, which only estimates the variance of the constructed optimistic value function $[\mathbb{V}_h \hat{V}_{k,h+1}](s_h^k, a_h^k)$. The purpose to estimate these two variances remains that we utilize Bernstein self-normalized tail inequality in Theorem 7.1 to bound the dominant term $[(\mathbb{P}_{k,h} - \mathbb{P}_h) V_{h+1}^*](s_h^k, a_h^k)$ and the correction term $[(\mathbb{P}_{k,h} - \mathbb{P}_h)(\hat{V}_{k,h+1} - V_{h+1}^*)](s_h^k, a_h^k)$ separately to remove the extra dependency of regrets, such that we need to estimate these two variance, which are illusated below.

Variance of Optimal Value Function We first consider the case where the transition matrix and optimal value function $V_{h+1}^*(\cdot)$ were given. In this case, the variance of the optimal value function is given by

$$[\mathbb{V}_h V_{h+1}^*](s_h^k, a_h^k) = \mathbb{P}_h V_{h+1}^{*2}(s_h^k, a_h^k) - [\mathbb{P}_h V_{h+1}^*](s_h^k, a_h^k)^2$$

However, only empirical estimation $\hat{\mu}_{k,h}$ and optimistic value function $\hat{V}_{k,h}$ are obtainable, which means we only have the empirical variance of the optimistic value function:

$$[\hat{\mathbb{V}}_{k,h} \hat{V}_{k,h+1}](s_h^k, a_h^k) = \hat{\mathbb{P}}_{k,h} \hat{V}_{k,h+1}^2(s_h^k, a_h^k)_{[0,H^2]} - [\hat{\mathbb{P}}_{k,h} \hat{V}_{k,h+1}](s_h^k, a_h^k)_{[0,H]}^2. \quad (9)$$

To ensure the accuracy of the estimation, we introduce an offset term $U_{k,h}$ to guarantee that $|\mathbb{V}_h V_{h+1}^*](s_h^k, a_h^k) - [\hat{\mathbb{V}}_{k,h} \hat{V}_{k,h+1}](s_h^k, a_h^k)| \leq U_{k,h}$ with high probability. Moreover, the exact form of offset term $U_{k,h}$ is specified in Lemma 7.5, which requires accessing the pessimistic value functions as detailed in Lemma C.13 in Appendix.

Variance of Sub-optimality Gap In particular, we try to build a upper bound for the variance of the sub-optimality gap, which is given as

$$\begin{aligned} & [\mathbb{V}_h (\hat{V}_{k,h+1} - V_{h+1}^*)](s_h^k, a_h^k) \\ &= [\mathbb{P}_h (\hat{V}_{k,h+1} - V_{h+1}^*)^2](s_h^k, a_h^k) - [\mathbb{P}_h (\hat{V}_{k,h+1} - V_{h+1}^*)](s_h^k, a_h^k)^2 \\ &\leq [\mathbb{P}_h (\hat{V}_{k,h+1} - V_{h+1}^*)^2](s_h^k, a_h^k) \leq H [\mathbb{P}_h (\hat{V}_{k,h+1} - V_{h+1}^*)](s_h^k, a_h^k) \\ &\leq H [\mathbb{P}_h (\hat{V}_{k,h+1} - \check{V}_{k,h+1})](s_h^k, a_h^k) \end{aligned}$$

where the second and last inequalities holds by the optimism and pessimism of $\hat{V}_{k,h+1}$ and $\check{V}_{k,h+1}$, respectively. Thus, it suffices to upper bound the deviation $[\mathbb{P}_h (\hat{V}_{k,h+1} - \check{V}_{k,h+1})](s_h^k, a_h^k)$. In addition, the upper bound of the variance of the sub-optimality gap is denoted as $E_{k,h}$, specified in Lemma 7.5.

Putting two variances together, the weight $\hat{\sigma}_{k,h}$ in Algorithm 1 is given by

$$\hat{\sigma}_{k,h} = \sqrt{\max\{\varsigma_{k,h}^2, dE_{k,h}, [\hat{\mathbb{V}}_{k,h} \hat{V}_{k,h+1}](s_h^k, a_h^k) + U_{k,h}\}},$$

which is the maximum over the weight lower bound $\varsigma_{k,h}$, the variance upper bound of the optimal value function, the variance of the sub-optimality gap with a factor d scaling. Here $\varsigma_{k,h}$ controls the lower bound of $\hat{\sigma}_{k,h}$ and is dynamically determined in lines 18-24. In particular, we try to keep the magnitude of the considered MDS to be small by adaptively enlarging $\varsigma_{k,h}$, which is detailed in Remark 7.4.

6. Main Results

This section presents the results of the statistical, space, and computational complexities of the LSVI-UCB⁺ algorithm. In particular, LSVI-UCB⁺ firstly reaches nearly minimax optimal regret in linear MDPs, while the space and computational complexities are no worse than prior works.

6.1. Statistical Complexity

We first present the regret upper bound of LSVI-UCB⁺ in Theorem 6.1.

Theorem 6.1 (Regret Upper Bound). *Set $\lambda = 1/(H^2\sqrt{d})$. Then, with probability at least $1 - 10\delta$, the regret of LSVI-UCB⁺ is upper bounded by*

$$\text{Regret}(K) = \tilde{O}\left(Hd\sqrt{T} + H^3d^6 + \sqrt{H^7d^7}\right), \quad (10)$$

where $T = KH$.

Proof Sketch. We prove the result conditioning on the conclusion of Lemma 7.5. Initially, with the standard regret decomposition, we can show that the total regret is bounded by the summation of the exploration bonus, i.e.,

$$\begin{aligned} \text{Regret}(K) &\leq \sum_{k=1}^K \sum_{h=1}^H \hat{\beta} \|\phi(s_h^k, a_h^k)\|_{\hat{\Lambda}_{k,h}^{-1}} \\ &= \sum_{k=1}^K \sum_{h=1}^H \hat{\beta} \hat{\sigma}_{k,h}^{-1} \|\hat{\sigma}_{k,h}^{-1} \phi(s_h^k, a_h^k)\|_{\hat{\Lambda}_{k,h}^{-1}} \\ &\leq \hat{\beta} \underbrace{\sqrt{\sum_{k=1}^K \sum_{h=1}^H \hat{\sigma}_{k,h}^2}}_{\tilde{O}(\sqrt{HT+c\sqrt{T}})} \underbrace{\sqrt{\sum_{k=1}^K \sum_{h=1}^H \|\hat{\sigma}_{k,h}^{-1} \phi(s_h^k, a_h^k)\|_{\hat{\Lambda}_{k,h}^{-1}}^2}}_{\tilde{O}(\sqrt{Hd})} \end{aligned} \quad (11)$$

where the second inequality holds by Cauchy-Schwarz inequality and c is a constant. The summation of $\|\hat{\sigma}_{k,h}^{-1} \phi(s_h^k, a_h^k)\|_{\hat{\Lambda}_{k,h}^{-1}}^2$ can be addressed by Elliptical Potential Lemma (Lemma F.5 in Appendix), and the summation of $\hat{\sigma}_{k,h}^2$ can be bounded by

$$\begin{aligned} \sum_{k=1}^K \sum_{h=1}^H \hat{\sigma}_{k,h}^2 &\leq \sum_{k=1}^K \sum_{h=1}^H \varsigma_{k,h}^2 + \sum_{k=1}^K \sum_{h=1}^H [dE_{k,h} + U_{k,h}] \\ &+ \sum_{k=1}^K \sum_{h=1}^H [\hat{V}_{k,h} \hat{V}_{k,h+1}](s_h^k, a_h^k) \leq \tilde{O}\left(HT + c\sqrt{T}\right), \end{aligned}$$

where the first inequality holds by definition of $\hat{\sigma}_{k,h}$, and the second inequality holds by $\sum_{k=1}^K \sum_{h=1}^H \varsigma_{k,h}^2 \leq \tilde{O}(HT)$ due to the conservatism of elliptical potentials, $\sum_{k=1}^K \sum_{h=1}^H [dE_{k,h} + U_{k,h}] \leq \tilde{O}(c\sqrt{T})$ due to the Elliptical Potential Lemma, and $\sum_{k=1}^K \sum_{h=1}^H [\hat{V}_{k,h} \hat{V}_{k,h+1}](s_h^k, a_h^k) \leq \tilde{O}(HT)$ due to the LTV. Besides, the exploration radius $\hat{\beta} = \tilde{O}(\sqrt{d})$, which determined by the upper bound of $(\mathbb{P}_{k,h} - \mathbb{P}_h)\hat{V}_{k,h+1}(s_h^k, a_h^k)$, detailed in Section 7.3. The full proof is given in Appendix D. \square

Theorem 6.1 is proved under the event that the optimistic confidence set $\hat{\mathcal{C}}_{k,h}$ holds, which is built in Lemma 7.5. In addition, we find that the exploration radius $\hat{\beta}$ of the optimistic confidence set $\hat{\mathcal{C}}_{k,h}$ determines the sharpness of the final regret, as shown in Eq. (11).

Remark 6.2. When² $T \geq \max\{H^4d^{10}, H^5d^5\}$, the regret in Eq. (10) can be simplified to $\tilde{O}(Hd\sqrt{T})$, which improves

²Large-sample regime conditions are required in many RL algorithms to obtain satisfactory statistical complexities, e.g. UCRL-VTR⁺ in (Zhou et al., 2021) requires $T \geq d^4H^2 + d^3H^3$.

the $\tilde{O}\left(\sqrt{H^3d^3T}\right)$ regret of LSVI-UCB (Jin et al., 2020) by a factor of \sqrt{Hd} . Moreover, our algorithm design an analytical tools including Theorem 7.1 and Lemma 7.3 in next sections can further improve the regret bound of UCRL-VTR⁺ in (Zhou et al., 2021) for linear mixture MDPs to $\tilde{O}(Hd\sqrt{T})$ from existing $\tilde{O}(\sqrt{H^2d^2T + H^3d^3T})$, such that it is minimax optimal up to logarithmic factor without large dimension assumption that $d \geq H$ in (Zhou et al., 2021).

Lower Bound We formalize a linear MDP instance in Appendix E to establish an $\Omega(Hd\sqrt{T})$ regret lower bound of linear MDPs. This linear MDP instance is firstly proposed in Remark 23 in (Zhou et al., 2021), which shares the same regret lower bound of a linear mixture MDP instance. This class of MDP is hard due to the intrinsic sparsity of reward and indistinguishability of large action space, which can be regarded as an extension of hard instances in linear bandits literature (Dani et al., 2008; Lattimore & Szepesvári, 2020). According to Theorem 8 in (Zhou et al., 2021), linear mixture MDPs have regret lower bound of $\Omega(Hd\sqrt{T})$. Thus linear MDPs have the same regret lower bound. The lower bound, together with the upper bound of LSVI-UCB⁺ in Theorem 6.1 show that LSVI-UCB⁺ is minimax optimal up to logarithmic factors when $T \geq \max\{H^4d^{10}, H^5d^5\}$.

6.2. Space and Computational Complexities

As stated above, LSVI-UCB⁺ reaches minimax optimal regret up to logarithmic factors, which is also computationally efficient. In particular, the space and computational complexities of LSVI-UCB⁺ are briefly stated below, which are both the same as LSVI-UCB in (Jin et al., 2020).

Space Complexity Although $\hat{\mu}_{k,h} \in \mathbb{R}^{d \times |\mathcal{S}|}$ and $|\mathcal{S}|$ can be infinitely large, we do not store it explicitly, as Algorithm 1 only utilizes indirect variables $\hat{\mu}_{k,h} \mathbf{V} = \hat{\Lambda}_{k,h}^{-1} \sum_{i=1}^{k-1} \hat{\sigma}_{i,h}^{-2} \phi(s_h^i, a_h^i) V(s_{h+1}^i)$ where $V \in \{\hat{V}_{k,h+1}, \check{V}_{k,h+1}, \hat{V}_{k,h+1}^2, \check{V}_{k,h+1}^2\}$. In episode $k \in [K]$, Algorithm 1 only stores $\hat{\mu}_{k,h} \mathbf{V}$, $\hat{\Lambda}_{k,h}$, $\tilde{\Lambda}_{k,h}$, $\hat{\sigma}_{k,h}$, $\check{\sigma}_{k,h}$ for all $h \in [H]$, and $\{\phi(s_h^{k'}, a)\}_{a \in \mathcal{A}}$ for all $(k', h) \in [k] \times [H]$, which means LSVI-UCB⁺ requires $O(d^2H + d|\mathcal{A}|T)$ space.

Computational Complexity Assume $\hat{\mu}_{k,h} \mathbf{V}$ is given for some $V \in \{\hat{V}_{k,h+1}, \check{V}_{k,h+1}, \hat{V}_{k,h+1}^2, \check{V}_{k,h+1}^2\}$, then each evaluation of $V(s)$ takes $O(d^2|\mathcal{A}|)$ operations. Thus, calculating $\hat{\mu}_{k,h} \mathbf{V}$ takes $O(d^2|\mathcal{A}|)K$ operations. Besides, $\hat{\Lambda}_{k,h}$, $\tilde{\Lambda}_{k,h}$ can be computed by Sherman-Morrison formula (Hager, 1989) with $O(d^2)$ operations and other steps take less operations. Thus, LSVI-UCB⁺ has a running time of $O(d^2|\mathcal{A}|KT)$, which is computationally efficient since its running time is polynomial on $d, K, H, |\mathcal{A}|$, and does not depend on $|\mathcal{S}|$, which can be possibly infinite.

7. Mechanism Towards Minimax Optimality

In this section, we highlight our technical contributions in building the sharp optimistic confidence set $\hat{C}_{k,h}$. We first present two novel analytical tools, a sharp Bernstein self-normalized tail inequality for vector-valued martingales in Section 7.1, the conservatism of elliptical potentials in Section 7.2. Together, these two analytical tools remove the additional dependency of regret on H in the regret of the LSVI-UCB⁺ algorithm. In addition, we also upper bound the correction term of the form $(\hat{\mathbb{P}} - \mathbb{P})(\hat{V} - V^*)$ to avoid extra cost from the covering net such that the additional dependency of regret on d is removed as well. Consequently, the sharp confidence set $\hat{C}_{k,h}$ is built in Lemma 7.5 in Section 7.3. These technical contributions together enable LSVI-UCB⁺ to achieve nearly minimax optimal regret and have the potential to improve other statistical results of algorithms for RL with linear function approximation.

7.1. Sharp Bernstein Self-normalized Bound

Most existing self-normalized concentrations used in prior works for RL with linear function approximation (Jin et al., 2020; Wang et al., 2020b;a; Ayoub et al., 2020) are all Hoeffding-type, i.e., they consider sub-Gaussian noises. Our self-normalized bound below considers sub-exponential noises, which is a Bernstein-type one.

Theorem 7.1 (Bernstein self-normalized bound). *Let $\{\mathcal{G}_t\}_{t=1}^\infty$ be a filtration, and $\{\mathbf{x}_t, \eta_t\}_{t \geq 1}$ be a stochastic process such that $\mathbf{x}_t \in \mathbb{R}^d$ is \mathcal{G}_t -measurable and $\eta_t \in \mathbb{R}$ is \mathcal{G}_{t+1} -measurable. Define $\mathbf{Z}_t = \lambda \mathbf{I} + \sum_{i=1}^t \mathbf{x}_i \mathbf{x}_i^\top$ for $t \geq 1$ and $\mathbf{Z}_0 = \lambda \mathbf{I}$. If $\|\mathbf{x}_t\|_2 \leq L$, and η_t satisfies $\mathbb{E}[\eta_t | \mathcal{G}_t] = 0$, $\mathbb{E}[\eta_t^2 | \mathcal{G}_t] \leq \sigma^2$, and $|\eta_t \cdot \min\{1, \|\mathbf{x}_t\|_{\mathbf{Z}_{t-1}^{-1}}\}| \leq R$ for all $t \geq 1$. Then, for any $0 < \delta < 1$, with probability at least $1 - \delta$, we have:*

$$\forall t > 0, \left\| \sum_{i=1}^t \mathbf{x}_i \eta_i \right\|_{\mathbf{Z}_t^{-1}} \leq \tilde{O}(\sigma \sqrt{d} + R).$$

Proof. Please refer to Appendix B. \square

Remark 7.2. The proof of Theorem 7.1 in Appendix B shows that bounding the self-normalized vector-valued martingales is equivalent to bounding a scalar-valued MDS $\{\eta_t \cdot \min\{1, \|\mathbf{x}_t\|_{\mathbf{Z}_{t-1}^{-1}}\}, \mathcal{G}_{t+1}\}$, where η_t is scaled by the factor of $\min\{1, \|\mathbf{x}_t\|_{\mathbf{Z}_{t-1}^{-1}}\}$. In particular, $\|\mathbf{x}_t\|_{\mathbf{Z}_{t-1}^{-1}} = \sqrt{\mathbf{x}_t^\top \mathbf{Z}_{t-1}^{-1} \mathbf{x}_t}$ is denoted as the elliptical potential, which is common in online learning literature (Cesa-Bianchi & Lugosi, 2006). Notice that Elliptical Potential Lemma shows that $\min\{1, \|\mathbf{x}_t\|_{\mathbf{Z}_{t-1}^{-1}}\}$ can be roughly regarded as an attenuated sequence. Theorem 7.1 looks similar to but is sharper than Theorem 2 in (Zhou et al., 2021), because it pay extra attentions on elliptical potentials $\|\mathbf{x}_t\|_{\mathbf{Z}_{t-1}^{-1}}$. However, the

scaling factor $\min\{1, \|\mathbf{x}_t\|_{\mathbf{Z}_{t-1}^{-1}}\}$ is crudely deflated to 1 in Theorem 2 in (Zhou et al., 2021), such that the attenuation of the MDS is neglected, which is highlighted in Lemma D.8.

7.2. Conservatism of Elliptical Potentials

Notice the self-normalized bound in Theorem 7.1 will determine the order of exploration radius $\hat{\beta}$. We try to keep the second term R , the magnitude of the MDS, in Theorem 7.1 smaller than the first $\sigma \sqrt{d}$ by utilizing the conservatism of elliptical potentials. Specifically, the following lemma characterizes the conservatism of elliptical potentials, i.e., elliptical potentials are usually small. This lemma is firstly proposed at Exercise 19.3 in (Lattimore & Szepesvári, 2020) for case $c = 1$, and we generalize it to case $c > 0$.

Lemma 7.3 (Elliptical Potentials are Usually Small). *Given $\lambda > 0$ and sequence $\{\mathbf{x}_t\}_{t=1}^T \subset \mathbb{R}^d$ with $\|\mathbf{x}_t\|_2 \leq L$ for all $t \in [T]$, define $\mathbf{Z}_t = \lambda \mathbf{I} + \sum_{i=1}^t \mathbf{x}_i \mathbf{x}_i^\top$ for $t \geq 1$ and $\mathbf{Z}_0 = \lambda \mathbf{I}$. The number of times $\|\mathbf{x}_t\|_{\mathbf{Z}_{t-1}^{-1}} \geq c$ is at most*

$$\frac{3d}{\log(1+c^2)} \log \left(1 + \frac{L^2}{\lambda \log(1+c^2)} \right)$$

for any $t \in [T]$, where $c > 0$ is a constant.

Proof. Please refer to Lemma D.8 in Appendix. \square

On the one hand, for some stage $h \in [H]$, the noise $\eta_k = \hat{\sigma}_{k,h}^{-1} \mathbf{V}^\top \epsilon_h^k$ for some value function V as detailed in Appendix C. On the other hand, in Theorem 7.1, R is the absolute bound of $|\eta_k \cdot \min\{1, \|\hat{\sigma}_{k,h}^{-1} \phi(s_h^k, a_h^k)\|_{\hat{\Lambda}_{k,h}^{-1}}\}|$. Lemma 7.3 reveals that R is intrinsically small since the elliptical potential $\|\hat{\sigma}_{k,h}^{-1} \phi(s_h^k, a_h^k)\|_{\hat{\Lambda}_{k,h}^{-1}}^2$ is small in most episodes. In addition, we only need to enlarge the lower bound of $\hat{\sigma}_{k,h}$, i.e., $\varsigma_{k,h}$, when $\|\hat{\sigma}_{k,h}^{-1} \phi(s_h^k, a_h^k)\|_{\hat{\Lambda}_{k,h}^{-1}}^2$ is large such that R can remain small uniformly, which is detailed in the following remark.

Remark 7.4. Lines 18-25 of Algorithm 1 ensure the following facts for any $(k, h) \in [K] \times [H]$ by introducing indicator variable $\|\tilde{\sigma}_{i,h}^{-1} \phi(s_h^i, a_h^i)\|_{\tilde{\Lambda}_{i,h}^{-1}}$:

(i) In most cases, we have $\|\tilde{\sigma}_{i,h}^{-1} \phi(s_h^i, a_h^i)\|_{\tilde{\Lambda}_{i,h}^{-1}} \leq 1/(Hd^3)$,

then $\varsigma_{k,h} = \sqrt{H}$ such that $\hat{\sigma}_{i,h} = \tilde{\sigma}_{i,h}$. We can prove that the elliptical potential $\|\hat{\sigma}_{k,h}^{-1} \phi(s_h^i, a_h^i)\|_{\hat{\Lambda}_{k,h}^{-1}}$ is small;

(ii) Otherwise, $\varsigma_{k,h} = H\sqrt{d^3}$, such that $\hat{\sigma}_{k,h} \geq H\sqrt{d^3}$. In this case, the R is still small since $\hat{\sigma}_{k,h}$ is large.

Notice that always enlarging $\varsigma_{k,h}$ for any $k \in [K]$ is a simple method to keep R small, but contributing to final regret linearly since it is an additive term in $\sum_{k=1}^K \sum_{h=1}^H \hat{\sigma}_{k,h}$. Nevertheless, the enlarging operation in case (ii) only contributes

an additive constant term to the regret, since the elliptical potential $\|\tilde{\sigma}_{k,h}^{-1} \phi(s_h^k, a_h^k)\|_{\hat{\Lambda}_{k,h}^{-1}}^2$ is small in most episodes such that the enlarging operation in case (ii) happens rarely.

As a consequence, the R in Theorem 7.1 can be controlled to be smaller than the $\sigma\sqrt{d}$ in the LSVI-UCB⁺ algorithm, such that the exploration radius $\hat{\beta}$ in LSVI-UCB⁺ is $\tilde{O}(\sqrt{d})$. However, analog Bernstein self-normalized bounds, such as Theorem 2 in (Zhou et al., 2021) and Theorem 1 in (Faury et al., 2020), cannot lead to such exploration radius in the SVI-UCB⁺ algorithm, while that of Theorem 2 in (Zhou et al., 2021) is $\tilde{O}(\sqrt{d} + \sqrt{H}d^2)$, and Theorem 1 in (Faury et al., 2020) is $\tilde{O}(\sqrt{H}\sqrt{d^5})$.

7.3. Building Confidence Set with Correction Term

This subsection explains critical steps of building a sharp optimistic confidence set $\hat{\mathcal{C}}_{k,h}$ with the correction term. Specifically, the exploration bonus $\hat{\beta}\|\phi(s_h^k, a_h^k)\|_{\hat{\Lambda}_{k,h}^{-1}}$ is the upper bound of the deviation term $[(\hat{\mathbb{P}}_{k,h} - \mathbb{P}_h)\hat{V}_{k,h+1}](s_h^k, a_h^k)$, which is the basis of optimistic learning. Specifically, the deviation term can be decomposed by triangle inequality as the sum of dominant term and correction term:

$$\begin{aligned} & \underbrace{[(\hat{\mathbb{P}}_{k,h} - \mathbb{P}_h)V_{h+1}^*](s_h^k, a_h^k)}_{\text{Dominant term with respect to } \hat{V}_{h+1}^*} \\ & + \underbrace{[(\hat{\mathbb{P}}_{k,h} - \mathbb{P}_h)(\hat{V}_{k,h+1} - V_{h+1}^*)](s_h^k, a_h^k)}_{\text{Correction Term}} \\ & \leq \|(\boldsymbol{\mu} - \hat{\boldsymbol{\mu}}_{k,h})\mathbf{V}_{h+1}^*\|_{\hat{\Lambda}_{k,h}} \|\phi(s_h^k, a_h^k)\|_{\hat{\Lambda}_{k,h}^{-1}} \\ & \quad + \|(\boldsymbol{\mu} - \hat{\boldsymbol{\mu}}_{k,h})(\hat{V}_{k,h+1} - V_{h+1}^*)\|_{\hat{\Lambda}_{k,h}} \|\phi(s_h^k, a_h^k)\|_{\hat{\Lambda}_{k,h}^{-1}} \\ & \leq \hat{\beta}^{(1)} \|\phi(s_h^k, a_h^k)\|_{\hat{\Lambda}_{k,h}^{-1}} + \hat{\beta}^{(2)} \|\phi(s_h^k, a_h^k)\|_{\hat{\Lambda}_{k,h}^{-1}} \\ & = \hat{\beta} \|\phi(s_h^k, a_h^k)\|_{\hat{\Lambda}_{k,h}^{-1}}, \end{aligned}$$

where the first inequality holds by Cauchy-Schwarz inequality, the second inequality holds since

$$\begin{aligned} \hat{\mathcal{C}}_{k,h}^{(1)} &= \{\boldsymbol{\mu} : \|(\boldsymbol{\mu} - \hat{\boldsymbol{\mu}}_{k,h})\mathbf{V}_{h+1}^*\|_{\hat{\Lambda}_{k,h}} \leq \hat{\beta}^{(1)}\} \\ \hat{\mathcal{C}}_{k,h}^{(2)} &= \{\boldsymbol{\mu} : \|(\boldsymbol{\mu} - \hat{\boldsymbol{\mu}}_{k,h})(\hat{V}_{k,h+1} - V_{h+1}^*)\|_{\hat{\Lambda}_{k,h}} \leq \hat{\beta}^{(2)}\}, \end{aligned}$$

and the last equality holds since $\hat{\beta} = \hat{\beta}^{(1)} + \hat{\beta}^{(2)}$. In the following, we briefly illustrate how to use Theorem 7.1 to build confidence sets $\hat{\mathcal{C}}_{k,h}^{(1)}$ and $\hat{\mathcal{C}}_{k,h}^{(2)}$. Initially, by

$$\|(\hat{\boldsymbol{\mu}}_{k,h} - \boldsymbol{\mu}_h)\mathbf{V}\|_{\hat{\Lambda}_{k,h}} \lesssim \left\| \sum_{i=1}^{k-1} \hat{\sigma}_{i,h}^{-2} \phi(s_h^i, a_h^i) \boldsymbol{\epsilon}_h^{i\top} \mathbf{V} \right\|_{\hat{\Lambda}_{k,h}^{-1}} \quad (12)$$

for some fixed function V in Lemma F.7, building a confidence set with respect to V is equivalent to building a self-normalized bound for $\left\| \sum_{i=1}^{k-1} \hat{\sigma}_{i,h}^{-2} \phi(s_h^i, a_h^i) \boldsymbol{\epsilon}_h^{i\top} \mathbf{V} \right\|_{\hat{\Lambda}_{k,h}^{-1}}$.

Building $\hat{\mathcal{C}}_{k,h}^{(1)}$: We build confidence set $\hat{\mathcal{C}}_{k,h}^{(1)}$ by applying the Bernstein self-normalized inequality in Theorem 7.1

with dynamic control of MDS magnitude, highlighted in Remark 7.2. Thus, $U_{k,h}$ is specified in Lemma 7.5 to guarantee that $\hat{\sigma}_{k,h}$ upper bounds $\mathbb{V}_h V_{h+1}^*(s_h^k, a_h^k)$, and $\varsigma_{k,h}$ is set dynamically. Besides, the uniform convergence argument by covering net in (Jin et al., 2020; Wang et al., 2020b;a) is not required, since now $V = V_{h+1}^*$ in Eq. (12) is a fixed function and there is no measurability issue. Consequently, we get $\hat{\beta}_k^{(1)} = \tilde{O}(\sqrt{d})$, which is detailed in Lemma C.16.

Building $\hat{\mathcal{C}}_{k,h}^{(2)}$: We apply also Theorem 7.1 with dynamic control of MDS magnitude to build $\hat{\mathcal{C}}_{k,h}^{(2)}$ as well. Similarly, $E_{k,h}$ is specified in Lemma 7.5 to guarantee that $\hat{\sigma}_{k,h}$ upper bounds $[\mathbb{V}_h(\hat{V}_{k,h+1} - V_{h+1}^*)](s_h^k, a_h^k)$, and $\varsigma_{k,h}$ is set dynamically to keep the MDS magnitude R in Theorem 7.1 small. Since now $V = \hat{V}_{k,h+1} - V_{h+1}^*$ in Eq. (12) suffers from the measurability issue, a uniform convergence argument by covering net is still required, which bring extra dependency on d in the exploration radius $\hat{\beta}^{(2)}$. That is why we enlarge $E_{k,h}$ with a d factor in $\hat{\sigma}_{k,h}$.

Putting everything together gives the following key technical lemma that builds the sharp optimistic confidence set $\hat{\mathcal{C}}_{k,h}$.

Lemma 7.5. *Set $\hat{\beta} = \hat{\beta}^{(1)} + \hat{\beta}^{(2)}$, then there exists an absolute constant $c > 0$ such that for any $\delta \in (0, 1)$, with probability at least $1 - 7\delta$, we have that simultaneously for any $k \in [K]$ and any $h \in [H]$,*

$$\boldsymbol{\mu}_h \in \hat{\mathcal{C}}_{k,h} \cap \check{\mathcal{C}}_{k,h},$$

and $|\left[\hat{\mathbb{V}}_{k,h} \hat{V}_{k,h+1}\right](s_h^k, a_h^k) - [\mathbb{V}_h V_{h+1}^*](s_h^k, a_h^k)| \leq U_{k,h}$, $[\mathbb{V}_h(\hat{V}_{k,h+1} - V_{h+1}^*)](s_h^k, a_h^k) \leq E_{k,h}$ where $\hat{\beta}^{(1)}$, $\hat{\beta}^{(2)}$, $U_{k,h}$, $E_{k,h}$ are specified in Lemma C.18 in Appendix.

Proof. Please refer to Appendix C.4. \square

8. Conclusion

This paper presents a computationally and statistically efficient algorithm, LSVI-UCB⁺, which builds on linear weighted ridge regression and upper confidence value iteration with a Bernstein-type exploration bonus. LSVI-UCB⁺ is the first algorithm to reach minimax optimal regret bound up to logarithmic factors for linear MDPs. Our sharp result builds on a novel Bernstein self-normalized bound with the conservatism of elliptical potentials, and refined analysis of the correction term, which serve as new analytical tools for RL with linear function approximation.

Acknowledgements

This work is supported by the Technology and Innovation Major Project of the Ministry of Science and Technology of China under Grant 2020AAA0108400 and 2020AAA0108403, and the Tsinghua University Initiative Scientific Research Program.

References

- Abbasi-Yadkori, Y., Pál, D., and Szepesvári, C. Improved algorithms for linear stochastic bandits. *Advances in neural information processing systems*, 24:2312–2320, 2011.
- Agarwal, A., Jiang, N., Kakade, S. M., and Sun, W. Reinforcement learning: Theory and algorithms. *CS Dept., UW Seattle, Seattle, WA, USA, Tech. Rep.*, 2019.
- Ayoub, A., Jia, Z., Szepesvari, C., Wang, M., and Yang, L. Model-based reinforcement learning with value-targeted regression. In *International Conference on Machine Learning*, pp. 463–474. PMLR, 2020.
- Azar, M. G., Munos, R., and Kappen, H. J. Minimax pac bounds on the sample complexity of reinforcement learning with a generative model. *Machine learning*, 91(3): 325–349, 2013.
- Azar, M. G., Osband, I., and Munos, R. Minimax regret bounds for reinforcement learning. In *International Conference on Machine Learning*, pp. 263–272. PMLR, 2017.
- Cai, Q., Yang, Z., Jin, C., and Wang, Z. Provably efficient exploration in policy optimization. In *International Conference on Machine Learning*, pp. 1283–1294. PMLR, 2020.
- Cesa-Bianchi, N. and Lugosi, G. *Prediction, learning, and games*. Cambridge university press, 2006.
- Dani, V., Hayes, T. P., and Kakade, S. M. Stochastic linear optimization under bandit feedback. In *Conference on Learning Theory-colt*, 2008.
- Dann, C., Jiang, N., Krishnamurthy, A., Agarwal, A., Langford, J., and Schapire, R. E. On oracle-efficient pac rl with rich observations. In *Proceedings of the 32nd International Conference on Neural Information Processing Systems*, pp. 1429–1439, 2018.
- Fauray, L., Abeille, M., Calauzènes, C., and Fercoq, O. Improved optimistic algorithms for logistic bandits. In *International Conference on Machine Learning*, pp. 3052–3060. PMLR, 2020.
- Freedman, D. A. On tail probabilities for martingales. *the Annals of Probability*, pp. 100–118, 1975.
- Hager, W. W. Updating the inverse of a matrix. *SIAM review*, 31(2):221–239, 1989.
- He, J., Zhou, D., and Gu, Q. Minimax optimal reinforcement learning for discounted mdps. In *Proceedings of the 32nd International Conference on Neural Information Processing Systems*, 2021.
- Jia, Z., Yang, L., Szepesvari, C., and Wang, M. Model-based reinforcement learning with value-targeted regression. In *Learning for Dynamics and Control*, pp. 666–686. PMLR, 2020.
- Jiang, N., Krishnamurthy, A., Agarwal, A., Langford, J., and Schapire, R. E. Contextual decision processes with low bellman rank are pac-learnable. In *International Conference on Machine Learning*, pp. 1704–1713. PMLR, 2017.
- Jin, C., Allen-Zhu, Z., Bubeck, S., and Jordan, M. I. Is q-learning provably efficient? In *Proceedings of the 32nd International Conference on Neural Information Processing Systems*, pp. 4868–4878, 2018.
- Jin, C., Yang, Z., Wang, Z., and Jordan, M. I. Provably efficient reinforcement learning with linear function approximation. In *Conference on Learning Theory*, pp. 2137–2143. PMLR, 2020.
- Kirschner, J. and Krause, A. Information directed sampling and bandits with heteroscedastic noise. In *Conference On Learning Theory*, pp. 358–384. PMLR, 2018.
- Lattimore, T. and Hutter, M. Pac bounds for discounted mdps. In *International Conference on Algorithmic Learning Theory*, pp. 320–334. Springer, 2012.
- Lattimore, T. and Szepesvári, C. *Bandit algorithms*. Cambridge University Press, 2020.
- Lattimore, T., Crammer, K., and Szepesvári, C. Linear multi-resource allocation with semi-bandit feedback. In *NIPS*, pp. 964–972, 2015.
- Li, Y., Wang, Y., Chen, X., and Zhou, Y. Tight regret bounds for infinite-armed linear contextual bandits. In *International Conference on Artificial Intelligence and Statistics*, pp. 370–378. PMLR, 2021.
- Modi, A., Jiang, N., Tewari, A., and Singh, S. Sample complexity of reinforcement learning using linearly combined model ensembles. In *International Conference on Artificial Intelligence and Statistics*, pp. 2010–2020. PMLR, 2020.
- Sun, W., Jiang, N., Krishnamurthy, A., Agarwal, A., and Langford, J. Model-based rl in contextual decision processes: Pac bounds and exponential improvements over model-free approaches. In *Conference on learning theory*, pp. 2898–2933. PMLR, 2019.
- Tossou, A., Basu, D., and Dimitrakakis, C. Near-optimal optimistic reinforcement learning using empirical bernstein inequalities. *arXiv preprint arXiv:1905.12425*, 2019.

- Victor, H., la Peña, D., Lai, T. L., and Shao, Q.-M. *Self-normalized processes: Limit theory and Statistical Applications*, volume 204. Springer, 2009.
- Wang, R., Salakhutdinov, R. R., and Yang, L. Reinforcement learning with general value function approximation: Provably efficient approach via bounded eluder dimension. *Advances in Neural Information Processing Systems*, 33, 2020a.
- Wang, Y., Wang, R., Du, S. S., and Krishnamurthy, A. Optimism in reinforcement learning with generalized linear function approximation. In *International Conference on Learning Representations*, 2020b.
- Yang, L. and Wang, M. Sample-optimal parametric q-learning using linearly additive features. In *International Conference on Machine Learning*, pp. 6995–7004. PMLR, 2019.
- Yang, L. and Wang, M. Reinforcement learning in feature space: Matrix bandit, kernels, and regret bound. In *International Conference on Machine Learning*, pp. 10746–10756. PMLR, 2020.
- Zanette, A. and Brunskill, E. Tighter problem-dependent regret bounds in reinforcement learning without domain knowledge using value function bounds. In *International Conference on Machine Learning*, pp. 7304–7312. PMLR, 2019.
- Zanette, A., Brandfonbrener, D., Brunskill, E., Pirotta, M., and Lazaric, A. Frequentist regret bounds for randomized least-squares value iteration. In *International Conference on Artificial Intelligence and Statistics*, pp. 1954–1964. PMLR, 2020a.
- Zanette, A., Lazaric, A., Kochenderfer, M., and Brunskill, E. Learning near optimal policies with low inherent bellman error. In *International Conference on Machine Learning*, pp. 10978–10989. PMLR, 2020b.
- Zhou, D., Gu, Q., and Szepesvari, C. Nearly minimax optimal reinforcement learning for linear mixture markov decision processes. In *Conference on Learning Theory*, pp. 4532–4576. PMLR, 2021.

Appendices

A	Additional Comparisons of Related Works	13
B	Sharp Bernstein Self-Normalized Bound	13
B.1	Proof of Theorem 7.1	17
C	High Probability Events	17
C.1	Independent Confidence Sets	18
C.2	Variance Upper Bound	22
C.3	Dependent Confidence Sets	25
C.4	Proof of Lemma 7.5	30
D	Regret Upper Bound	31
D.1	Monotonicity	32
D.2	Suboptimality Gap	34
D.3	Summation of Estimated Variances	38
D.4	Proof of Theorem 6.1	43
E	Lower Bound	44
F	Auxiliary Lemmas	45
F.1	Concentration Inequality	46
F.2	Elliptical Potentials	46
F.3	Linear MDP Property	47
F.4	Covering Net	47

In the appendix, we present some additional results and supporting materials to supplement the statements, theorems and proofs in the main papers. There are 6 sections in appendix:

- Appendix A presents additional comparisons with related works.
- Appendix B presents the proof of our proposed sharp Bernstein tail inequality for self-normalized vector-valued martingales.
- Appendix C presents the construction of several high probability confidences sets.
- Appendix D presents the former proof of the regret upper bound of LSVI-UCB⁺ algorithm, with conclusion given in Theorem 6.1.
- Appendix E constructs a hard-to-learn MDP to build a regret lower bound for linear MDPs.
- Appendix F presents auxiliary lemmas necessary for proofs in above sections and important properties that will be helpful in algorithm design.

A. Additional Comparisons of Related Works

Table 2 serves as a more complete table compared to Table 1 in the main paper, which lists some representative works in RL with linear function approximation. The top part of Table 2 lists representative works for the linear MDP and its generalizations and the bottom part is for the linear mixture MDP and its generalizations.

Table 2. Theoretical results on RL with linear function approximation, where † denotes that rewards are adversarial, and d_h is the dimension of the feature mapping at the h -th stage within the episodes and K is the number of episodes.

Setting	Algorithm	Technique	Regret
Linear MDP	OPT-RLSVI (Zanette et al., 2020a)	Hoeffding+Covering	$\tilde{O}(H^2 d^2 \sqrt{T})$
Linear MDP	LSVI-UCB (Jin et al., 2020)	Hoeffding+Covering	$\tilde{O}(\sqrt{H^3 d^3 T})$
Linear MDP	LSVI-UCB ⁺ (this paper)	Bernstein+Covering	$\tilde{O}(H d \sqrt{T})$
Linear Q Function	LSVI-UCB* (Wang et al., 2020b)	Hoeffding+Covering	$\tilde{O}(\sqrt{d^3 T})$
Low Bellman Error	ELEANOR (Zanette et al., 2020b)	Hoeffding+Covering	$\tilde{O}(\sum_{h=1}^H d_h \sqrt{K})$
Bounded Eluder Dimension	\mathcal{F} -LSVI(δ) (Wang et al., 2020a)	Hoeffding+Covering	$\tilde{O}(\text{poly}(dH) \sqrt{T})$
Linear Mixture MDP	UCRL-VTR (Jia et al., 2020; Ayoub et al., 2020)	Hoeffding	$\tilde{O}(\sqrt{H^3 d^2 T})$
Linear Mixture MDP	UCRL-VTR ⁺ (Zhou et al., 2021)	Bernstein	$\tilde{O}(\sqrt{H^2 d^2 T + H^3 d T})$
Feature Space	MatrixRL (Yang & Wang, 2020)	Hoeffding	$\tilde{O}(H^2 d \log T \sqrt{T})$
Linear Mixture MDP [†]	OPPO (Cai et al., 2020)	Hoeffding	$\tilde{O}(\sqrt{d^2 H^3 T})$

From Table 2, we can find that existing algorithms for the linear MDP and its generalizations all use a classical Hoeffding self-normalized bound such as Theorem 1 in (Abbasi-Yadkori et al., 2011) with the covering net argument, while our work is the first to introduce a Bernstein self-normalized bound with a covering net argument. Moreover, building a covering net argument in our work does not bring extra dependency on feature space dimension d since we only consider covering net argument in bounding the correction term which can be made small.

As for the linear mixture MDP and its generalizations, the covering net argument is not required due to the structure of the linear mixture MDP. In addition, prior works (Yang & Wang, 2020; Jia et al., 2020; Ayoub et al., 2020; Cai et al., 2020) utilize Hoeffding self-normalized bound to build confidence sets, while (Zhou et al., 2021) considers the Bernstein self-normalized bound for the first time in the setting of linear mixture MDP. Compared with regret bound of $\tilde{O}(\sqrt{H^3 d^2 T})$ obtained in (Jia et al., 2020; Ayoub et al., 2020) for linear mixture MDPs, the regret bound $\tilde{O}(\sqrt{H^2 d^2 T + H^3 d T})$ in (Zhou et al., 2021) is better and a \sqrt{H} factor is further saved if $d \geq H$.

B. Sharp Bernstein Self-Normalized Bound

In this section, we prove the proposed sharp Bernstein tail inequality for self-normalized vector-valued martingales. Our proof diagram is based on the proof of Theorem 1 in (Zhou et al., 2021), which is firstly proposed in the proof of Lemma 14 in (Dani et al., 2008). However, our Bernstein self-normalized bound is sharper than Theorem 1 in (Zhou et al., 2021) with critical changes of the attenuation of the martingale difference sequence.

Specifically, $\{\min\{1, \|\mathbf{x}_t\|_{\mathbf{Z}_{t-1}^{-1}}\}\}_{t \in [T]}$ can be roughly considered as an attenuated sequence since we can prove $\sum_{t=1}^T \min\{1, \|\mathbf{x}_t\|_{\mathbf{Z}_{t-1}^{-1}}\} = O(d \log T)$ by Elliptical Potential Lemma. On the contrary, $\min\{1, \|\mathbf{x}_t\|_{\mathbf{Z}_{t-1}^{-1}}\}$ is deflated to 1 in Theorem 1, (Zhou et al., 2021) such that the bound is looser than ours. In the following proof, we do not deflate $\min\{1, \|\mathbf{x}_t\|_{\mathbf{Z}_{t-1}^{-1}}\}$ to 1, and we take into account the elliptical potential $\|\mathbf{x}_t\|_{\mathbf{Z}_{t-1}^{-1}}$ in our algorithm design, which is one of the major contributions in this paper.

Firstly, we give the following definitions to simplifying notations during the proof.

Definition B.1.

$$\mathbf{d}_t := \sum_{i=1}^t \mathbf{x}_i \eta_i,$$

$$\mathbf{Z}_t := \|\mathbf{d}_t\|_{\mathbf{Z}_t^{-1}},$$

$$\begin{aligned}
 w_t &:= \|\mathbf{x}_t\|_{\mathbf{Z}_{t-1}^{-1}}, \\
 \beta_t &:= 8\sigma\sqrt{d\log(1+tL^2/(d\lambda))\log(4t^2/\delta)} + 4R\log(4t^2/\delta), \\
 \mathcal{E}_t &:= \mathbb{1}\{0 \leq s \leq t, Z_s \leq \beta_s\},
 \end{aligned}$$

for $t \geq 1$ and $\mathbf{d}_0 = 0, Z_0 = 0, \beta_0 = 0$ for $t = 0$.

Measurability With the assumptions in Theorem 7.1, x_t is \mathcal{G}_t -measurable and η_t is \mathcal{G}_{t+1} -measurable. Thus, w_t is \mathcal{G}_t -measurable, and d_t, Z_t and \mathcal{E}_t are \mathcal{G}_{t+1} -measurable.

Our goal is to upper bound Z_t . By definition of Z_t , we have

$$\begin{aligned}
 Z_t^2 &= (\mathbf{d}_{t-1} + \mathbf{x}_t\eta_t)^\top \mathbf{Z}_t^{-1} (\mathbf{d}_{t-1} + \mathbf{x}_t\eta_t) \\
 &= \mathbf{d}_{t-1}^\top \mathbf{Z}_t^{-1} \mathbf{d}_{t-1} + 2\eta_t \mathbf{x}_t^\top \mathbf{Z}_t^{-1} \mathbf{d}_{t-1} + \eta_t^2 \mathbf{x}_t^\top \mathbf{Z}_t^{-1} \mathbf{x}_t \\
 &\leq Z_{t-1}^2 + \underbrace{2\eta_t \mathbf{x}_t^\top \mathbf{Z}_t^{-1} \mathbf{d}_{t-1}}_{I_1} + \underbrace{\eta_t^2 \mathbf{x}_t^\top \mathbf{Z}_t^{-1} \mathbf{x}_t}_{I_2}
 \end{aligned}$$

where the inequality holds since $\mathbf{Z}_t \geq \mathbf{Z}_{t-1}$.

Since $\mathbf{Z}_t = \mathbf{Z}_{t-1} + \mathbf{x}_t \mathbf{x}_t^\top$, by the Sherman–Morrison formula (Hager, 1989), we obtain

$$\mathbf{Z}_t^{-1} = \mathbf{Z}_{t-1}^{-1} - \frac{\mathbf{Z}_{t-1}^{-1} \mathbf{x}_t \mathbf{x}_t^\top \mathbf{Z}_{t-1}^{-1}}{1 + w_t^2}, \quad \text{for } t \geq 1.$$

Subsequently,

$$\begin{aligned}
 I_1 &= 2\eta_t \left(\mathbf{x}_t^\top \mathbf{Z}_{t-1}^{-1} \mathbf{d}_{t-1} - \frac{\mathbf{x}_t^\top \mathbf{Z}_{t-1}^{-1} \mathbf{x}_t \mathbf{x}_t^\top \mathbf{Z}_{t-1}^{-1} \mathbf{d}_{t-1}}{1 + w_t^2} \right) = 2\eta_t \left(\mathbf{x}_t^\top \mathbf{Z}_{t-1}^{-1} \mathbf{d}_{t-1} - \frac{w_t^2 \mathbf{x}_t^\top \mathbf{Z}_{t-1}^{-1} \mathbf{d}_{t-1}}{1 + w_t^2} \right) = \frac{2\eta_t \mathbf{x}_t^\top \mathbf{Z}_{t-1}^{-1} \mathbf{d}_{t-1}}{1 + w_t^2} \\
 I_2 &= \eta_t^2 \left(\mathbf{x}_t^\top \mathbf{Z}_{t-1}^{-1} \mathbf{x}_t - \frac{\mathbf{x}_t^\top \mathbf{Z}_{t-1}^{-1} \mathbf{x}_t \mathbf{x}_t^\top \mathbf{Z}_{t-1}^{-1} \mathbf{x}_t}{1 + w_t^2} \right) = \eta_t^2 \left(w_t^2 - \frac{w_t^4}{1 + w_t^2} \right) = \frac{\eta_t^2 w_t^2}{1 + w_t^2}
 \end{aligned}$$

Therefore, we have

$$Z_t^2 \leq \sum_{i=1}^t \frac{2\eta_i \mathbf{x}_i^\top \mathbf{Z}_{i-1}^{-1} \mathbf{d}_{i-1}}{1 + w_i^2} + \sum_{i=1}^t \frac{\eta_i^2 w_i^2}{1 + w_i^2} \tag{13}$$

Now we try to bound the two summation terms on the r.h.s. of Eq. (13) in Lemma B.3 and Lemma B.4, respectively. Before that, we present a uniform Bernstein bound required for proving Lemma B.3 and Lemma B.4.

Lemma B.2 (Uniform Bernstein Bound). *Let $\{x_i, \mathcal{F}_i\}$ be a martingale difference sequence with $\forall i \geq 1, \mathbb{E}(x_i | \mathcal{F}_{i-1}) = 0, \mathbb{E}(x_i^2 | \mathcal{F}_{i-1}) = \sigma_i^2, V_i^2 = \sum_{j=1}^i \sigma_j^2$. Furthermore, assume that $\mathbb{P}(|x_i| \leq c | \mathcal{F}_{i-1}) = 1$ for any $0 < c < \infty$.*

Then, for any $\delta > 0$, with probability at least $1 - \delta$, simultaneously for any $t \geq 1$, it holds that

$$\sum_{i=1}^t d_i \leq \sqrt{2V_t^2 \log(2t^2/\delta)} + \frac{2c \log(2t^2/\delta)}{3}.$$

Proof. By Freedman's inequality in Lemma F.2, for any $\delta > 0$ and some $t \geq 1$, with probability at least $1 - \delta/(2t^2)$, we have:

$$\sum_{i=1}^t d_i \leq \sqrt{2V_t^2 \log(1/\delta)} + \frac{2}{3}c \log(1/\delta). \tag{14}$$

Taking a union bound for Eq. (14) from $t = 1$ to ∞ and using the fact that $\sum_{t=1}^{\infty} t^{-2} < 2$ complete the proof. \square

Next, we bound the first summation term on the r.h.s. of Eq. (13) in Lemma B.3.

Lemma B.3. *Under assumptions in Theorem 7.1 and Definition B.1, with probability at least $1 - \delta/2$, simultaneously for all $t \geq 1$ it holds that*

$$\sum_{i=1}^t \frac{2\eta_i \mathbf{x}_i^\top \mathbf{Z}_{i-1}^{-1} \mathbf{d}_{i-1}}{1 + w_i^2} \mathcal{E}_{i-1} \leq 3\beta_t^2/4$$

Proof. Define

$$\ell_i = \frac{2\eta_i \mathbf{x}_i^\top \mathbf{Z}_{i-1}^{-1} \mathbf{d}_{i-1}}{1 + w_i^2} \mathcal{E}_{i-1},$$

and we will give a uniform upper bound of $\sum_{i=1}^t \ell_i$ by Lemma B.2 below.

Firstly, for any $1 \leq i \leq t$,

$$\mathbb{E}[\ell_i \mid \mathcal{G}_i] = \frac{2\mathbf{x}_i^\top \mathbf{Z}_{i-1}^{-1} \mathbf{d}_{i-1}}{1 + w_i^2} \mathcal{E}_{i-1} \mathbb{E}[\eta_i \mid \mathcal{G}_i] = 0.$$

Besides, we have

$$\left| \frac{2\mathbf{x}_i^\top \mathbf{Z}_{i-1}^{-1} \mathbf{d}_{i-1}}{1 + w_i^2} \mathcal{E}_{i-1} \right| \leq \frac{2 \|\mathbf{x}_i\|_{\mathbf{Z}_{i-1}^{-1}} \left[\|\mathbf{d}_{i-1}\|_{\mathbf{Z}_{i-1}^{-1}} |\mathcal{E}_{i-1}| \right]}{1 + w_i^2} \leq \frac{2w_i \beta_{i-1}}{1 + w_i^2} \leq 2 \min\{1, w_i\} \beta_{i-1} \quad (15)$$

where the first inequality holds due to Cauchy-Schwarz inequality, the second inequality holds due to the definition of \mathcal{E}_{i-1} , and the last inequality holds by algebra. Thus,

$$|\ell_i| = |2\eta_i \cdot \min\{1, w_i\} \beta_i| \leq 2R\beta_i \leq 2R\beta_t$$

where the last inequality holds since $\{\beta_i\}_{i \in [t]}$ is an increasing sequence.

Secondly, we also have

$$\begin{aligned} \sum_{i=1}^t \mathbb{E}[\ell_i^2 \mid \mathcal{G}_i] &\leq \sigma^2 \sum_{i=1}^t \left(\frac{2\mathbf{x}_i^\top \mathbf{Z}_{i-1}^{-1} \mathbf{d}_{i-1}}{1 + w_i^2} \mathcal{E}_{i-1} \right)^2 \\ &\leq \sigma^2 \sum_{i=1}^t [2 \min\{1, w_i\} \beta_{i-1}]^2 \\ &\leq 4\sigma^2 \beta_t^2 \sum_{i=1}^t \min\{1, w_i^2\} \\ &\leq 8\sigma^2 \beta_t^2 d \log(1 + tL^2/(d\lambda)) \end{aligned}$$

where the first inequality holds since $\mathbb{E}[\eta_i^2 \mid \mathcal{G}_i] \leq \sigma^2$, the second inequality holds due to Eq. (15), the third inequality holds since $\{\beta_i\}_{i \in [t]}$ is an increasing sequence, and the last inequality holds due to Lemma F.5.

Therefore, using Lemma B.2, simultaneously for all $t \geq 1$, with probability at least $1 - \delta/(4t^2)$, it holds that

$$\begin{aligned} \sum_{i=1}^t \ell_i &\leq \sqrt{16\sigma^2 \beta_t^2 d \log(1 + tL^2/(d\lambda)) \log(4t^2/\delta)} + 2/3 \cdot 2R\beta_t \log(4t^2/\delta) \\ &\leq \frac{\beta_t^2}{4} + 16\sigma^2 d \log(1 + tL^2/(d\lambda)) \log(4t^2/\delta) + \frac{\beta_t^2}{4} + 4R^2 \log^2(4t^2/\delta) \\ &\leq \beta_t^2/2 + \frac{1}{4} \left(8\sigma \sqrt{d \log(1 + tL^2/(d\lambda)) \log(4t^2/\delta)} + 4R \log(4t^2/\delta) \right)^2 \\ &= 3\beta_t^2/4 \end{aligned}$$

where the first inequality holds due to Lemma B.2, the second inequality holds due to $2\sqrt{|ab|} \leq |a| + |b|$, and the last equality follows from the definition of β_t . \square

Accordingly, the second summation terms on the r.h.s. of Eq. (13) is bounded in Lemma B.4.

Lemma B.4. *Under assumptions in Theorem 7.1 and Definition B.1, with probability at least $1 - \delta/2$, simultaneously for all $t \geq 1$ it holds that*

$$\sum_{i=1}^t \frac{\eta_i^2 w_i^2}{1 + w_i^2} \leq \beta_t^2/4$$

Proof. Define

$$\ell_i = \frac{\eta_i^2 w_i^2}{1 + w_i^2} - \mathbb{E} \left[\frac{\eta_i^2 w_i^2}{1 + w_i^2} \mid \mathcal{G}_i \right].$$

We will give a uniform upper bound of $\sum_{i=1}^t \ell_i$ by Lemma B.2, similar to Lemma B.3.

Clearly, for any $1 \leq i \leq t$, we have $\mathbb{E}[\ell_i \mid \mathcal{G}_i] = 0$. We further have that

$$\begin{aligned} \sum_{i=1}^t \mathbb{E}[\ell_i^2 \mid \mathcal{G}_i] &\leq \sum_{i=1}^t \mathbb{E} \left[\frac{\eta_i^4 w_i^4}{(1 + w_i^2)^2} \mid \mathcal{G}_i \right] \\ &\leq R^2 \sum_{i=1}^t \mathbb{E} \left[\frac{\eta_i^2 w_i^2}{1 + w_i^2} \mid \mathcal{G}_i \right] \\ &\leq R^2 \sigma^2 \sum_{i=1}^t \frac{w_i^2}{1 + w_i^2} \\ &\leq 2R^2 \sigma^2 d \log(1 + tL^2/(d\lambda)) \end{aligned}$$

where the first inequality holds due to the fact $\mathbb{E}(X - \mathbb{E}X)^2 \leq \mathbb{E}X^2$, the second inequality holds since $|\eta_i \cdot \min\{1, w_i\}| \leq R$, the third inequality holds since $\mathbb{E}[\eta_i^2 \mid \mathcal{G}_i] \leq \sigma^2$, and the fourth inequality holds due to the fact $w_i^2/(1 + w_i^2) \leq \min\{1, w_i^2\}$ and Lemma F.5.

Furthermore, using the fact that $|\eta_i \cdot \min\{1, w_i\}| \leq R$ holds almost surely under filtration \mathcal{G}_i , we obtain

$$\left| \frac{\eta_i^2 w_i^2}{1 + w_i^2} \right| \leq |\eta_i \cdot \min\{1, w_i^2\}| = |\eta_i \cdot \min\{1, w_i\}|^2 = R^2,$$

and

$$|\ell_i| \leq \left| \frac{\eta_i^2 w_i^2}{1 + w_i^2} \right| + \left| \mathbb{E} \left[\frac{\eta_i^2 w_i^2}{1 + w_i^2} \mid \mathcal{G}_i \right] \right| \leq 2R^2.$$

Therefore, by Lemma B.2, simultaneously for all $t \geq 1$, with probability at least $1 - \delta/(4t^2)$, it holds that

$$\begin{aligned} \sum_{i=1}^t \frac{\eta_i^2 w_i^2}{1 + w_i^2} &\leq \sum_{i=1}^t \mathbb{E} \left[\frac{\eta_i^2 w_i^2}{1 + w_i^2} \mid \mathcal{G}_i \right] + \sqrt{4R^2 \sigma^2 d \log(1 + tL^2/(d\lambda)) \log(4t^2/\delta)} + 4/3 \cdot R^2 \log(4t^2/\delta) \\ &\leq \sigma^2 \sum_{i=1}^t \frac{w_i^2}{1 + w_i^2} + 2R\sigma \sqrt{d \log(1 + tL^2/(d\lambda)) \log(4t^2/\delta)} + 2R^2 \log(4t^2/\delta) \\ &\leq 2\sigma^2 d \log(1 + tL^2/(d\lambda)) + 2R\sigma \sqrt{d \log(1 + tL^2/(d\lambda)) \log(4t^2/\delta)} + 2R^2 \log(4t^2/\delta) \\ &\leq 1/4 \cdot \left(8\sigma \sqrt{d} \sqrt{\log(1 + tL^2/(d\lambda)) \log(4t^2/\delta)} + 4R \log(4t^2/\delta) \right)^2 \\ &= \beta_t^2/4 \end{aligned}$$

where the first inequality holds due to Lemma B.2, the second inequality holds due to $\mathbb{E}[\eta_i^2 \mid \mathcal{G}_i] \leq \sigma^2$, the third inequality holds due to the fact $w_i^2/(1 + w_i^2) \leq \min\{1, w_i^2\}$ and Lemma 12, and the last inequality holds due to the definition of β_t . \square

B.1. Proof of Theorem 7.1

Proof of Theorem 7.1. Consider the case when conclusions of Lemma B.3 and Lemma B.4 hold. Conditioning on this event, we claim $Z_t \leq \beta_t$ for any $t \geq 0$.

We prove this by induction on t . Initially, the base case of $t = 0$ holds since $\beta_0 = 0 = Z_0$ by definition. Now fix some $t \geq 1$ and assume that for all $0 \leq i \leq t - 1$, we have $Z_i \leq \beta_i$. This implies that $\mathcal{E}_1 = \mathcal{E}_2 = \dots = \mathcal{E}_{t-1} = 1$. Then, by Eq. (13), we have

$$Z_t^2 \leq \sum_{i=1}^t \frac{2\eta_i \mathbf{x}_i^\top \mathbf{Z}_{i-1}^{-1} \mathbf{d}_{i-1}}{1 + w_i^2} + \sum_{i=1}^t \frac{\eta_i^2 w_i^2}{1 + w_i^2} = \sum_{i=1}^t \frac{2\eta_i \mathbf{x}_i^\top \mathbf{Z}_{i-1}^{-1} \mathbf{d}_{i-1}}{1 + w_i^2} \mathcal{E}_{i-1} + \sum_{i=1}^t \frac{\eta_i^2 w_i^2}{1 + w_i^2}. \quad (16)$$

Since the conclusions of Lemma B.3 and Lemma B.4 hold, we have

$$\sum_{i=1}^t \frac{2\eta_i \mathbf{x}_i^\top \mathbf{Z}_{i-1}^{-1} \mathbf{d}_{i-1}}{1 + w_i^2} \mathcal{E}_{i-1} \leq 3\beta_t^2/4, \quad (17)$$

$$\sum_{i=1}^t \frac{\eta_i^2 w_i^2}{1 + w_i^2} \leq \beta_t^2/4. \quad (18)$$

Therefore, substituting Eq. (17) and (18) into Eq. (16), we have $Z_t \leq \beta_t$, which ends the induction. Taking the union bound of the events in Lemma B.3 and Lemma B.4 implies that with probability at least $1 - \delta$, $Z_t \leq \beta_t$ holds for any $t \geq 1$. \square

C. High Probability Events

In this section, we define some high probability events, i.e., confidence sets concerning the parameter $\boldsymbol{\mu}_h$, and show how to build them. The goal of this section is to build the sharp optimistic confidence set $\hat{\mathcal{C}}_{k,h}$ in Lemma 7.5 for all $(k, h) \in [K] \times [H]$.

We lists all confidence sets encountered during the proof in the following. Confidence sets $\bar{\mathcal{C}}_{k,h}, \dot{\mathcal{C}}_{k,h}, \tilde{\mathcal{C}}_{k,h}, \check{\mathcal{C}}_{k,h}$ are called independent confidence sets, since they can be built by applying self-normalized concentration inequality directly without conditioning on other events. Instead, confidence sets $\hat{\mathcal{C}}_{k,h}^{(1)}, \hat{\mathcal{C}}_{k,h}^{(2)}, \hat{\mathcal{C}}_{k,h}$ are called dependent confidence sets since they can only be built by conditioning on other confidence sets, apart from self-normalized concentration inequalities.

Definition C.1 (Confidence Set). • Independent Confidence Sets:

$$\begin{aligned} \bar{\mathcal{C}}_{k,h} &= \left\{ \boldsymbol{\mu} : \left\| (\boldsymbol{\mu} - \hat{\boldsymbol{\mu}}_{k,h}) \hat{\mathbf{V}}_{k,h+1} \right\|_{\hat{\boldsymbol{\Lambda}}_{k,h}} \leq \bar{\beta} \right\} \\ \dot{\mathcal{C}}_{k,h} &= \left\{ \boldsymbol{\mu} : \left\| (\boldsymbol{\mu} - \hat{\boldsymbol{\mu}}_{k,h}) \dot{\mathbf{V}}_{k,h+1} \right\|_{\hat{\boldsymbol{\Lambda}}_{k,h}} \leq \dot{\beta} \right\} \\ \tilde{\mathcal{C}}_{k,h} &= \left\{ \boldsymbol{\mu} : \left\| (\boldsymbol{\mu} - \hat{\boldsymbol{\mu}}_{k,h}) \hat{\mathbf{V}}_{k,h+1}^2 \right\|_{\hat{\boldsymbol{\Lambda}}_{k,h}} \leq \tilde{\beta} \right\} \\ \check{\mathcal{C}}_{k,h} &= \left\{ \boldsymbol{\mu} : \left\| (\boldsymbol{\mu} - \hat{\boldsymbol{\mu}}_{k,h}) \check{\mathbf{V}}_{k,h+1} \right\|_{\hat{\boldsymbol{\Lambda}}_{k,h}} \leq \check{\beta} \right\} \end{aligned}$$

• Dependent Confidence Sets:

$$\begin{aligned} \hat{\mathcal{C}}_{k,h}^{(1)} &= \left\{ \boldsymbol{\mu} : \left\| (\boldsymbol{\mu} - \hat{\boldsymbol{\mu}}_{k,h}) \mathbf{V}_{h+1}^* \right\|_{\hat{\boldsymbol{\Lambda}}_{k,h}} \leq \hat{\beta}^{(1)} \right\} \\ \hat{\mathcal{C}}_{k,h}^{(2)} &= \left\{ \boldsymbol{\mu} : \left\| (\boldsymbol{\mu} - \hat{\boldsymbol{\mu}}_{k,h}) \left(\hat{\mathbf{V}}_{k,h+1} - \mathbf{V}_{h+1}^* \right) \right\|_{\hat{\boldsymbol{\Lambda}}_{k,h}} \leq \hat{\beta}^{(2)} \right\} \\ \hat{\mathcal{C}}_{k,h} &= \left\{ \boldsymbol{\mu} : \left\| (\boldsymbol{\mu} - \hat{\boldsymbol{\mu}}_{k,h}) \hat{\mathbf{V}}_{k,h+1} \right\|_{\hat{\boldsymbol{\Lambda}}_{k,h}} \leq \hat{\beta}^{(1)} + \hat{\beta}^{(2)} = \hat{\beta} \right\} \end{aligned}$$

To simplify notations during the proof, we further define the following events that optimistic and pessimistic confidence sets hold in multiple stages under some episode $k \in [K]$ or all episodes.

Definition C.2 (Optimism Event).

$$\begin{aligned}\widehat{\Psi}_{k,h} &:= \left\{ \forall h \leq h' \leq H : \boldsymbol{\mu}_{h'} \in \widehat{\mathcal{C}}_{k,h'} \right\} \\ \widehat{\Psi}_h &:= \left\{ \forall i \in [K], \forall h \leq h' \leq H : \boldsymbol{\mu}_{h'} \in \widehat{\mathcal{C}}_{i,h'} \right\}\end{aligned}\quad (19)$$

Definition C.3 (Pessimism Event).

$$\begin{aligned}\check{\Psi}_{k,h} &:= \left\{ \forall h \leq h' \leq H : \boldsymbol{\mu}_{h'} \in \check{\mathcal{C}}_{k,h'} \right\} \\ \check{\Psi}_h &:= \left\{ \forall i \in [K], \forall h \leq h' \leq H : \boldsymbol{\mu}_{h'} \in \check{\mathcal{C}}_{i,h'} \right\}\end{aligned}\quad (20)$$

In this section, independent confidence sets $\bar{\mathcal{C}}_{k,h}, \hat{\mathcal{C}}_{k,h}, \check{\mathcal{C}}_{k,h}, \check{\check{\mathcal{C}}}_{k,h}$ are built in Lemma C.8, C.9, C.10, and C.11 respectively in Appendix C.1. These independent confidence sets are built to upper bound the variance of the considered value function. Specifically, the difference between the estimated variance of the constructed optimistic value function and the real variance of the optimal value function, i.e., $|\mathbb{V}_h V_{h+1}^*(s_h^k, a_h^k) - [\widehat{\mathbb{V}}_{k,h} \widehat{V}_{k,h+1}](s_h^k, a_h^k)|$, is upper bounded in high probability in Lemma C.13. In addition, the variance $[\mathbb{V}_h(\widehat{V}_{k,h+1} - V_{h+1}^*)](s_h^k, a_h^k)$ is also upper bounded in Lemma C.14. Subsequently, dependent confidence sets $\widehat{\mathcal{C}}_{k,h}^{(1)}, \widehat{\mathcal{C}}_{k,h}^{(2)}$ can be built based on the independent confidence sets in Lemma C.16, C.17 respectively in Appendix C.3. Thus, the confidence set $\widehat{\mathcal{C}}_{k,h}$, the goal of this section, holds trivially if $\widehat{\mathcal{C}}_{k,h}^{(1)}, \widehat{\mathcal{C}}_{k,h}^{(2)}$ both hold. Finally, Lemma 7.5 in the main paper is proved in Appendix C.4.

Before the formal proof begins, we give some necessary definitions. We first give definitions about-measurable space and filtration required for our proofs.

Measurable Space Note that the stochasticity in the transition probability of the MDP are the only source of randomness. Denote \mathbb{P} as the gather of the distributions over state-action pair sequence $(\mathcal{S} \times \mathcal{A})^{\mathbb{N}}$, induced by the interconnection of policy obtained from LSVI-UCB⁺ algorithm and the episodic linear MDP \mathcal{M} . Denote \mathbb{E} as the corresponding expectation operator. Hence, all random variables can be defined over the sample space $\Omega = (\mathcal{S} \times \mathcal{A})^{\mathbb{N}}$. Thus, we work with the probability space given by the triplet $(\Omega, \mathcal{F}, \mathbb{P})$, where \mathcal{F} is the product σ -algebra generated by the discrete σ -algebras underlying \mathcal{S} and \mathcal{A} .

Definition C.4 (Filtration). For any $k \in [K]$ and any $h \in [H]$, let $\mathcal{F}_{k,h}$ be the σ -algebra generated by the random variables representing the state-action pairs up to and including that appears in stage h of episode k .

Measurability Thus, $[\widehat{\mathbb{V}}_{k,h} \widehat{V}_{k,h+1}](s_h^k, a_h^k), U_{k,h}, E_{k,h}, \varsigma_{k,h}, \tilde{\sigma}_{k,h}, \hat{\sigma}_{k,h}, \widehat{\Lambda}_{k+1,h}$ are $\mathcal{F}_{k,h}$ -measurable, $\widehat{\boldsymbol{\mu}}_{k+1,h}$ is $\mathcal{F}_{k,h+1}$ -measurable, $\widehat{Q}_{k,h}, \widehat{\check{Q}}_{k,h}, \widehat{\check{Q}}_{k,h}, \widehat{V}_{k,h}, \widehat{\check{V}}_{k,h}, \widehat{\check{V}}_{k,h}, \pi_h^k$ are $\mathcal{F}_{k-1,H}$ -measurable, but not $\mathcal{F}_{k-1,h}$ -measurable due to their backwards construction.

C.1. Independent Confidence Sets

In this subsection, independent confidence sets $\bar{\mathcal{C}}_{k,h}, \hat{\mathcal{C}}_{k,h}, \check{\mathcal{C}}_{k,h}, \check{\check{\mathcal{C}}}_{k,h}$ are built in Lemma C.8, C.9, C.10, C.11, respectively. During the construction of these confidence sets, it is unavoidable to build a uniform convergence argument by covering net of the encountered function class. Thus, we also present the definition of possibly encountered function classes in the following.

Definition C.5 (Optimistic Value Function Class). Let $\widehat{\mathcal{V}}$ denote a class of functions mapping from \mathcal{S} to \mathbb{R} with following parametric form

$$\widehat{\mathcal{V}}(\cdot) = \min \left\{ \max_a \boldsymbol{w}^\top \boldsymbol{\phi}(\cdot, a) + \beta \sqrt{\boldsymbol{\phi}(\cdot, a)^\top \boldsymbol{\Lambda}^{-1} \boldsymbol{\phi}(\cdot, a)}, H \right\},$$

where the parameters $(\boldsymbol{w}, \beta, \boldsymbol{\Lambda})$ satisfy $\|\boldsymbol{w}\|_2 \leq L, \beta \in [0, B]$, the minimum eigenvalue satisfies $\lambda_{\min}(\boldsymbol{\Lambda}) \geq \lambda$, and $\sup_{s,a} \|\boldsymbol{\phi}(s, a)\|_2 \leq 1$.

Definition C.6 (Squared Optimistic Value Function Class). Let $\widehat{\mathcal{V}}^2$ denote a class of functions mapping from \mathcal{S} to \mathbb{R} with following parametric form

$$\widehat{\mathcal{V}}^2(\cdot) = \left[\min \left\{ \max_a \boldsymbol{w}^\top \boldsymbol{\phi}(\cdot, a) + \beta \sqrt{\boldsymbol{\phi}(\cdot, a)^\top \boldsymbol{\Lambda}^{-1} \boldsymbol{\phi}(\cdot, a)}, H \right\} \right]^2,$$

where the parameters $(\mathbf{w}, \beta, \Lambda)$ satisfy $\|\mathbf{w}\|_2 \leq L, \beta \in [0, B]$, the minimum eigenvalue satisfies $\lambda_{\min}(\Lambda) \geq \lambda$, and $\sup_{s,a} \|\phi(s, a)\|_2 \leq 1$.

Definition C.7 (Pessimistic Value Function Class). Let $\check{\mathcal{V}}$ denote a class of functions mapping from \mathcal{S} to \mathbb{R} with following parametric form

$$\check{\mathcal{V}}(\cdot) = \max \left\{ \max_a \mathbf{w}^\top \phi(\cdot, a) - \beta \sqrt{\phi(\cdot, a)^\top \Lambda^{-1} \phi(\cdot, a)}, 0 \right\},$$

where the parameters $(\mathbf{w}, \beta, \Lambda)$ satisfy $\|\mathbf{w}\|_2 \leq L, \beta \in [0, B]$, the minimum eigenvalue satisfies $\lambda_{\min}(\Lambda) \geq \lambda$, and $\sup_{s,a} \|\phi(s, a)\|_2 \leq 1$.

Now we are ready to build four independent confidence sets $\bar{\mathcal{C}}_{k,h}, \hat{\mathcal{C}}_{k,h}, \check{\mathcal{C}}_{k,h}, \check{\mathcal{C}}_{k,h}$. Since radius of independent confidence sets, i.e., $\bar{\beta}, \hat{\beta}, \check{\beta}, \check{\beta}$, will not become dominant terms in the final regret bound, we build these four confidence sets with traditional Hoeffding inequality (Lemma F.3) with covering net arguments.

Lemma C.8. In Algorithm 1, for any $\delta \in (0, 1)$, any $k \in [K]$ fixed $h \in [H]$, with probability at least $1 - \delta/H$:

$$\boldsymbol{\mu}_h \in \bar{\mathcal{C}}_{k,h} = \left\{ \boldsymbol{\mu} : \left\| (\boldsymbol{\mu} - \hat{\boldsymbol{\mu}}_{k,h}) \hat{\mathbf{V}}_{k,h+1} \right\|_{\hat{\Lambda}_{k,h}} \leq \bar{\beta} \right\},$$

where

$$\bar{\beta} = \sqrt{H} \sqrt{d \log \left(1 + \frac{K}{Hd\lambda} \right) + \log \left(\frac{H}{\delta} \right) + d \log \left(1 + \frac{4KL}{H\sqrt{\lambda}} \right) + d^2 \log \left(1 + \frac{8K^2 \hat{B}^2 \sqrt{d}}{H^2 \lambda^2} \right)} + H\sqrt{\lambda d} + 2.$$

Here $L = W + K/\lambda$ and \hat{B} is a constant satisfying $\hat{\beta} \leq \hat{B}$ with $\hat{\beta}$ given in Lemma C.18.

Proof. Initially, note that we have

$$\begin{aligned} & \left\| (\hat{\boldsymbol{\mu}}_{k,h} - \boldsymbol{\mu}_h) \hat{\mathbf{V}}_{k,h+1} \right\|_{\hat{\Lambda}_{k,h}} \\ &= \left\| \hat{\Lambda}_{k,h}^{-1} \left[-\lambda \boldsymbol{\mu}_h + \sum_{i=1}^{k-1} \hat{\sigma}_{i,h}^{-2} \phi(s_h^i, a_h^i) \boldsymbol{\epsilon}_h^i \right] \hat{\mathbf{V}}_{k,h+1} \right\|_{\hat{\Lambda}_{k,h}} \\ &= \left\| -\lambda \boldsymbol{\mu}_h \hat{\mathbf{V}}_{k,h+1} + \sum_{i=1}^{k-1} \hat{\sigma}_{i,h}^{-2} \phi(s_h^i, a_h^i) \boldsymbol{\epsilon}_h^i \hat{\mathbf{V}}_{k,h+1} \right\|_{\hat{\Lambda}_{k,h}^{-1}} \\ &\leq \left\| -\lambda \boldsymbol{\mu}_h \hat{\mathbf{V}}_{k,h+1} \right\|_{\hat{\Lambda}_{k,h}^{-1}} + \left\| \sum_{i=1}^{k-1} \hat{\sigma}_{i,h}^{-2} \phi(s_h^i, a_h^i) \boldsymbol{\epsilon}_h^i \hat{\mathbf{V}}_{k,h+1} \right\|_{\hat{\Lambda}_{k,h}^{-1}} \\ &\leq \frac{1}{\sqrt{\lambda}} \cdot \lambda H \sqrt{d} + \left\| \sum_{i=1}^{k-1} \hat{\sigma}_{i,h}^{-2} \phi(s_h^i, a_h^i) \boldsymbol{\epsilon}_h^i \hat{\mathbf{V}}_{k,h+1} \right\|_{\hat{\Lambda}_{k,h}^{-1}}, \end{aligned} \tag{21}$$

where the first equality is due to Eq. (62) in Lemma F.7, the first inequality is due to triangle inequality, and the second inequality holds since $\|\boldsymbol{\mu}_h \hat{\mathbf{V}}_{k,h+1}\|_2 \leq H\sqrt{d}$ and the minimum eigenvalue of $\hat{\Lambda}_{k,h}$ is no less than λ .

Thus, we bound $\left\| \sum_{i=1}^{k-1} \hat{\sigma}_{i,h}^{-2} \phi(s_h^i, a_h^i) \boldsymbol{\epsilon}_h^i \hat{\mathbf{V}}_{k,h+1} \right\|_{\hat{\Lambda}_{k,h}^{-1}}$ in the following. However, $\hat{\mathbf{V}}_{k,h+1}$ is $\mathcal{F}_{k,h}$ -measurable, which brings obstacles in directly applying self-normalized bound for martingales. We need to build a uniform convergence argument for $\hat{\mathbf{V}}_{k,h+1}$.

For any $(k, h) \in \times [K] \times [H]$, $\hat{\mathbf{V}}_{k,h}(\cdot) = \min\{\max_a \langle \boldsymbol{\theta}_h + \hat{\boldsymbol{\mu}}_{k,h} \hat{\mathbf{V}}_{k,h+1}, \phi(\cdot, a) \rangle + \hat{\beta} \|\phi(\cdot, a)\|_{\hat{\Lambda}_{k,h}^{-1}}, H\}$ in Algorithm 1. Moreover, we have

$$\left\| \boldsymbol{\theta}_h + \hat{\boldsymbol{\mu}}_{k,h} \hat{\mathbf{V}}_{k,h+1} \right\|_2 = \left\| \boldsymbol{\theta}_h + \hat{\Lambda}_{k,h}^{-1} \sum_{i=1}^{k-1} \hat{\sigma}_{i,h}^{-2} \phi(s_h^i, a_h^i) \hat{\mathbf{V}}_{k,h+1}(s_{h+1}^i) \right\|_2$$

$$\begin{aligned}
 &\leq W + \left\| \widehat{\Lambda}_{k,h}^{-1} \sum_{i=1}^{k-1} \widehat{\sigma}_{i,h}^{-2} \phi(s_h^i, a_h^i) \widehat{V}_{k,h+1}(s_{h+1}^i) \right\|_2 \\
 &\leq W + \frac{H}{(\sqrt{H})^2} \left\| \widehat{\Lambda}_{k,h}^{-1} \sum_{i=1}^{k-1} \phi(s_h^i, a_h^i) \right\|_2 \leq W + K/\lambda,
 \end{aligned}$$

where the first inequality holds due to triangle inequality, the second inequality holds since $\widehat{V}_{k,h+1}(\cdot) \leq H$ and $\widehat{\sigma}_{i,h} \geq \sqrt{H}$ for any $i \in [k]$, and the last inequality holds since $\lambda_{\min}(\Lambda_{k,h}) \geq \lambda$ and $\sup_{s,a} \|\phi(s,a)\|_2 \leq 1$. Subsequently, we claim $\widehat{V}_{k,h} \in \widehat{\mathcal{V}}$, where $\widehat{\mathcal{V}}$ is defined in Definition C.5, with $L = W + K/\lambda$ and $B = \widehat{B}$. Here, \widehat{B} is a constant satisfying $\widehat{\beta} \leq \widehat{B}$ with $\widehat{\beta}$ specified in Lemma C.18.

Then, we fix a function $V(\cdot) \in \widehat{\mathcal{V}} : \mathcal{S} \mapsto [0, H]$. Let $\mathcal{G}_i = \mathcal{F}_{i,h}$, $\mathbf{x}_i = \widehat{\sigma}_{i,h}^{-1} \phi(s_h^i, a_h^i)$ and $\eta_i = \widehat{\sigma}_{i,h}^{-1} \boldsymbol{\epsilon}_h^i \top \mathbf{V} = \widehat{\sigma}_{i,h}^{-1} \langle \boldsymbol{\mu}_h \mathbf{V}, \phi(s_h^i, a_h^i) \rangle - \widehat{\sigma}_{i,h}^{-1} V(s_{h+1}^i)$. It is clear that \mathbf{x}_i is \mathcal{G}_i -measurable and η_i is \mathcal{G}_{i+1} -measurable. Since $\varsigma_i \geq \sqrt{H}$, we have $\widehat{\sigma}_{i,h} \geq \sqrt{H}$. Besides, we have $\|\mathbf{x}_i\|_2 \leq 1/\sqrt{H}$, $\mathbb{E}[\eta_i | \mathcal{G}_i] = 0$, $|\eta_i| \leq \sqrt{H}$ and $\mathbb{E}[\eta_i^2 | \mathcal{G}_i] \leq H$. By Lemma F.3, we obtain that, with probability at least $1 - \delta/H$, for any $k \in [K]$ and fixed $h \in [H]$,

$$\left\| \sum_{i=1}^{k-1} \widehat{\sigma}_{i,h}^{-2} \phi(s_h^i, a_h^i) \boldsymbol{\epsilon}_h^i \top \mathbf{V} \right\|_{\widehat{\Lambda}_{k,h}^{-1}} \leq \sqrt{H} \sqrt{d \log \left(1 + \frac{K}{Hd\lambda} \right) + \log \left(\frac{H}{\delta} \right)}.$$

Denote the ε -cover of function class $\widehat{\mathcal{V}}$ as $\widehat{\mathcal{N}}_\varepsilon$. Consider an arbitrary $f \in \widehat{\mathcal{V}}$. From the definition of ε -cover, we know that for f , there exists a $V \in \widehat{\mathcal{N}}_\varepsilon$, such that $\|f - V\|_\infty \leq \varepsilon$. Since $\|\boldsymbol{\epsilon}_h^i \top (f - V)\|_2 \leq \|\boldsymbol{\epsilon}_h^i\|_1 \|f - V\|_\infty \leq 2\varepsilon$ and $\|\sum_{i=1}^{k-1} \widehat{\sigma}_{i,h}^{-2} \phi(s_h^i, a_h^i)\|_{\widehat{\Lambda}_{k,h}^{-1}} \leq K/(H\sqrt{\lambda})$, we have

$$\left\| \sum_{i=1}^{k-1} \widehat{\sigma}_{i,h}^{-2} \phi(s_h^i, a_h^i) \boldsymbol{\epsilon}_h^i \top (f - V) \right\|_{\widehat{\Lambda}_{k,h}^{-1}} \leq \frac{2\varepsilon K}{H\sqrt{\lambda}}. \quad (22)$$

This further implies the following inequality holds with probability at least $1 - \delta/H$:

$$\begin{aligned}
 &\left\| \sum_{i=1}^{k-1} \widehat{\sigma}_{i,h}^{-2} \phi(s_h^i, a_h^i) \boldsymbol{\epsilon}_h^i \top f \right\|_{\widehat{\Lambda}_{k,h}^{-1}} \\
 &\leq \left\| \sum_{i=1}^{k-1} \widehat{\sigma}_{i,h}^{-2} \phi(s_h^i, a_h^i) \boldsymbol{\epsilon}_h^i \top \mathbf{V} \right\|_{\widehat{\Lambda}_{k,h}^{-1}} + \left\| \sum_{i=1}^{k-1} \widehat{\sigma}_{i,h}^{-2} \phi(s_h^i, a_h^i) \boldsymbol{\epsilon}_h^i \top (f - V) \right\|_{\widehat{\Lambda}_{k,h}^{-1}} \\
 &\leq \left\| \sum_{i=1}^{k-1} \widehat{\sigma}_{i,h}^{-2} \phi(s_h^i, a_h^i) \boldsymbol{\epsilon}_h^i \top \mathbf{V} \right\|_{\widehat{\Lambda}_{k,h}^{-1}} + \frac{2\varepsilon K}{H\sqrt{\lambda}} \\
 &\leq \sqrt{H} \sqrt{d \log \left(1 + \frac{K}{Hd\lambda} \right) + \log \left(\frac{H}{\delta} \right) + \log |\widehat{\mathcal{N}}_\varepsilon|} + \frac{2\varepsilon K}{H\sqrt{\lambda}}.
 \end{aligned} \quad (23)$$

where the first inequality is due to the triangle inequality, the second one holds by Eq. (22), and the third inequality holds by a union bound over all functions in $\widehat{\mathcal{N}}_\varepsilon$ with

$$\log |\widehat{\mathcal{N}}_\varepsilon| \leq d \log \left(1 + \frac{4L}{\varepsilon} \right) + d^2 \log \left(1 + \frac{8\widehat{B}^2 \sqrt{d}}{\lambda \varepsilon^2} \right)$$

according to Lemma F.9.

Note that Eq. (21) holds and $\widehat{V}_{k,h+1}(\cdot) \in \widehat{\mathcal{V}}$. We have with probability at least $1 - \delta/H$, for any $k \in [K]$ and fixed $h \in [H]$:

$$\left\| (\widehat{\boldsymbol{\mu}}_{k,h} - \boldsymbol{\mu}_h) \widehat{V}_{k,h+1} \right\|_{\widehat{\Lambda}_{k,h}} \leq \left\| \sum_{i=1}^{k-1} \widehat{\sigma}_{i,h}^{-2} \phi(s_h^i, a_h^i) \boldsymbol{\epsilon}_h^i \top \widehat{V}_{k,h+1} \right\|_{\widehat{\Lambda}_{k,h}^{-1}} + H\sqrt{\lambda d}$$

$$\begin{aligned} &\leq \sqrt{H} \sqrt{d \log \left(1 + \frac{K}{Hd\lambda}\right) + \log \left(\frac{H}{\delta}\right) + d \log \left(1 + \frac{4KL}{H\sqrt{\lambda}}\right) + d^2 \log \left(1 + \frac{8K^2 \hat{B}^2 \sqrt{d}}{H^2 \lambda^2}\right) + H\sqrt{\lambda d} + 2} \\ &= \bar{\beta}, \end{aligned}$$

where the last inequality holds by Eq. (23) and setting $\varepsilon = H\sqrt{\lambda}/K$. \square

After building the confidence set $\bar{\mathcal{C}}_{k,h}$ in Lemma C.8, confidence sets $\dot{\mathcal{C}}_{k,h}$, $\tilde{\mathcal{C}}_{k,h}$ and $\check{\mathcal{C}}_{k,h}$ can be built similarly.

Lemma C.9. *In Algorithm 1, for any $\delta \in (0, 1)$, any $k \in [K]$ and fixed $h \in [H]$, with probability at least $1 - \delta/H$:*

$$\boldsymbol{\mu}_h \in \dot{\mathcal{C}}_{k,h} = \left\{ \boldsymbol{\mu} : \left\| (\boldsymbol{\mu} - \hat{\boldsymbol{\mu}}_{k,h}) \hat{\mathbf{V}}_{k,h+1} \right\|_{\hat{\boldsymbol{\Lambda}}_{k,h}} \leq \hat{\beta} \right\},$$

where

$$\hat{\beta} = \sqrt{H} \sqrt{d \log \left(1 + \frac{K}{Hd\lambda}\right) + \log \left(\frac{H}{\delta}\right) + d \log \left(1 + \frac{4KL}{H\sqrt{\lambda}}\right) + d^2 \log \left(1 + \frac{8K^2 \hat{B}^2 \sqrt{d}}{H^2 \lambda^2}\right) + H\sqrt{\lambda d} + 2}.$$

Here $L = W + K/\lambda$ and \hat{B} is the solution of the inequality $3\hat{\beta} \leq \hat{B}$ concerning \hat{B} , where $\hat{\beta}$ is specified in Lemma C.18.

Proof. Notice $\hat{\mathbf{V}}_{k,h}(\cdot) = \min\{\max_a \langle \boldsymbol{\theta}_h + \hat{\boldsymbol{\mu}}_{k,h} \hat{\mathbf{V}}_{k,h+1}, \boldsymbol{\phi}(\cdot, a) \rangle + 3\hat{\beta} \|\boldsymbol{\phi}(\cdot, a)\|_{\hat{\boldsymbol{\Lambda}}_{k,h}^{-1}}, H\}$ in Algorithm 1 for any $(k, h) \in [K] \times [H]$. The proof is the same as that for Lemma C.8, except the considered function class $\hat{\mathcal{V}}$ is slightly different, i.e., the B is set as \hat{B} , which is the solution of the inequality $3\hat{\beta} \leq \hat{B}$ concerning \hat{B} , where $\hat{\beta}$ is specified in Lemma C.18. \square

Lemma C.10. *In Algorithm 1, for any $\delta \in (0, 1)$, any $k \in [K]$ and fixed $h \in [H]$, with probability at least $1 - \delta/H$:*

$$\boldsymbol{\mu}_h \in \tilde{\mathcal{C}}_{k,h} = \left\{ \boldsymbol{\mu} : \left\| (\boldsymbol{\mu} - \hat{\boldsymbol{\mu}}_{k,h}) \hat{\mathbf{V}}_{k,h+1}^2 \right\|_{\hat{\boldsymbol{\Lambda}}_{k,h}} \leq \tilde{\beta} \right\},$$

where

$$\tilde{\beta} = \sqrt{H^3} \sqrt{d \log \left(1 + \frac{K}{Hd\lambda}\right) + \log \left(\frac{H}{\delta}\right) + d \log \left(1 + \frac{8KL}{\sqrt{\lambda}}\right) + d^2 \log \left(1 + \frac{32K^2 \hat{B}^2 \sqrt{d}}{\lambda^2}\right) + H^2 \sqrt{\lambda d} + 2}.$$

Here $L = W + K/\lambda$ and \hat{B} is a constant satisfying $\tilde{\beta} \leq \hat{B}$ with $\hat{\beta}$ given in Lemma C.18.

Proof. The proof of this lemma is almost the same as that of Lemma C.8, except for replacing $\hat{\mathbf{V}}_{k,h+1}$, $\hat{\mathcal{V}}$, and $\hat{\mathcal{N}}_\varepsilon$ by $\hat{\mathbf{V}}_{k,h+1}^2$, $\hat{\mathcal{V}}^2$, and $\hat{\mathcal{N}}_\varepsilon^2$, respectively. Here

$$\log |\hat{\mathcal{N}}_\varepsilon^2| \leq d \log \left(1 + \frac{8LH}{\varepsilon}\right) + d^2 \log \left(1 + \frac{32\hat{B}^2 H^2 \sqrt{d}}{\lambda \varepsilon^2}\right),$$

where $L = W + k/\lambda$ and \hat{B} is a constant satisfying $\tilde{\beta} \leq \hat{B}$ with $\hat{\beta}$ given in Lemma C.18.

After setting $\varepsilon = H\sqrt{\lambda}/K$, it can be proved that with probability at least $1 - \delta/H$, for any $k \in [K]$ and fixed $h \in [H]$:

$$\begin{aligned} &\left\| (\hat{\boldsymbol{\mu}}_{k,h} - \boldsymbol{\mu}_h) \hat{\mathbf{V}}_{k,h+1}^2 \right\|_{\hat{\boldsymbol{\Lambda}}_{k,h}} \\ &\leq \sqrt{H^3} \sqrt{d \log \left(1 + \frac{K}{Hd\lambda}\right) + \log \left(\frac{H}{\delta}\right) + d \log \left(1 + \frac{8KL}{\sqrt{\lambda}}\right) + d^2 \log \left(1 + \frac{32K^2 \hat{B}^2 \sqrt{d}}{\lambda^2}\right) + H^2 \sqrt{\lambda d} + 2} \\ &= \tilde{\beta}. \end{aligned}$$

\square

Lemma C.11. *In Algorithm 1, for any $\delta \in (0, 1)$, any $k \in [K]$ and fixed $h \in [H]$, with probability at least $1 - \delta/H$:*

$$\boldsymbol{\mu}_h \in \check{\mathcal{C}}_{k,h} = \left\{ \boldsymbol{\mu} : \left\| (\boldsymbol{\mu} - \hat{\boldsymbol{\mu}}_{k,h}) \check{\mathbf{V}}_{k,h+1} \right\|_{\hat{\boldsymbol{\Lambda}}_{k,h}} \leq \check{\beta} \right\},$$

where

$$\check{\beta} = \sqrt{H} \sqrt{d \log \left(1 + \frac{K}{Hd\lambda} \right) + \log \left(\frac{H}{\delta} \right) + d \log \left(1 + \frac{4KL}{H\sqrt{\lambda}} \right) + d^2 \log \left(1 + \frac{8K^2 \check{B}^2 \sqrt{d}}{H^2 \lambda^2} \right) + H\sqrt{\lambda d} + 2}. \quad (24)$$

Here $L = W + K/\lambda$ and \check{B} is a constant satisfying $\check{\beta} \leq \check{B}$.

Proof. The proof of this lemma is similar to that of Lemma C.8 except for replacing $\hat{V}_{k,h+1}$, $\hat{\mathcal{V}}$, and $\hat{\mathcal{N}}_\varepsilon$ by $\check{V}_{k,h+1}$, $\check{\mathcal{V}}$, and $\check{\mathcal{N}}_\varepsilon$, respectively, where

$$\log |\check{\mathcal{N}}_\varepsilon| \leq d \log \left(1 + \frac{4L}{\varepsilon} \right) + d^2 \log \left(1 + \frac{8\check{B}^2 \sqrt{d}}{\lambda \varepsilon^2} \right).$$

Here $L = W + K/\lambda$ and \check{B} is a constant satisfying $\check{\beta} \leq \check{B}$ with $\check{\beta}$ specified in Eq. (24).

After setting $\varepsilon = H\sqrt{\lambda}/K$, it can be proved that with probability at least $1 - \delta/H$, for any $k \in [K]$ and fixed $h \in [H]$:

$$\begin{aligned} & \left\| (\hat{\boldsymbol{\mu}}_{k,h} - \boldsymbol{\mu}_h) \check{\mathbf{V}}_{k,h+1} \right\|_{\hat{\boldsymbol{\Lambda}}_{k,h}} \\ & \leq \sqrt{H} \sqrt{d \log \left(1 + \frac{K}{Hd\lambda} \right) + \log \left(\frac{H}{\delta} \right) + d \log \left(1 + \frac{4KL}{H\sqrt{\lambda}} \right) + d^2 \log \left(1 + \frac{8K^2 \check{B}^2 \sqrt{d}}{H^2 \lambda^2} \right) + H\sqrt{\lambda d} + 2} \\ & = \check{\beta}. \end{aligned}$$

□

C.2. Variance Upper Bound

In this section, we prove some necessary lemmas to build upper bounds of value function variances, including variances of $V_{h+1}^*(\cdot)$ and $[\hat{V}_{k,h+1} - V_{h+1}^*](\cdot)$. Specifically, we present Lemma C.13 in Appendix C.2.1 to bound the difference between the estimated variance of the constructed optimistic value function and the real variance of the optimal value function, i.e., $|\mathbb{V}_h V_{h+1}^*(s_h^k, a_h^k) - [\hat{\mathbb{V}}_{k,h} \hat{V}_{k,h+1}](s_h^k, a_h^k)|$, with high probability. In addition, we also present Lemma C.14 in Appendix C.2.2 to upper bound the variance $[\mathbb{V}_h (\hat{V}_{k,h+1} - V_{h+1}^*)](s_h^k, a_h^k)$.

C.2.1. VARIANCE OF $V_{h+1}^*(\cdot)$

Before proving Lemma C.13, we first present Lemma C.12, which upper bounds $|\mathbb{P}_h(\hat{V}_{k,h+1} - V_{h+1}^*)](s_h^k, a_h^k)|$ under optimism and pessimism events $\hat{\Psi}_{k,h+1} \cap \check{\Psi}_{k,h+1}$, and serves as the building block for Lemma C.13.

Lemma C.12. *In Algorithm 1, for any $k \in [K]$ and any $h \in [H]$, under $\hat{\Psi}_{k,h+1} \cap \check{\Psi}_{k,h+1}$, we have*

$$\begin{aligned} & \left| \langle \boldsymbol{\mu}_h \hat{\mathbf{V}}_{k,h+1}, \boldsymbol{\phi}(s_h^k, a_h^k) \rangle - \langle \boldsymbol{\mu}_h \mathbf{V}_{h+1}^*, \boldsymbol{\phi}(s_h^k, a_h^k) \rangle \right| \leq \left| \langle (\hat{\boldsymbol{\mu}}_{k,h} - \boldsymbol{\mu}_h) \hat{\mathbf{V}}_{k,h+1}, \boldsymbol{\phi}(s_h^k, a_h^k) \rangle \right| \\ & \quad + \left| \langle \hat{\boldsymbol{\mu}}_{k,h} \hat{\mathbf{V}}_{k,h+1} - \hat{\boldsymbol{\mu}}_{k,h} \check{\mathbf{V}}_{k,h+1}, \boldsymbol{\phi}(s_h^k, a_h^k) \rangle \right| + \left| \langle (\hat{\boldsymbol{\mu}}_{k,h} - \boldsymbol{\mu}_h) \check{\mathbf{V}}_{k,h+1}, \boldsymbol{\phi}(s_h^k, a_h^k) \rangle \right|. \end{aligned}$$

Proof. By definition, we have

$$\begin{aligned} & \left| \langle \boldsymbol{\mu}_h \hat{\mathbf{V}}_{k,h+1}, \boldsymbol{\phi}(s_h^k, a_h^k) \rangle - \langle \boldsymbol{\mu}_h \mathbf{V}_{h+1}^*, \boldsymbol{\phi}(s_h^k, a_h^k) \rangle \right| \\ & = \left| \mathbb{P}_h \hat{V}_{k,h+1}(s_h^k, a_h^k) - \mathbb{P}_h V_{h+1}^*(s_h^k, a_h^k) \right| = \mathbb{P}_h \hat{V}_{k,h+1}(s_h^k, a_h^k) - \mathbb{P}_h V_{h+1}^*(s_h^k, a_h^k) \end{aligned}$$

$$\leq \mathbb{P}_h \widehat{V}_{k,h+1}(s_h^k, a_h^k) - \mathbb{P}_h \check{V}_{k,h+1}(s_h^k, a_h^k) = \left| \mathbb{P}_h \widehat{V}_{k,h+1}(s_h^k, a_h^k) - \mathbb{P}_h \check{V}_{k,h+1}(s_h^k, a_h^k) \right|,$$

where the second equality holds since \mathbb{P}_h is a valid distribution and $\widehat{V}_{k,h+1}(\cdot) \geq V_{h+1}^*(\cdot)$ under $\widehat{\Psi}_{k,h+1}$ by Lemma D.1, the first inequality holds since \mathbb{P}_h is a valid distribution plus $V_{h+1}^*(\cdot) \geq \check{V}_{k,h+1}(\cdot)$ under $\check{\Psi}_{k,h+1}$ by Lemma D.1, and the last equality holds since \mathbb{P}_h is a valid distribution plus $\widehat{V}_{k,h+1}(\cdot) \geq \check{V}_{k,h+1}(\cdot)$ under $\widehat{\Psi}_{k,h+1} \cap \check{\Psi}_{k,h+1}$ by Lemma D.1.

Therefore,

$$\begin{aligned} & \left| \mathbb{P}_h \widehat{V}_{k,h+1}(s_h^k, a_h^k) - \mathbb{P}_h V_{h+1}^*(s_h^k, a_h^k) \right| \leq \left| \mathbb{P}_h \widehat{V}_{k,h+1}(s_h^k, a_h^k) - \mathbb{P}_h \check{V}_{k,h+1}(s_h^k, a_h^k) \right| \\ &= \left| \mathbb{P}_h \widehat{V}_{k,h+1}(s_h^k, a_h^k) - \widehat{\mathbb{P}}_{k,h} \widehat{V}_{k,h+1}(s_h^k, a_h^k) + \widehat{\mathbb{P}}_{k,h} \widehat{V}_{k,h+1}(s_h^k, a_h^k) - \widehat{\mathbb{P}}_{k,h} \check{V}_{k,h+1}(s_h^k, a_h^k) \right. \\ & \quad \left. + \widehat{\mathbb{P}}_{k,h} \check{V}_{k,h+1}(s_h^k, a_h^k) - \mathbb{P}_h \check{V}_{k,h+1}(s_h^k, a_h^k) \right| \\ &\leq \left| \mathbb{P}_h \widehat{V}_{k,h+1}(s_h^k, a_h^k) - \widehat{\mathbb{P}}_{k,h} \widehat{V}_{k,h+1}(s_h^k, a_h^k) \right| + \left| \widehat{\mathbb{P}}_{k,h} \left[\widehat{V}_{k,h+1} - \check{V}_{k,h+1} \right](s_h^k, a_h^k) \right| \\ & \quad + \left| \widehat{\mathbb{P}}_{k,h} \check{V}_{k,h+1}(s_h^k, a_h^k) - \mathbb{P}_h \check{V}_{k,h+1}(s_h^k, a_h^k) \right| \\ &= \left| \langle (\widehat{\mu}_{k,h} - \mu_h) \widehat{V}_{k,h+1}, \phi(s_h^k, a_h^k) \rangle \right| + \left| \langle \widehat{\mu}_{k,h} (\widehat{V}_{k,h+1} - \check{V}_{k,h+1}), \phi(s_h^k, a_h^k) \rangle \right| \\ & \quad + \left| \langle (\widehat{\mu}_{k,h} - \mu_h) \check{V}_{k,h+1}, \phi(s_h^k, a_h^k) \rangle \right|. \end{aligned}$$

□

Based on Lemma C.12, we are ready to present Lemma C.13.

Lemma C.13. *In Algorithm 1, for any $k \in [K]$ and any $h \in [H]$, under $\widehat{\Psi}_{k,h+1} \cap \check{\Psi}_{k,h+1}$, we have*

$$\left| [\mathbb{V}_h V_{h+1}^*](s_h^k, a_h^k) - [\widehat{\mathbb{V}}_{k,h} \widehat{V}_{k,h+1}](s_h^k, a_h^k) \right| \leq \min \left\{ \left\| (\widehat{\mu}_{k,h} - \mu_h) \widehat{V}_{k,h+1} \right\|_{\widehat{\Lambda}_{k,h}} \left\| \phi(s_h^k, a_h^k) \right\|_{\widehat{\Lambda}_{k,h}^{-1}} + 4H \Delta_{k,h}, 2H^2 \right\},$$

where

$$\begin{aligned} \Delta_{k,h} &= \left\| (\widehat{\mu}_{k,h} - \mu_h) \widehat{V}_{k,h+1} \right\|_{\widehat{\Lambda}_{k,h}} \left\| \phi(s_h^k, a_h^k) \right\|_{\widehat{\Lambda}_{k,h}^{-1}} + \left| \langle \widehat{\mu}_{k,h} (\widehat{V}_{k,h+1} - \check{V}_{k,h+1}), \phi(s_h^k, a_h^k) \rangle \right| \\ & \quad + \left\| (\widehat{\mu}_{k,h} - \mu_h) \check{V}_{k,h+1} \right\|_{\widehat{\Lambda}_{k,h}} \left\| \phi(s_h^k, a_h^k) \right\|_{\widehat{\Lambda}_{k,h}^{-1}}. \end{aligned}$$

Proof. By definition, we have

$$\begin{aligned} & \left| [\mathbb{V}_h V_{h+1}^*](s_h^k, a_h^k) - [\widehat{\mathbb{V}}_{k,h} \widehat{V}_{k,h+1}](s_h^k, a_h^k) \right| \\ &= \left| \langle \mu_h \mathbf{V}_{h+1}^*, \phi(s_h^k, a_h^k) \rangle - \left[\langle \widehat{\mu}_{k,h} \widehat{V}_{k,h+1}, \phi(s_h^k, a_h^k) \rangle \right]_{[0, H^2]} \right| \\ & \quad + \left\{ \left[\langle \widehat{\mu}_{k,h} \widehat{V}_{k,h+1}, \phi(s_h^k, a_h^k) \rangle \right]_{[0, H]} \right\}^2 - \left[\langle \mu_h \mathbf{V}_{h+1}^*, \phi(s_h^k, a_h^k) \rangle \right]^2 \Big| \\ &\leq \underbrace{\left| \left[\langle \widehat{\mu}_{k,h} \widehat{V}_{k,h+1}, \phi(s_h^k, a_h^k) \rangle \right]_{[0, H^2]} - \langle \mu_h \mathbf{V}_{h+1}^*, \phi(s_h^k, a_h^k) \rangle \right|}_{I_1} \\ & \quad + \underbrace{\left| \left\{ \left[\langle \widehat{\mu}_{k,h} \widehat{V}_{k,h+1}, \phi(s_h^k, a_h^k) \rangle \right]_{[0, H]} \right\}^2 - \left[\langle \mu_h \mathbf{V}_{h+1}^*, \phi(s_h^k, a_h^k) \rangle \right]^2 \right|}_{I_2} \end{aligned}$$

where the inequality holds due to the triangle inequality. We bound I_1 first.

$$\begin{aligned}
 I_1 &= \left| \left[\langle \hat{\boldsymbol{\mu}}_{k,h} \widehat{\mathbf{V}}_{k,h+1}^2, \boldsymbol{\phi}(s_h^k, a_h^k) \rangle \right]_{[0,H^2]} - \langle \boldsymbol{\mu}_h \mathbf{V}_{h+1}^{*2}, \boldsymbol{\phi}(s_h^k, a_h^k) \rangle \right| \\
 &= \left| \left[\langle \hat{\boldsymbol{\mu}}_{k,h} \widehat{\mathbf{V}}_{k,h+1}^2, \boldsymbol{\phi}(s_h^k, a_h^k) \rangle \right]_{[0,H^2]} - \langle \boldsymbol{\mu}_h \widehat{\mathbf{V}}_{k,h+1}^2, \boldsymbol{\phi}(s_h^k, a_h^k) \rangle \right. \\
 &\quad \left. + \langle \boldsymbol{\mu}_h \widehat{\mathbf{V}}_{k,h+1}^2, \boldsymbol{\phi}(s_h^k, a_h^k) \rangle - \langle \boldsymbol{\mu}_h \mathbf{V}_{h+1}^{*2}, \boldsymbol{\phi}(s_h^k, a_h^k) \rangle \right| \\
 &\leq \left| \left[\langle \hat{\boldsymbol{\mu}}_{k,h} \widehat{\mathbf{V}}_{k,h+1}^2, \boldsymbol{\phi}(s_h^k, a_h^k) \rangle \right]_{[0,H^2]} - \langle \boldsymbol{\mu}_h \widehat{\mathbf{V}}_{k,h+1}^2, \boldsymbol{\phi}(s_h^k, a_h^k) \rangle \right| \\
 &\quad + \left| \langle \boldsymbol{\mu}_h \widehat{\mathbf{V}}_{k,h+1}^2, \boldsymbol{\phi}(s_h^k, a_h^k) \rangle - \langle \boldsymbol{\mu}_h \mathbf{V}_{h+1}^{*2}, \boldsymbol{\phi}(s_h^k, a_h^k) \rangle \right| \\
 &\leq \left| \langle (\hat{\boldsymbol{\mu}}_{k,h} - \boldsymbol{\mu}_h) \widehat{\mathbf{V}}_{k,h+1}^2, \boldsymbol{\phi}(s_h^k, a_h^k) \rangle \right| + 2H \left| \langle \boldsymbol{\mu}_h \widehat{\mathbf{V}}_{k,h+1}, \boldsymbol{\phi}(s_h^k, a_h^k) \rangle - \langle \boldsymbol{\mu}_h \mathbf{V}_{h+1}^*, \boldsymbol{\phi}(s_h^k, a_h^k) \rangle \right| \\
 &\leq \left| \langle (\hat{\boldsymbol{\mu}}_{k,h} - \boldsymbol{\mu}_h) \widehat{\mathbf{V}}_{k,h+1}^2, \boldsymbol{\phi}(s_h^k, a_h^k) \rangle \right| + 2H \left[\left| \langle (\hat{\boldsymbol{\mu}}_{k,h} - \boldsymbol{\mu}_h) \widehat{\mathbf{V}}_{k,h+1}, \boldsymbol{\phi}(s_h^k, a_h^k) \rangle \right| \right. \\
 &\quad \left. + \left| \langle \hat{\boldsymbol{\mu}}_{k,h} (\widehat{\mathbf{V}}_{k,h+1} - \check{\mathbf{V}}_{k,h+1}), \boldsymbol{\phi}(s_h^k, a_h^k) \rangle \right| + \left| \langle (\hat{\boldsymbol{\mu}}_{k,h} - \boldsymbol{\mu}_h) \check{\mathbf{V}}_{k,h+1}, \boldsymbol{\phi}(s_h^k, a_h^k) \rangle \right| \right] \\
 &\leq \left\| (\hat{\boldsymbol{\mu}}_{k,h} - \boldsymbol{\mu}_h) \widehat{\mathbf{V}}_{k,h+1}^2 \right\|_{\widehat{\boldsymbol{\Lambda}}_{k,h}} \left\| \boldsymbol{\phi}(s_h^k, a_h^k) \right\|_{\widehat{\boldsymbol{\Lambda}}_{k,h}^{-1}} + 2H \left[\left\| (\hat{\boldsymbol{\mu}}_{k,h} - \boldsymbol{\mu}_h) \widehat{\mathbf{V}}_{k,h+1} \right\|_{\widehat{\boldsymbol{\Lambda}}_{k,h}} \left\| \boldsymbol{\phi}(s_h^k, a_h^k) \right\|_{\widehat{\boldsymbol{\Lambda}}_{k,h}^{-1}} \right. \\
 &\quad \left. + \left| \langle \hat{\boldsymbol{\mu}}_{k,h} (\widehat{\mathbf{V}}_{k,h+1} - \check{\mathbf{V}}_{k,h+1}), \boldsymbol{\phi}(s_h^k, a_h^k) \rangle \right| + \left\| (\hat{\boldsymbol{\mu}}_{k,h} - \boldsymbol{\mu}_h) \check{\mathbf{V}}_{k,h+1} \right\|_{\widehat{\boldsymbol{\Lambda}}_{k,h}} \left\| \boldsymbol{\phi}(s_h^k, a_h^k) \right\|_{\widehat{\boldsymbol{\Lambda}}_{k,h}^{-1}} \right]
 \end{aligned} \tag{25}$$

where the first inequality holds due to the triangle inequality, the second inequality holds since $\mathbb{P}_h(\cdot \mid s_h^k, a_h^k) = \langle \boldsymbol{\mu}_h(\cdot), \boldsymbol{\phi}(s_h^k, a_h^k) \rangle$ is valid distribution and $\widehat{\mathbf{V}}_{k,h+1}(\cdot), \mathbf{V}_{h+1}^*(\cdot) \in [0, H]$, the third inequality holds due to Lemma C.12 under $\widehat{\Psi}_{k,h+1} \cap \check{\Psi}_{k,h+1}$, and the last inequality holds due to the Cauchy-Schwarz inequality.

For I_2 , we have

$$\begin{aligned}
 I_2 &= \left| \left\{ \left[\langle \hat{\boldsymbol{\mu}}_{k,h} \widehat{\mathbf{V}}_{k,h+1}, \boldsymbol{\phi}(s_h^k, a_h^k) \rangle \right]_{[0,H]} \right\}^2 - \left[\langle \boldsymbol{\mu}_h \mathbf{V}_{h+1}^*, \boldsymbol{\phi}(s_h^k, a_h^k) \rangle \right]^2 \right| \\
 &= \left| \left[\langle \hat{\boldsymbol{\mu}}_{k,h} \widehat{\mathbf{V}}_{k,h+1}, \boldsymbol{\phi}(s_h^k, a_h^k) \rangle \right]_{[0,H]} + \langle \boldsymbol{\mu}_h \mathbf{V}_{h+1}^*, \boldsymbol{\phi}(s_h^k, a_h^k) \rangle \right| \\
 &\quad \cdot \left| \left[\langle \hat{\boldsymbol{\mu}}_{k,h} \widehat{\mathbf{V}}_{k,h+1}, \boldsymbol{\phi}(s_h^k, a_h^k) \rangle \right]_{[0,H]} - \langle \boldsymbol{\mu}_h \mathbf{V}_{h+1}^*, \boldsymbol{\phi}(s_h^k, a_h^k) \rangle \right| \\
 &\leq 2H \left| \left[\langle \hat{\boldsymbol{\mu}}_{k,h} \widehat{\mathbf{V}}_{k,h+1}, \boldsymbol{\phi}(s_h^k, a_h^k) \rangle \right]_{[0,H]} - \langle \boldsymbol{\mu}_h \mathbf{V}_{h+1}^*, \boldsymbol{\phi}(s_h^k, a_h^k) \rangle \right| \\
 &\leq 2H \left[\left| \langle (\hat{\boldsymbol{\mu}}_{k,h} - \boldsymbol{\mu}_h) \widehat{\mathbf{V}}_{k,h+1}, \boldsymbol{\phi}(s_h^k, a_h^k) \rangle \right| + \left| \langle \hat{\boldsymbol{\mu}}_{k,h} (\widehat{\mathbf{V}}_{k,h+1} - \check{\mathbf{V}}_{k,h+1}), \boldsymbol{\phi}(s_h^k, a_h^k) \rangle \right| \right. \\
 &\quad \left. + \left| \langle (\hat{\boldsymbol{\mu}}_{k,h} - \boldsymbol{\mu}_h) \check{\mathbf{V}}_{k,h+1}, \boldsymbol{\phi}(s_h^k, a_h^k) \rangle \right| \right] \\
 &\leq 2H \left[\left\| (\hat{\boldsymbol{\mu}}_{k,h} - \boldsymbol{\mu}_h) \widehat{\mathbf{V}}_{k,h+1} \right\|_{\widehat{\boldsymbol{\Lambda}}_{k,h}} \left\| \boldsymbol{\phi}(s_h^k, a_h^k) \right\|_{\widehat{\boldsymbol{\Lambda}}_{k,h}^{-1}} + \left| \langle \hat{\boldsymbol{\mu}}_{k,h} (\widehat{\mathbf{V}}_{k,h+1} - \check{\mathbf{V}}_{k,h+1}), \boldsymbol{\phi}(s_h^k, a_h^k) \rangle \right| \right. \\
 &\quad \left. + \left\| (\hat{\boldsymbol{\mu}}_{k,h} - \boldsymbol{\mu}_h) \check{\mathbf{V}}_{k,h+1} \right\|_{\widehat{\boldsymbol{\Lambda}}_{k,h}} \left\| \boldsymbol{\phi}(s_h^k, a_h^k) \right\|_{\widehat{\boldsymbol{\Lambda}}_{k,h}^{-1}} \right]
 \end{aligned} \tag{26}$$

where the first inequality holds since $0 \leq \left[\langle \hat{\boldsymbol{\mu}}_{k,h} \widehat{\mathbf{V}}_{k,h+1}, \boldsymbol{\phi}(s_h^k, a_h^k) \rangle \right]_{[0,H]} \leq H$, $0 \leq \mathbb{P}_h \mathbf{V}_{h+1}^*(s_h^k, a_h^k) = \langle \boldsymbol{\mu}_h \mathbf{V}_{h+1}^*, \boldsymbol{\phi}(s_h^k, a_h^k) \rangle \leq H$, the second inequality holds due to Lemma C.12 under $\widehat{\Psi}_{k,h+1} \cap \check{\Psi}_{k,h+1}$, and the third inequality holds due to the Cauchy-Schwarz inequality. Combining Eq. (25) and (26) and using the fact that I_1 and I_2 are both bounded by H^2 give the final result. \square

C.2.2. VARIANCE OF $[\hat{V}_{k,h+1} - V_{h+1}^*](\cdot)$

Lemma C.14. *In Algorithm 1, for any $k \in [K]$ and any $h \in [H]$, under $\hat{\Psi}_{k,h+1} \cap \check{\Psi}_{k,h+1}$, we have $V_{h+1}^*(\cdot) \leq \hat{V}_{k,h+1}(\cdot)$. Moreover, for any function $V : \mathcal{S} \mapsto [0, H]$ satisfying $V_{h+1}^*(\cdot) \leq V(\cdot) \leq \hat{V}_{k,h+1}(\cdot)$, we have*

$$\begin{aligned} [\mathbb{V}_h(V - V_{h+1}^*)](s_h^k, a_h^k) &\leq \min \left\{ H \left[\langle \hat{\mu}_{k,h} \hat{V}_{k,h+1}, \phi(s_h^k, a_h^k) \rangle - \langle \hat{\mu}_{k,h} \check{V}_{k,h+1}, \phi(s_h^k, a_h^k) \rangle \right. \right. \\ &\quad \left. \left. + \left\| (\hat{\mu}_{k,h} - \mu_h) \hat{V}_{k,h+1} \right\|_{\hat{\Lambda}_{k,h}} \left\| \phi(s_h^k, a_h^k) \right\|_{\hat{\Lambda}_{k,h}^{-1}} \right] \right. \\ &\quad \left. + \left\| (\hat{\mu}_{k,h} - \mu_h) \check{V}_{k,h+1} \right\|_{\hat{\Lambda}_{k,h}} \left\| \phi(s_h^k, a_h^k) \right\|_{\hat{\Lambda}_{k,h}^{-1}} \right], H^2 \}. \end{aligned}$$

Proof. Initially, we have $V_{h+1}^*(\cdot) \leq \hat{V}_{k,h+1}(\cdot) \leq \hat{V}_{k,h+1}(\cdot)$, where the first inequality holds by Lemma D.1 under $\hat{\Psi}_{k,h+1}$, the second inequality holds by definitions of $\hat{V}_{k,h+1}(\cdot)$ and $\check{V}_{k,h+1}(\cdot)$ in Algorithm 1.

Denote $\tilde{V}(\cdot) = V(\cdot) - V_{h+1}^*(\cdot)$. By definition of the variance, we have

$$\begin{aligned} [\mathbb{V}_h(V - V_{h+1}^*)](s_h^k, a_h^k) &= [\mathbb{V}_h \tilde{V}](s_h^k, a_h^k) = \mathbb{P}_h \tilde{V}^2(s_h^k, a_h^k) - [\mathbb{P}_h \tilde{V}(s_h^k, a_h^k)]^2 \leq \mathbb{P}_h \tilde{V}^2(s_h^k, a_h^k) \\ &\leq H \cdot \mathbb{P}_h \tilde{V}(s_h^k, a_h^k) \leq H \cdot \left[\mathbb{P}_h \hat{V}_{k,h+1}(s_h^k, a_h^k) - \mathbb{P}_h V_{h+1}^*(s_h^k, a_h^k) \right] \leq H \cdot \left[\mathbb{P}_h \hat{V}_{k,h+1}(s_h^k, a_h^k) - \mathbb{P}_h \check{V}_{k,h+1}(s_h^k, a_h^k) \right], \end{aligned} \quad (27)$$

where the first inequality holds trivially, the second inequality holds since $V_{h+1}^*(\cdot) \leq V(\cdot) \leq \hat{V}_{k,h+1}(\cdot) \leq H$, the third inequality holds since $V(\cdot) \leq \hat{V}_{k,h+1}(\cdot)$, and the last inequality holds since $\check{V}_{k,h+1}(\cdot) \leq V_{h+1}^*(\cdot)$ under $\check{\Psi}_{k,h+1}$ by Lemma D.1.

On the other hand, by Cauchy-Schwarz inequality, we have

$$\left| \mathbb{P}_h \hat{V}_{k,h+1}(s_h^k, a_h^k) - \hat{\mathbb{P}}_{k,h} \hat{V}_{k,h+1}(s_h^k, a_h^k) \right| \leq \left\| (\hat{\mu}_{k,h} - \mu_h) \hat{V}_{k,h+1} \right\|_{\hat{\Lambda}_{k,h}} \left\| \phi(s_h^k, a_h^k) \right\|_{\hat{\Lambda}_{k,h}^{-1}} \quad (28)$$

$$\left| \mathbb{P}_h \check{V}_{k,h+1}(s_h^k, a_h^k) - \hat{\mathbb{P}}_{k,h} \check{V}_{k,h+1}(s_h^k, a_h^k) \right| \leq \left\| (\hat{\mu}_{k,h} - \mu_h) \check{V}_{k,h+1} \right\|_{\hat{\Lambda}_{k,h}} \left\| \phi(s_h^k, a_h^k) \right\|_{\hat{\Lambda}_{k,h}^{-1}} \quad (29)$$

Combining Eq (27), Eq (28), Eq (29) and using the fact that $[\mathbb{V}_h(V - V_{h+1}^*)](s_h^k, a_h^k) \leq H^2$ give the final result. \square

C.3. Dependent Confidence Sets

Based on independent confidence sets $\bar{\mathcal{C}}_{k,h}, \hat{\mathcal{C}}_{k,h}, \check{\mathcal{C}}_{k,h}, \check{\mathcal{C}}_{k,h}$ built above and Lemma C.13, C.14, dependent confidence sets $\hat{\mathcal{C}}_{k,h}^{(1)}, \hat{\mathcal{C}}_{k,h}^{(2)}$ are built in Lemma C.16, C.17, respectively. As a results, the confidence set $\hat{\mathcal{C}}_{k,h}$, the goal of this section, holds trivially if $\hat{\mathcal{C}}_{k,h}^{(1)}, \hat{\mathcal{C}}_{k,h}^{(2)}$ both hold. We build confidence sets $\hat{\mathcal{C}}_{k,h}^{(1)}$ and $\hat{\mathcal{C}}_{k,h}^{(2)}$ elegantly because the radius of the confidence set $\hat{\mathcal{C}}_{k,h}$ will exactly determine the sharpness of the regret obtained by LSVI-UCB⁺ algorithm. In particular, we utilize the conservatism of elliptical potentials, which is detailed in Remark 7.4 in the main paper. To formally utilize this property, we first present Lemma C.15 to keep the magnitude of the considered MDS small with the conservatism of elliptical potentials.

Lemma C.15. *In Algorithm 1, for any $k \in [K]$ and any $h \in [H]$, we have*

$$\hat{\sigma}_{k,h}^{-1} \cdot \min \left\{ 1, \left\| \hat{\sigma}_{k,h}^{-1} \phi(s_h^i, a_h^i) \right\|_{\hat{\Lambda}_{k,h}^{-1}} \right\} \leq \frac{1}{H\sqrt{d^3}}.$$

Proof. In Algorithm 1, for any $k \in [K]$ and any $h \in [H]$, we have following two cases:

- If $\left\| \hat{\sigma}_{i,h}^{-1} \phi(s_h^i, a_h^i) \right\|_{\hat{\Lambda}_{i,h}^{-1}} \leq 1/(Hd^3)$, then $s_{k,h} = \sqrt{H}$ such that $\hat{\sigma}_{i,h} = \tilde{\sigma}_{i,h}$. In this case, we have

$$\left\| \hat{\sigma}_{k,h}^{-1} \phi(s_h^i, a_h^i) \right\|_{\hat{\Lambda}_{k,h}^{-1}} = \left\| \tilde{\sigma}_{k,h}^{-1} \phi(s_h^i, a_h^i) \right\|_{\hat{\Lambda}_{k,h}^{-1}} \leq \sqrt{Hd^3} \left\| \tilde{\sigma}_{k,h}^{-1} \phi(s_h^i, a_h^i) \right\|_{\hat{\Lambda}_{k,h}^{-1}} \leq \sqrt{Hd^3}/(Hd^3) \leq 1/\sqrt{Hd^3}.$$

where the inequality holds since $Hd^3 \cdot \widehat{\Lambda}_{k,h} \geq \widetilde{\Lambda}_{k,h}$ by following facts:

$$Hd^3 \cdot \widehat{\Lambda}_{k,h} = Hd^3 \cdot \left(\sum_{i=1}^{k-1} \widehat{\sigma}_{i,h}^{-2} \phi(s_h^i, a_h^i) \phi(s_h^i, a_h^i)^\top + \lambda \mathbf{I} \right) \geq \sum_{i=1}^{k-1} \widetilde{\sigma}_{i,h}^{-2} \phi(s_h^i, a_h^i) \phi(s_h^i, a_h^i)^\top + \lambda \mathbf{I} = \widetilde{\Lambda}_{k,h},$$

where the inequality holds since $\widehat{\sigma}_{i,h} \leq \sqrt{Hd^3} \widetilde{\sigma}_{i,h}$ in Algorithm 1, which implies $Hd^3 \cdot \widehat{\Lambda}_{k,h} - \widetilde{\Lambda}_{k,h}$ is a semi-positive definite matrix.

Therefore, the conclusion holds in this case since $\widehat{\sigma}_{k,h} \geq \sqrt{H}$.

- Otherwise, $\varsigma_{k,h} = H\sqrt{d^3}$, such that $\widehat{\sigma}_{k,h} \geq H\sqrt{d^3}$. In this case, the conclusion still holds since $\min\{1, \|\widehat{\sigma}_{k,h}^{-1} \phi(s_h^i, a_h^i)\|_{\widehat{\Lambda}_{k,h}^{-1}}\} \leq 1$.

□

Now we are ready to prove Lemma C.16 which builds the dependent confidence set $\widehat{\mathcal{C}}_{k,h}^{(1)}$, based on independent confidence sets $\bar{\mathcal{C}}_{k,h} \cap \widetilde{\mathcal{C}}_{k,h} \cap \check{\mathcal{C}}_{k,h}$ and Lemma C.13. Indeed, the confidence set $\widehat{\mathcal{C}}_{k,h}^{(1)}$ corresponds to the deviation term of the form $[(\widehat{\mathbb{P}}_{k,h} - \mathbb{P}_h) V_{h+1}^*](s_h^k, a_h^k)$ in the main paper.

Lemma C.16. *In Algorithm 1, for any $\delta \in (0, 1)$, any $k \in [K]$ and fixed $h \in [H]$, under $\widehat{\Psi}_{h+1} \cap \check{\Psi}_{h+1}$, with probability at least $1 - 4\delta/H$:*

$$\boldsymbol{\mu}_h \in \widehat{\mathcal{C}}_{k,h}^{(1)} \cap \bar{\mathcal{C}}_{k,h} \cap \widetilde{\mathcal{C}}_{k,h} \cap \check{\mathcal{C}}_{k,h},$$

where

$$\begin{aligned} \widehat{\mathcal{C}}_{k,h}^{(1)} &= \left\{ \boldsymbol{\mu} : \|(\boldsymbol{\mu} - \widehat{\boldsymbol{\mu}}_{k,h}) \mathbf{V}_{h+1}^*\|_{\widehat{\Lambda}_{k,h}} \leq \widehat{\beta}^{(1)} \right\}, \\ \widehat{\beta}^{(1)} &= 8\sqrt{d \log \left(1 + \frac{K}{Hd\lambda} \right) \log \left(\frac{4K^2H}{\delta} \right) + 4 \log \left(\frac{4K^2H}{\delta} \right) + H\sqrt{\lambda d}}. \end{aligned} \quad (30)$$

Proof. Let $\mathcal{G}_i = \mathcal{F}_{i,h}$, $\mathbf{x}_i = \widehat{\sigma}_{i,h}^{-1} \phi(s_h^i, a_h^i)$, $\mathbf{Z}_i = \lambda \mathbf{I} + \sum_{j=1}^i \mathbf{x}_j \mathbf{x}_j^\top$, and $\eta_i = \widehat{\sigma}_{i,h}^{-1} \boldsymbol{\epsilon}_h^\top \mathbf{V}_{h+1}^* \cdot \mathbf{1}\{\boldsymbol{\mu}_h \in \bar{\mathcal{C}}_{i,h} \cap \widetilde{\mathcal{C}}_{i,h} \cap \check{\mathcal{C}}_{i,h}\} \cdot \mathbf{1}\{\widehat{\Psi}_{i,h+1} \cap \check{\Psi}_{i,h+1}\}$. Now that V_{h+1}^* is a fixed function and $\mathbf{1}\{\boldsymbol{\mu}_h \in \bar{\mathcal{C}}_{i,h} \cap \widetilde{\mathcal{C}}_{i,h} \cap \check{\mathcal{C}}_{i,h}\} \cdot \mathbf{1}\{\widehat{\Psi}_{i,h+1} \cap \check{\Psi}_{i,h+1}\}$ is \mathcal{G}_i -measurable, it is clear that \mathbf{x}_i are \mathcal{G}_i -measurable and η_i is \mathcal{G}_{i+1} -measurable.

Besides, we have $\mathbb{E}[\eta_i | \mathcal{G}_i] = 0$. Since $\widehat{\sigma}_{i,h} \geq \varsigma_{i,h} \geq \sqrt{H}$, $|\eta_i| \leq \sqrt{H}$ and $\|\mathbf{x}_i\|_2 \leq 1/\sqrt{H}$. In particular, we claim $|\eta_i \cdot \min\{1, \|\mathbf{x}_i\|_{\mathbf{Z}_{i-1}^{-1}}\}| \leq 1$ because of the following three facts: (i) $|\boldsymbol{\epsilon}_h^\top \mathbf{V}_{h+1}^* \cdot \mathbf{1}\{\boldsymbol{\mu}_h \in \bar{\mathcal{C}}_{i,h} \cap \widetilde{\mathcal{C}}_{i,h} \cap \check{\mathcal{C}}_{i,h}\} \cdot \mathbf{1}\{\widehat{\Psi}_{i,h+1} \cap \check{\Psi}_{i,h+1}\}| \leq H$ holds by $|V_{h+1}^* \cdot (\cdot)| \leq H$; (ii) $\widehat{\sigma}_{i,h}^{-1} \cdot \min\{1, \|\mathbf{x}_i\|_{\mathbf{Z}_{i-1}^{-1}}\} \leq 1/(H\sqrt{d^3})$ holds by Lemma C.15; and (iii) $|\mathbf{1}\{\boldsymbol{\mu}_h \in \bar{\mathcal{C}}_{i,h} \cap \widetilde{\mathcal{C}}_{i,h} \cap \check{\mathcal{C}}_{i,h}\} \cdot \mathbf{1}\{\widehat{\Psi}_{i,h+1} \cap \check{\Psi}_{i,h+1}\}| \leq 1$.

Furthermore, it holds that

$$\begin{aligned} \mathbb{E}[\eta_i^2 | \mathcal{G}_i] &= \widehat{\sigma}_{i,h}^{-2} \cdot \mathbf{1}\{\boldsymbol{\mu}_h \in \bar{\mathcal{C}}_{i,h} \cap \widetilde{\mathcal{C}}_{i,h} \cap \check{\mathcal{C}}_{i,h}\} \cdot \mathbf{1}\{\widehat{\Psi}_{i,h+1} \cap \check{\Psi}_{i,h+1}\} [\mathbb{V}_h V_{h+1}^*](s_h^i, a_h^i) \\ &\leq \widehat{\sigma}_{i,h}^{-2} \cdot \mathbf{1}\{\boldsymbol{\mu}_h \in \bar{\mathcal{C}}_{i,h} \cap \widetilde{\mathcal{C}}_{i,h} \cap \check{\mathcal{C}}_{i,h}\} \left[\left[\widehat{\mathbb{V}}_{i,h} \widehat{V}_{i,h+1} \right] (s_h^i, a_h^i) \right. \\ &\quad \left. + \min \left\{ \left\| (\widehat{\boldsymbol{\mu}}_{i,h} - \boldsymbol{\mu}_h) \widehat{\mathbf{V}}_{i,h+1} \right\|_{\widehat{\Lambda}_{i,h}} \left\| \phi(s_h^i, a_h^i) \right\|_{\widehat{\Lambda}_{i,h}^{-1}} + 4H\Delta_{i,h}, 2H^2 \right\} \right] \\ &\leq \widehat{\sigma}_{i,h}^{-2} \left[\left[\widehat{\mathbb{V}}_{i,h} \widehat{V}_{i,h+1} \right] (s_h^i, a_h^i) + \min \left\{ \widetilde{\beta} \left\| \phi(s_h^i, a_h^i) \right\|_{\widehat{\Lambda}_{i,h}^{-1}} + 4H\delta_{i,h}, 2H^2 \right\} \right] \\ &\leq 1, \end{aligned}$$

where the first inequality holds due to Lemma C.13 under $\widehat{\Psi}_{i,h+1} \cap \check{\Psi}_{i,h+1}$, the second inequality holds due to the definition of indicator function, and the last inequality holds due to the definition of $\widehat{\sigma}_{i,h}$. Here $\Delta_{i,h}$ and $\delta_{i,h}$ are given by

$$\Delta_{i,h} = \left\| (\widehat{\boldsymbol{\mu}}_{i,h} - \boldsymbol{\mu}_h) \widehat{\mathbf{V}}_{i,h+1} \right\|_{\widehat{\Lambda}_{i,h}} \left\| \phi(s_h^i, a_h^i) \right\|_{\widehat{\Lambda}_{i,h}^{-1}} + \left| \langle \widehat{\boldsymbol{\mu}}_{i,h} (\widehat{\mathbf{V}}_{i,h+1} - \check{\mathbf{V}}_{i,h+1}), \phi(s_h^i, a_h^i) \rangle \right|$$

$$\begin{aligned}
 & + \left\| (\hat{\boldsymbol{\mu}}_{i,h} - \boldsymbol{\mu}_h) \check{\mathbf{V}}_{i,h+1} \right\|_{\hat{\Lambda}_{i,h}} \left\| \phi(s_h^i, a_h^i) \right\|_{\hat{\Lambda}_{i,h}^{-1}}, \\
 \delta_{i,h} & = \bar{\beta} \left\| \phi(s_h^i, a_h^i) \right\|_{\hat{\Lambda}_{i,h}^{-1}} + \left| \langle \hat{\boldsymbol{\mu}}_{i,h} (\hat{\mathbf{V}}_{i,h+1} - \check{\mathbf{V}}_{i,h+1}), \phi(s_h^i, a_h^i) \rangle \right| + \check{\beta} \left\| \phi(s_h^i, a_h^i) \right\|_{\hat{\Lambda}_{i,h}^{-1}}.
 \end{aligned}$$

Then, by Lemma F.4, with probability at least $1 - \delta/H$, for all $k \in [K]$ and fixed $h \in [H]$,

$$\begin{aligned}
 & \left\| \sum_{i=1}^{k-1} \hat{\sigma}_{i,h}^{-2} \phi(s_h^i, a_h^i) \boldsymbol{\epsilon}_h^{i \top} \mathbf{V}_{h+1}^* \mathbb{1} \left\{ \boldsymbol{\mu}_h \in \bar{\mathcal{C}}_{i,h} \cap \check{\mathcal{C}}_{i,h} \cap \check{\check{\mathcal{C}}}_{i,h} \right\} \mathbb{1} \left\{ \hat{\Psi}_{i,h+1} \cap \check{\Psi}_{i,h+1} \right\} \right\|_{\hat{\Lambda}_{k,h}^{-1}} \\
 & \leq 8 \sqrt{d \log \left(1 + \frac{K}{Hd\lambda} \right) \log \left(\frac{4K^2H}{\delta} \right) + 4 \log \left(\frac{4K^2H}{\delta} \right)}.
 \end{aligned} \tag{31}$$

Denote $\mathcal{E}_h^{(1)}$ as the event that $\boldsymbol{\mu}_h \in \cap_{k \in [K]} \bar{\mathcal{C}}_{k,h} \cap \check{\mathcal{C}}_{k,h} \cap \check{\check{\mathcal{C}}}_{k,h}$ and Eq. (31) hold, which happens with probability at least $1 - 4\delta/H$ by taking a union bound. In addition, we claim that with probability at least $1 - 4\delta/H$, for all $k \in [K]$ and fixed $h \in [H]$,

$$\begin{aligned}
 & \left\| \sum_{i=1}^{k-1} \hat{\sigma}_{i,h}^{-2} \phi(s_h^i, a_h^i) \boldsymbol{\epsilon}_h^{i \top} \mathbf{V}_{h+1}^* \cdot \mathbb{1} \left\{ \boldsymbol{\mu}_h \in \bar{\mathcal{C}}_{i,h} \cap \check{\mathcal{C}}_{i,h} \cap \check{\check{\mathcal{C}}}_{i,h} \right\} \cdot \mathbb{1} \left\{ \hat{\Psi}_{i,h+1} \cap \check{\Psi}_{i,h+1} \right\} \right\|_{\hat{\Lambda}_{k,h}^{-1}} \\
 & = \left\| \sum_{i=1}^{k-1} \hat{\sigma}_{i,h}^{-2} \phi(s_h^i, a_h^i) \boldsymbol{\epsilon}_h^{i \top} \mathbf{V}_{h+1}^* \cdot \mathbb{1} \left\{ \hat{\Psi}_{i,h+1} \cap \check{\Psi}_{i,h+1} \right\} \right\|_{\hat{\Lambda}_{k,h}^{-1}} \\
 & \leq 8 \sqrt{d \log \left(1 + \frac{K}{Hd\lambda} \right) \log \left(\frac{4K^2H}{\delta} \right) + 4 \log \left(\frac{4K^2H}{\delta} \right)},
 \end{aligned}$$

where the equality holds since under event $\mathcal{E}_h^{(1)}$, for any $i \in [K]$, $\mathbb{1} \left\{ \boldsymbol{\mu}_h \in \bar{\mathcal{C}}_{i,h} \cap \check{\mathcal{C}}_{i,h} \cap \check{\check{\mathcal{C}}}_{i,h} \right\} = 1$. Moreover, if we further assume $\hat{\Psi}_{h+1} \cap \check{\Psi}_{h+1}$ holds, which means $\mathbb{1} \left\{ \hat{\Psi}_{i,h+1} \cap \check{\Psi}_{i,h+1} \right\} = 1$ for any $i \in [k]$, then with probability at least $1 - 4\delta/H$, for any $k \in [K]$ and fixed $h \in [H]$:

$$\left\| (\hat{\boldsymbol{\mu}}_{k,h} - \boldsymbol{\mu}_h) \mathbf{V}_{h+1}^* \right\|_{\hat{\Lambda}_{k,h}} \leq 8 \sqrt{d \log \left(1 + \frac{K}{Hd\lambda} \right) \log \left(\frac{4K^2H}{\delta} \right) + 4 \log \left(\frac{4K^2H}{\delta} \right) + H\sqrt{\lambda d}} = \hat{\beta}_k^{(1)}.$$

since

$$\left\| (\hat{\boldsymbol{\mu}}_{k,h} - \boldsymbol{\mu}_h) \mathbf{V}_{h+1}^* \right\|_{\hat{\Lambda}_{k,h}} \leq H\sqrt{\lambda d} + \left\| \sum_{i=1}^{k-1} \hat{\sigma}_{i,h}^{-2} \phi(s_h^i, a_h^i) \boldsymbol{\epsilon}_h^{i \top} \mathbf{V}_{h+1}^* \right\|_{\hat{\Lambda}_{k,h}^{-1}}$$

with a similar argument as in Eq. (21). Thus, we conclude that for any $k \in [K]$ and fixed $h \in [H]$, under $\hat{\Psi}_{h+1} \cap \check{\Psi}_{h+1}$, with probability at least $1 - 4\delta/H$:

$$\boldsymbol{\mu}_h \in \hat{\mathcal{C}}_{k,h}^{(1)} \cap \bar{\mathcal{C}}_{k,h} \cap \check{\mathcal{C}}_{k,h} \cap \check{\check{\mathcal{C}}}_{k,h}.$$

□

Subsequently, we prove Lemma C.17 which builds the dependent confidence set $\hat{\mathcal{C}}_{k,h}^{(2)}$, based on independent confidence sets $\hat{\mathcal{C}}_{k,h} \cap \check{\mathcal{C}}_{k,h}$ and Lemma C.14. The confidence set $\hat{\mathcal{C}}_{k,h}^{(2)}$ corresponds to the deviation term of the form $[(\hat{\mathbb{P}}_{k,h} - \mathbb{P}_h)(\hat{\mathbf{V}}_{k,h+1} - \mathbf{V}_{h+1}^*)](s_h^k, a_h^k)$ in main paper, which is controlled to be small in LSVI-UCB⁺.

Lemma C.17. *In Algorithm 1, for any $\delta \in (0, 1)$, any $k \in [K]$ and fixed $h \in [H]$, under $\hat{\Psi}_{h+1} \cap \check{\Psi}_{h+1}$, with probability at least $1 - 3\delta/H$:*

$$\boldsymbol{\mu}_h \in \hat{\mathcal{C}}_{k,h}^{(2)} \cap \hat{\mathcal{C}}_{k,h} \cap \check{\mathcal{C}}_{k,h}.$$

where

$$\begin{aligned} \hat{\mathcal{C}}_{k,h}^{(2)} &= \left\{ \boldsymbol{\mu} : \left\| (\boldsymbol{\mu} - \hat{\boldsymbol{\mu}}_{k,h}) \left(\hat{\mathbf{V}}_{k,h+1} - \mathbf{V}_{h+1}^* \right) \right\|_{\hat{\boldsymbol{\Lambda}}_{k,h}} \leq \hat{\beta}^{(2)} \right\}, \\ \hat{\beta}^{(2)} &= 8 \sqrt{\frac{1}{d} \log \left(1 + \frac{K}{Hd\lambda} \right) \left[\log \left(\frac{4K^2H}{\delta} \right) + d \log \left(1 + \frac{4KL}{H\sqrt{\lambda}} \right) + d^2 \log \left(1 + \frac{8K^2\hat{B}^2\sqrt{d}}{H^2\lambda^2} \right) \right]} \\ &\quad + \frac{4}{\sqrt{d^3}} \left[\log \left(\frac{4K^2H}{\delta} \right) + d \log \left(1 + \frac{4KL}{H\sqrt{\lambda}} \right) + d^2 \log \left(1 + \frac{8K^2\hat{B}^2\sqrt{d}}{H^2\lambda^2} \right) \right] + H\sqrt{\lambda d} + 2. \end{aligned} \quad (32)$$

Here $L = W + K/\lambda$ and \hat{B} is a constant satisfying $\hat{\beta} \leq \hat{B}$ with $\hat{\beta}$ given in Lemma C.18.

Proof. It suffices to upper bound $\left\| \sum_{i=1}^{k-1} \hat{\sigma}_{i,h}^{-2} \phi(s_h^i, a_h^i) \boldsymbol{\epsilon}_h^{i\top} \left(\hat{\mathbf{V}}_{k,h+1} - \mathbf{V}_{h+1}^* \right) \right\|_{\hat{\boldsymbol{\Lambda}}_{k,h}^{-1}}$ with a similar argument as Eq. (21). Besides, we need to build a uniform convergence argument by covering net since $\hat{\mathbf{V}}_{k,h+1}(\cdot)$ is $\mathcal{F}_{k-1,h}$ -measurable. As stated in the proof of Lemma C.8, $\hat{\mathbf{V}}_{k,h} \in \hat{\mathcal{V}}$, where $\hat{\mathcal{V}}$ is defined in Definition C.5, with $L = W + K/\lambda$ and $B = \hat{B}$. Here, \hat{B} is a constant satisfying $\hat{\beta} \leq \hat{B}$ with $\hat{\beta}$ specified in Lemma C.18.

Then, for a fixed function $V(\cdot) \in \hat{\mathcal{V}} : \mathcal{S} \mapsto [0, H]$ and a constant $\zeta > 0$, let $\mathcal{G}_i = \mathcal{F}_{i,h}$, $\mathbf{x}_i = \hat{\sigma}_{i,h}^{-1} \phi(s_h^i, a_h^i)$, $\mathbf{Z}_i = \lambda \mathbf{I} + \sum_{j=1}^i \mathbf{x}_j \mathbf{x}_j^\top$, and $\eta_i = \hat{\sigma}_{i,h}^{-1} \boldsymbol{\epsilon}_h^{i\top} (\mathbf{V} - \mathbf{V}_{h+1}^*) \cdot \mathbb{1}\{\boldsymbol{\mu}_h \in \hat{\mathcal{C}}_{i,h} \cap \check{\mathcal{C}}_{i,h}\} \cdot \mathbb{1}\{\mathbf{V}_{h+1}^* - \zeta \leq \mathbf{V} \leq \hat{\mathbf{V}}_{i,h+1} + \zeta\} \cdot \mathbb{1}\{\hat{\Psi}_{i,h+1} \cap \check{\Psi}_{i,h+1}\}$.

Since $V(\cdot)$ and $V(\cdot)_{h+1}^*$ are fixed functions, and $\hat{\mathbf{V}}_{i,h+1}$ and $\mathbb{1}\{\boldsymbol{\mu}_h \in \hat{\mathcal{C}}_{i,h} \cap \check{\mathcal{C}}_{i,h}\} \cdot \mathbb{1}\{\hat{\Psi}_{i,h+1} \cap \check{\Psi}_{i,h+1}\}$ are \mathcal{G}_i -measurable, it is clear that \mathbf{x}_i is \mathcal{G}_i -measurable and η_i is \mathcal{G}_{i+1} -measurable. Besides, we have $\mathbb{E}[\eta_i | \mathcal{G}_i] = 0$. Since $\hat{\sigma}_{i,h} \geq \varsigma_{i,h} \geq \sqrt{H}$, $|\eta_i| \leq \sqrt{H}$ and $\|\mathbf{x}_i\|_2 \leq 1/\sqrt{H}$.

Similar to the proof in Lemma C.16, we claim $|\eta_i \cdot \min\{1, \|\mathbf{x}_i\|_{\mathbf{Z}_{i-1}^{-1}}\}| \leq 1/\sqrt{d^3}$ because of the following three facts:

- (i) $|\boldsymbol{\epsilon}_h^{i\top} (\mathbf{V} - \mathbf{V}_{h+1}^*) \cdot \mathbb{1}\{\mathbf{V}_{h+1}^* - \zeta \leq \mathbf{V} \leq \hat{\mathbf{V}}_{i,h+1} + \zeta\} \cdot \mathbb{1}\{\hat{\Psi}_{i,h+1} \cap \check{\Psi}_{i,h+1}\}| \leq H$ holds by $|(V(\cdot) - V_{h+1}^*(\cdot)) \cdot \mathbb{1}\{\mathbf{V}_{h+1}^* - \zeta \leq \mathbf{V} \leq \hat{\mathbf{V}}_{i,h+1} + \zeta\}| \leq H$; (ii) $\hat{\sigma}_{i,h}^{-1} \cdot \min\{1, \|\mathbf{x}_i\|_{\mathbf{Z}_{i-1}^{-1}}\} \leq 1/(H\sqrt{d^3})$ holds by Lemma C.15; and (iii) $|\mathbb{1}\{\boldsymbol{\mu}_h \in \hat{\mathcal{C}}_{i,h} \cap \check{\mathcal{C}}_{i,h}\} \cdot \mathbb{1}\{\hat{\Psi}_{i,h+1} \cap \check{\Psi}_{i,h+1}\}| \leq 1$.

Furthermore, it holds that

$$\begin{aligned} \mathbb{E}[\eta_i^2 | \mathcal{G}_i] &= \hat{\sigma}_{i,h}^{-2} \cdot \mathbb{1}\{\boldsymbol{\mu}_h \in \hat{\mathcal{C}}_{i,h} \cap \check{\mathcal{C}}_{i,h}\} \cdot \mathbb{1}\{\mathbf{V}_{h+1}^* - \zeta \leq \mathbf{V} \leq \hat{\mathbf{V}}_{i,h+1} + \zeta\} \cdot \mathbb{1}\{\hat{\Psi}_{i,h+1} \cap \check{\Psi}_{i,h+1}\} \\ &\quad \cdot [\mathbf{V}_h(V - V_{h+1}^*)(s_h^i, a_h^i)] \\ &\leq \hat{\sigma}_{i,h}^{-2} \cdot \mathbb{1}\{\boldsymbol{\mu}_h \in \hat{\mathcal{C}}_{i,h} \cap \check{\mathcal{C}}_{i,h}\} \cdot H \left[\hat{\mathbb{P}}_{k,h} \hat{\mathbf{V}}_{k,h+1}(s_h^k, a_h^k) - \hat{\mathbb{P}}_{k,h} \check{\mathbf{V}}_{k,h+1}(s_h^k, a_h^k) \right. \\ &\quad \left. + \left\| (\hat{\boldsymbol{\mu}}_{k,h} - \boldsymbol{\mu}_h) \hat{\mathbf{V}}_{k,h+1} \right\|_{\hat{\boldsymbol{\Lambda}}_{k,h}} \left\| \phi(s_h^k, a_h^k) \right\|_{\hat{\boldsymbol{\Lambda}}_{k,h}^{-1}} \right. \\ &\quad \left. + \left\| (\hat{\boldsymbol{\mu}}_{k,h} - \boldsymbol{\mu}_h) \check{\mathbf{V}}_{k,h+1} \right\|_{\hat{\boldsymbol{\Lambda}}_{k,h}} \left\| \phi(s_h^k, a_h^k) \right\|_{\hat{\boldsymbol{\Lambda}}_{k,h}^{-1}} \right] \\ &\leq \hat{\sigma}_{i,h}^{-2} \cdot H \left[\hat{\mathbb{P}}_{k,h} \hat{\mathbf{V}}_{k,h+1}(s_h^k, a_h^k) - \hat{\mathbb{P}}_{k,h} \check{\mathbf{V}}_{k,h+1}(s_h^k, a_h^k) + \hat{\beta} \left\| \phi(s_h^k, a_h^k) \right\|_{\hat{\boldsymbol{\Lambda}}_{k,h}^{-1}} + \check{\beta} \left\| \phi(s_h^k, a_h^k) \right\|_{\hat{\boldsymbol{\Lambda}}_{k,h}^{-1}} \right] \\ &\leq \frac{1}{d}, \end{aligned}$$

where the first inequality holds due to Lemma C.14 under $\hat{\Psi}_{i,h+1} \cap \check{\Psi}_{i,h+1}$, the second inequality holds due to the definition of indicator function, and the last inequality holds due to the definition of $\hat{\sigma}_{i,h}$.

Then, by Lemma F.4, for all $k \in [K]$ and fixed $h \in [H]$, with probability at least $1 - \delta/H$:

$$\begin{aligned} &\left\| \sum_{i=1}^{k-1} \hat{\sigma}_{i,h}^{-2} \phi(s_h^i, a_h^i) \boldsymbol{\epsilon}_h^{i\top} (\mathbf{V} - \mathbf{V}_{h+1}^*) \right. \\ &\quad \left. \cdot \mathbb{1}\{\boldsymbol{\mu}_h \in \hat{\mathcal{C}}_{i,h} \cap \check{\mathcal{C}}_{i,h}\} \cdot \mathbb{1}\{\mathbf{V}_{h+1}^* - \zeta \leq \mathbf{V} \leq \hat{\mathbf{V}}_{i,h+1} + \zeta\} \cdot \mathbb{1}\{\hat{\Psi}_{i,h+1} \cap \check{\Psi}_{i,h+1}\} \right\|_{\hat{\boldsymbol{\Lambda}}_{k,h}^{-1}} \end{aligned}$$

$$\leq 8\sqrt{\frac{1}{d} \log\left(1 + \frac{K}{Hd\lambda}\right) \log\left(\frac{4K^2H}{\delta}\right) + \frac{4}{\sqrt{d^3}} \log\left(\frac{4K^2H}{\delta}\right)}.$$

We further proceed our proof under the event that $\hat{\Psi}_{h+1} \cap \check{\Psi}_{h+1}$ holds, which implies $\mathbb{1}\{\hat{\Psi}_{i,h+1} \cap \check{\Psi}_{i,h+1}\} = 1$ for any $i \in [k]$. Denote $\mathcal{E}_h^{(2)}$ as the event that $\boldsymbol{\mu}_h \in \cap_{k \in [K]} \hat{\mathcal{C}}_{k,h} \cap \check{\mathcal{C}}_{k,h}$ and the above inequality holds, which happens with probability at least $1 - 3\delta/H$ by taking a union bound. In addition, we claim that under $\hat{\Psi}_{h+1} \cap \check{\Psi}_{h+1}$, with probability at least $1 - 3\delta/H$, for all $k \in [K]$ and fixed $h \in [H]$,

$$\begin{aligned} & \left\| \sum_{i=1}^{k-1} \hat{\sigma}_{i,h}^{-2} \phi(s_h^i, a_h^i) \boldsymbol{\epsilon}_h^i \top (\mathbf{V} - \mathbf{V}_{h+1}^*) \right. \\ & \quad \cdot \mathbb{1}\{\boldsymbol{\mu}_h \in \hat{\mathcal{C}}_{i,h} \cap \check{\mathcal{C}}_{i,h}\} \cdot \mathbb{1}\{\mathbf{V}_{h+1}^* - \zeta \leq \mathbf{V} \leq \hat{\mathbf{V}}_{i,h+1} + \zeta\} \cdot \mathbb{1}\{\hat{\Psi}_{i,h+1} \cap \check{\Psi}_{i,h+1}\} \left. \right\|_{\hat{\Lambda}_{k,h}^{-1}} \\ &= \left\| \sum_{i=1}^{k-1} \hat{\sigma}_{i,h}^{-2} \phi(s_h^i, a_h^i) \boldsymbol{\epsilon}_h^i \top (\mathbf{V} - \mathbf{V}_{h+1}^*) \cdot \mathbb{1}\{\mathbf{V}_{h+1}^* - \zeta \leq \mathbf{V} \leq \hat{\mathbf{V}}_{i,h+1} + \zeta\} \right\|_{\hat{\Lambda}_{k,h}^{-1}} \\ &\leq 8\sqrt{\frac{1}{d} \log\left(1 + \frac{K}{Hd\lambda}\right) \log\left(\frac{4K^2H}{\delta}\right) + \frac{4}{\sqrt{d^3}} \log\left(\frac{4K^2H}{\delta}\right)}, \end{aligned}$$

where the equality holds since under event $\mathcal{E}_h^{(2)} \cap \hat{\Psi}_{h+1} \cap \check{\Psi}_{h+1}$, for any $i \in [k]$, $\mathbb{1}\{\boldsymbol{\mu}_h \in \hat{\mathcal{C}}_{i,h} \cap \check{\mathcal{C}}_{i,h}\} \cdot \mathbb{1}\{\hat{\Psi}_{i,h+1} \cap \check{\Psi}_{i,h+1}\} = 1$.

Denote the ε -cover of function class $\hat{\mathcal{V}}$ as $\hat{\mathcal{N}}_\varepsilon$. Since $\hat{\mathbf{V}}_{k,h+1}(\cdot) \in \hat{\mathcal{V}}$, for any $\hat{\mathbf{V}}_{k,h+1}$, there exists a $\mathbf{V}' \in \hat{\mathcal{N}}_\varepsilon$, such that $\|\hat{\mathbf{V}}_{k,h+1} - \mathbf{V}'\|_\infty \leq \varepsilon$. This implies $\mathbf{V}_{h+1}^* - \varepsilon \leq \hat{\mathbf{V}}_{k,h+1} - \varepsilon \leq \mathbf{V}' \leq \hat{\mathbf{V}}_{k,h+1} + \varepsilon \leq \hat{\mathbf{V}}_{i,h+1} + \varepsilon$, where the first inequality holds by Lemma D.1 under $\hat{\Psi}_{h+1}$, and the last inequality holds by Lemma D.2 under $\hat{\Psi}_{h+1}$.

In addition, setting $\zeta = \varepsilon$ makes $\mathbb{1}\{\mathbf{V}_{h+1}^* - \zeta \leq \mathbf{V}' \leq \hat{\mathbf{V}}_{i,h+1} + \zeta\} = 1$ for any $i \in [k]$. Moreover, since $\|\boldsymbol{\epsilon}_h^i \top (\hat{\mathbf{V}}_{k,h+1} - \mathbf{V}')\|_2 \leq \|\boldsymbol{\epsilon}_h^i\|_1 \|\hat{\mathbf{V}}_{k,h+1} - \mathbf{V}'\|_\infty = 2\varepsilon$ and $\|\sum_{i=1}^{k-1} \hat{\sigma}_{i,h}^{-2} \phi(s_h^i, a_h^i)\|_{\hat{\Lambda}_{k,h}^{-1}} \leq K/(H\sqrt{\lambda})$, we have

$$\left\| \sum_{i=1}^{k-1} \hat{\sigma}_{i,h}^{-2} \phi(s_h^i, a_h^i) \boldsymbol{\epsilon}_h^i \top (\hat{\mathbf{V}}_{k,h+1} - \mathbf{V}') \cdot \mathbb{1}\{\mathbf{V}_{h+1}^* - \zeta \leq \mathbf{V}' \leq \hat{\mathbf{V}}_{i,h+1} + \zeta\} \right\|_{\hat{\Lambda}_{k,h}^{-1}} \leq \frac{2\varepsilon K}{H\sqrt{\lambda}}. \quad (33)$$

This further implies that the following inequality holds with probability at least $1 - 3\delta/H$:

$$\begin{aligned} & \left\| \sum_{i=1}^{k-1} \hat{\sigma}_{i,h}^{-2} \phi(s_h^i, a_h^i) \boldsymbol{\epsilon}_h^i \top (\hat{\mathbf{V}}_{k,h+1} - \mathbf{V}_{h+1}^*) \right\|_{\hat{\Lambda}_{k,h}^{-1}} \\ &= \left\| \sum_{i=1}^{k-1} \hat{\sigma}_{i,h}^{-2} \phi(s_h^i, a_h^i) \boldsymbol{\epsilon}_h^i \top (\hat{\mathbf{V}}_{k,h+1} - \mathbf{V}_{h+1}^*) \cdot \mathbb{1}\{\mathbf{V}_{h+1}^* - \zeta \leq \mathbf{V}' \leq \hat{\mathbf{V}}_{i,h+1} + \zeta\} \right\|_{\hat{\Lambda}_{k,h}^{-1}} \\ &\leq \left\| \sum_{i=1}^{k-1} \hat{\sigma}_{i,h}^{-2} \phi(s_h^i, a_h^i) \boldsymbol{\epsilon}_h^i \top (\mathbf{V}' - \mathbf{V}_{h+1}^*) \cdot \mathbb{1}\{\mathbf{V}_{h+1}^* - \zeta \leq \mathbf{V}' \leq \hat{\mathbf{V}}_{i,h+1} + \zeta\} \right\|_{\hat{\Lambda}_{k,h}^{-1}} \\ & \quad + \left\| \sum_{i=1}^{k-1} \hat{\sigma}_{i,h}^{-2} \phi(s_h^i, a_h^i) \boldsymbol{\epsilon}_h^i \top (\hat{\mathbf{V}}_{k,h+1} - \mathbf{V}') \cdot \mathbb{1}\{\mathbf{V}_{h+1}^* - \zeta \leq \mathbf{V}' \leq \hat{\mathbf{V}}_{i,h+1} + \zeta\} \right\|_{\hat{\Lambda}_{k,h}^{-1}} \\ &\leq \left\| \sum_{i=1}^{k-1} \hat{\sigma}_{i,h}^{-2} \phi(s_h^i, a_h^i) \boldsymbol{\epsilon}_h^i \top (\mathbf{V}' - \mathbf{V}_{h+1}^*) \cdot \mathbb{1}\{\mathbf{V}_{h+1}^* - \zeta \leq \mathbf{V}' \leq \hat{\mathbf{V}}_{i,h+1} + \zeta\} \right\|_{\hat{\Lambda}_{k,h}^{-1}} + \frac{2\varepsilon K}{H\sqrt{\lambda}} \\ &\leq 8\sqrt{\frac{1}{d} \log\left(1 + \frac{K}{Hd\lambda}\right) \left[\log\left(\frac{4K^2H}{\delta}\right) + \log|\hat{\mathcal{N}}_\varepsilon| \right] + \frac{4}{\sqrt{d^3}} \left[\log\left(\frac{4K^2H}{\delta}\right) + \log|\hat{\mathcal{N}}_\varepsilon| \right] + \frac{2\varepsilon K}{H\sqrt{\lambda}}} \end{aligned}$$

where the first inequality is due to triangle inequality, the second inequality holds by Eq. (33), and the third inequality holds by a union bound over all functions in $\widehat{\mathcal{N}}_\varepsilon$ with

$$\log |\widehat{\mathcal{N}}_\varepsilon| \leq d \log \left(1 + \frac{4L}{\varepsilon} \right) + d^2 \log \left(1 + \frac{8\widehat{B}^2\sqrt{d}}{\lambda\varepsilon^2} \right)$$

according to Lemma F.9.

Similar to Eq. (21), for any $k \in [K]$ and fixed $h \in [H]$, under $\widehat{\Psi}_{h+1} \cap \check{\Psi}_{h+1}$, we have that, with probability at least $1 - 3\delta/H$:

$$\begin{aligned} & \left\| (\widehat{\boldsymbol{\mu}}_{k,h} - \boldsymbol{\mu}_h) \left(\widehat{\mathbf{V}}_{k,h+1} - \mathbf{V}_{h+1}^* \right) \right\|_{\widehat{\boldsymbol{\Lambda}}_{k,h}} \\ & \leq \frac{1}{\sqrt{\lambda}} \cdot \lambda H \sqrt{d} + \left\| \sum_{i=1}^{k-1} \widehat{\sigma}_{i,h}^{-2} \boldsymbol{\phi}(s_h^i, a_h^i) \boldsymbol{\epsilon}_h^{i\top} \left(\widehat{\mathbf{V}}_{k,h+1} - \mathbf{V}_{h+1}^* \right) \right\|_{\widehat{\boldsymbol{\Lambda}}_{k,h}^{-1}} \\ & \leq 8 \sqrt{\frac{1}{d} \log \left(1 + \frac{K}{Hd\lambda} \right) \left[\log \left(\frac{4K^2H}{\delta} \right) + d \log \left(1 + \frac{4KL}{H\sqrt{\lambda}} \right) + d^2 \log \left(1 + \frac{8K^2\widehat{B}^2\sqrt{d}}{H^2\lambda^2} \right) \right]} \\ & \quad + \frac{4}{\sqrt{d^3}} \left[\log \left(\frac{4K^2H}{\delta} \right) + d \log \left(1 + \frac{4KL}{H\sqrt{\lambda}} \right) + d^2 \log \left(1 + \frac{8K^2\widehat{B}^2\sqrt{d}}{H^2\lambda^2} \right) \right] + H\sqrt{\lambda d} + 2 \\ & = \widehat{\beta}^{(2)}, \end{aligned} \tag{34}$$

where the last inequality holds by the above proved self-normalized bound and setting $\varepsilon = H\sqrt{\lambda}/K$. Thus, we conclude that for any $k \in [K]$ and fixed $h \in [H]$, under $\widehat{\Psi}_{h+1} \cap \check{\Psi}_{h+1}$, with probability at least $1 - 3\delta/H$:

$$\boldsymbol{\mu}_h \in \widehat{\mathcal{C}}_{k,h}^{(2)} \cap \dot{\mathcal{C}}_{k,h} \cap \check{\mathcal{C}}_{k,h}.$$

□

C.4. Proof of Lemma 7.5

Now we are ready to prove Lemma 7.5, i.e., building the sharp confidence set $\widehat{\mathcal{C}}_{k,h}$, in the main paper, based on above building blocks including confidence sets $\bar{\mathcal{C}}_{k,h}, \dot{\mathcal{C}}_{k,h}, \check{\mathcal{C}}_{k,h}, \check{\mathcal{C}}_{k,h}, \widehat{\mathcal{C}}_{k,h}^{(1)}, \widehat{\mathcal{C}}_{k,h}^{(2)}$, and Lemma C.13 and C.14 for upper bounding variances of value functions. In the following, we present Lemma C.18, which is the full version of Lemma 7.5 in the main paper.

Lemma C.18. *Set $\widehat{\beta} = \widehat{\beta}^{(1)} + \widehat{\beta}^{(2)}$ with $\widehat{\beta}^{(1)}$ and $\widehat{\beta}^{(2)}$ given in Lemma C.16 and C.17, respectively. Then for any $\delta \in (0, 1)$, with probability at least $1 - 7\delta$, we have that simultaneously for any $k \in [K]$ and any $h \in [H]$,*

$$\boldsymbol{\mu}_h \in \widehat{\mathcal{C}}_{k,h} \cap \widehat{\mathcal{C}}_{k,h}^{(1)} \cap \widehat{\mathcal{C}}_{k,h}^{(2)} \cap \bar{\mathcal{C}}_{k,h} \cap \dot{\mathcal{C}}_{k,h} \cap \check{\mathcal{C}}_{k,h} \cap \check{\mathcal{C}}_{k,h}$$

and

$$\begin{aligned} & \left| [\widehat{\mathbf{V}}_{k,h} \widehat{\mathbf{V}}_{k,h+1}] (s_h^k, a_h^k) - [\mathbb{V}_h \mathbf{V}_{h+1}^*] (s_h^k, a_h^k) \right| \leq U_{k,h} \\ & \quad [\mathbb{V}_h (\mathbf{V} - \mathbf{V}_{h+1}^*)] (s_h^k, a_h^k) \leq E_{k,h} \end{aligned}$$

where

$$\begin{aligned} U_{k,h} = \min & \left\{ \widetilde{\beta} \left\| \boldsymbol{\phi}(s_h^k, a_h^k) \right\|_{\widehat{\boldsymbol{\Lambda}}_{k,h}^{-1}} + 4H \left[\left| \langle \widehat{\boldsymbol{\mu}}_{k,h} (\widehat{\mathbf{V}}_{k,h+1} - \check{\mathbf{V}}_{k,h+1}), \boldsymbol{\phi}(s_h^k, a_h^k) \rangle \right| + \bar{\beta} \left\| \boldsymbol{\phi}(s_h^k, a_h^k) \right\|_{\widehat{\boldsymbol{\Lambda}}_{k,h}^{-1}} \right. \right. \\ & \left. \left. + \widetilde{\beta} \left\| \boldsymbol{\phi}(s_h^k, a_h^k) \right\|_{\widehat{\boldsymbol{\Lambda}}_{k,h}^{-1}} \right], 2H^2 \right\} \end{aligned} \tag{35}$$

$$E_{k,h} = \min \left\{ H \left[\langle \widehat{\boldsymbol{\mu}}_{k,h} \widehat{\mathbf{V}}_{k,h+1}, \boldsymbol{\phi}(s_h^k, a_h^k) \rangle - \langle \widehat{\boldsymbol{\mu}}_{k,h} \check{\mathbf{V}}_{k,h+1}, \boldsymbol{\phi}(s_h^k, a_h^k) \rangle + \widehat{\beta} \left\| \boldsymbol{\phi}(s_h^k, a_h^k) \right\|_{\widehat{\boldsymbol{\Lambda}}_{k,h}^{-1}} + \widetilde{\beta} \left\| \boldsymbol{\phi}(s_h^k, a_h^k) \right\|_{\widehat{\boldsymbol{\Lambda}}_{k,h}^{-1}} \right], H^2 \right\}. \tag{36}$$

Here $\bar{\beta}, \widehat{\beta}, \widetilde{\beta}, \check{\beta}$ are specified in Lemma C.8, C.9, C.10, C.11, respectively.

Proof. We first prove the following claim:

For any $\delta \in (0, 1)$, any $k \in [K]$ and fixed $h \in [H]$, with probability at least $1 - 7(H - h)\delta/H$, for any h' such that $h \leq h' \leq H$:

$$\boldsymbol{\mu}_{h'} \in \hat{\mathcal{C}}_{k,h'} \cap \hat{\mathcal{C}}_{k,h'}^{(1)} \cap \hat{\mathcal{C}}_{k,h'}^{(2)} \cap \bar{\mathcal{C}}_{k,h'} \cap \dot{\mathcal{C}}_{k,h'} \cap \tilde{\mathcal{C}}_{k,h'} \cap \check{\mathcal{C}}_{k,h'},$$

and

$$\begin{aligned} \left| [\hat{\mathbb{V}}_{k,h} \hat{V}_{k,h+1}](s_h^k, a_h^k) - [\mathbb{V}_h V_{h+1}^*](s_h^k, a_h^k) \right| &\leq U_{k,h} \\ [\mathbb{V}_h (V - V_{h+1}^*)](s_h^k, a_h^k) &\leq \cdot E_{k,h} \end{aligned}$$

hold simultaneously.

We prove this claim by induction.

- We first prove the claim for $h = H$. Since $\check{V}_{k,H+1}(\cdot) = V_{H+1}^*(\cdot) = \hat{V}_{k,H+1}(\cdot) = 0$ in Algorithm 1 for any $k \in [K]$, the conclusion holds for sure.
- Assume the claim holds for $h + 1 \leq H$. Then, for any $k \in [K]$ and fixed $h + 1 \in [H]$, with probability at least $1 - 7(H - h - 1)\delta/H$, $\boldsymbol{\mu}_{h'} \in \hat{\mathcal{C}}_{k,h'} \cap \check{\mathcal{C}}_{k,h'}$ for any $h + 1 \leq h' \leq H$, which implies $\check{\Psi}_{h+1} \cap \hat{\Psi}_{h+1}$ holds.

Combined with conclusions from Lemma C.16 and Lemma C.17, for any $k \in [K]$, the following events holds with probability at least $1 - 7\delta/H$:

$$\boldsymbol{\mu}_h \in \hat{\mathcal{C}}_{k,h}^{(1)} \cap \hat{\mathcal{C}}_{k,h}^{(2)} \cap \bar{\mathcal{C}}_{k,h} \cap \dot{\mathcal{C}}_{k,h} \cap \tilde{\mathcal{C}}_{k,h} \cap \check{\mathcal{C}}_{k,h}.$$

Moreover, we have

$$\begin{aligned} \left| [\hat{\mathbb{V}}_{k,h} \hat{V}_{k,h+1}](s_h^k, a_h^k) - [\mathbb{V}_h V_{h+1}^*](s_h^k, a_h^k) \right| &\leq U_{k,h} \\ [\mathbb{V}_h (V - V_{h+1}^*)](s_h^k, a_h^k) &\leq E_{k,h} \end{aligned}$$

under the event $\hat{\Psi}_{h+1} \cap \check{\Psi}_{h+1}$ and $\bar{\mathcal{C}}_{k,h} \cap \dot{\mathcal{C}}_{k,h} \cap \tilde{\mathcal{C}}_{k,h} \cap \check{\mathcal{C}}_{k,h}$ by Lemma C.13 and C.14.

Considering $\boldsymbol{\mu}_h \in \hat{\mathcal{C}}_{k,h}^{(1)} \cap \hat{\mathcal{C}}_{k,h}^{(2)}$, we have

$$\begin{aligned} \left\| (\boldsymbol{\mu}_h - \hat{\boldsymbol{\mu}}_{k,h}) \hat{V}_{k,h+1} \right\|_{\hat{\Lambda}_{k,h}} &\leq \left\| (\boldsymbol{\mu}_h - \hat{\boldsymbol{\mu}}_{k,h}) \mathbf{V}_{h+1}^* \right\|_{\hat{\Lambda}_{k,h}} + \left\| (\boldsymbol{\mu}_h - \hat{\boldsymbol{\mu}}_{k,h}) (\hat{V}_{k,h+1} - \mathbf{V}_{h+1}^*) \right\|_{\hat{\Lambda}_{k,h}} \\ &\leq \hat{\beta}^{(1)} + \hat{\beta}^{(2)} = \hat{\beta}, \end{aligned}$$

which implies $\boldsymbol{\mu}_h \in \hat{\mathcal{C}}_{k,h}$ since $\hat{\mathcal{C}}_{k,h} = \{\boldsymbol{\mu} : \|(\boldsymbol{\mu} - \hat{\boldsymbol{\mu}}_{k,h}) \hat{V}_{k,h+1}\|_{\hat{\Lambda}_{k,h}} \leq \hat{\beta}\}$. In other words, $\boldsymbol{\mu}_h \in \hat{\mathcal{C}}_{k,h} \cap \hat{\mathcal{C}}_{k,h}^{(1)} \cap \hat{\mathcal{C}}_{k,h}^{(2)} \cap \bar{\mathcal{C}}_{k,h} \cap \dot{\mathcal{C}}_{k,h} \cap \tilde{\mathcal{C}}_{k,h} \cap \check{\mathcal{C}}_{k,h}$.

Thus, by taking a union bound over these two events, we claim that with probability at least $1 - 7(H - h)\delta/H$, the claim holds for h .

Therefore, the claim is proved by induction and setting $h = 1$ gives the desired results in Lemma C.18. \square

D. Regret Upper Bound

In this section, we upper bound the final regret, where we show that the total regret is roughly bounded by the summation of the exploration bonus, i.e.,

$$\text{Regret}(K) \leq \sum_{k=1}^K \left[\hat{V}_{k,1}(s_1^k) - V_1^{\pi^k}(s_1^k) \right] \leq \sum_{k=1}^K \left[\hat{V}_{k,1}(s_1^k) - V_1^{\pi^k}(s_1^k) \right] \lesssim \sum_{k=1}^K \sum_{h=1}^H \hat{\beta} \|\phi(s_h^k, a_h^k)\|_{\hat{\Lambda}_{k,h}^{-1}}$$

$$\lesssim \hat{\beta} \underbrace{\sqrt{\sum_{k=1}^K \sum_{h=1}^H \hat{\sigma}_{k,h}^2}}_{\tilde{O}(\sqrt{HT})}} \underbrace{\sqrt{\sum_{k=1}^K \sum_{h=1}^H \|\hat{\sigma}_{k,h}^{-1} \phi(s_h^k, a_h^k)\|_{\hat{\Lambda}_{k,h}^{-1}}^2}}_{\tilde{O}(\sqrt{Hd})}} \leq \tilde{O}(Hd\sqrt{T}).$$

- The first and second inequalities hold by optimism and over-optimism of constructed value functions. Specifically, we first build the optimism and pessimism of constructed value functions in Lemma D.1. Then, the over-optimism is established in Lemma D.2.
- The third inequality is proved in Lemma D.5, which bounds the cumulative difference between the over-optimistic value function $\hat{V}_{k,h}$ and the value function associated with policy π^k value function $V_h^{\pi^k}$ in Lemma D.5. We further bound the cumulative difference between the over-optimistic value function $\hat{V}_{k,h}$ and the pessimistic value function $\check{V}_{k,h}$ in Lemma D.6.
- The fourth inequality holds by Cauchy-Schwarz inequality, where first summation of estimated variance is bounded in Lemma D.10 in Appendix D.3. The summation of estimated variance $\sum_{k=1}^K \sum_{h=1}^H \hat{\sigma}_{k,h}^2 = \tilde{O}(HT)$, utilizing the Law of Total Variance in (Lattimore & Hutter, 2012), detailed in Lemma D.7. The second summation can be bounded by classical Elliptical Potential Lemma, presented in Lemma F.5 in Appendix F.

Putting these building blocks together, we are finally ready to upper bound the regret in Appendix D.4 as $\tilde{O}(Hd\sqrt{T} + H^4d^4 + H^3d^6)$. Before the formal proof begins, denote the event when the conclusion of Lemma 7.5 holds as Υ . Also denote the event that the conclusion of Lemma D.3, D.4 and D.7 holds as Ξ_1, Ξ_2 and Ξ_3 , respectively. The final regret bound is a high probability bound builds under event $\Upsilon \cap \Xi_1 \cap \Xi_2 \cap \Xi_3$.

D.1. Monotonicity

In this subsection, we build the optimism of the constructed optimistic value function $\hat{V}_{k,h}$, and the pessimism of the constructed pessimistic value function $\check{V}_{k,h}$ over the optimal value function V_h^* in Lemma D.1. We also build over-optimism of $\hat{V}_{k,h}$ over $\check{V}_{k,h}$ in Lemma D.2. These are the preliminaries for our later proofs.

Lemma D.1 (Optimism and Pessimism). *In Algorithm 1, if $\check{\Psi}_{k,h} \cap \hat{\Psi}_{k,h}$ holds, then for any $k \in [K]$ and any $h \in [H]$, we have*

$$\check{V}_{k,h}(s) \stackrel{(a)}{\leq} V_h^*(s) \stackrel{(b)}{\leq} \hat{V}_{k,h}(s), \quad \forall s \in \mathcal{S}. \quad (37)$$

Proof. We prove two inequalities by induction on respective hypotheses.

(a) Pessimism: For any $k \in [K]$, the statement holds for $h = H + 1$ since $\check{V}_{k,H+1}(\cdot) = V_{H+1}^*(\cdot) = 0$.

Assume the statement holds for $h + 1$, which means $\check{V}_{k,h+1}(\cdot) \leq V_{h+1}^*(\cdot)$ under $\check{\Psi}_{k,h+1}$. Since $\check{Q}_{k,h}(\cdot, \cdot) = r_h(\cdot, \cdot) + \langle \hat{\mu}_{k,h} \check{V}_{k,h+1}, \phi(\cdot, \cdot) \rangle - \check{\beta} \|\phi(\cdot, \cdot)\|_{\hat{\Lambda}_{k,h}^{-1}}$, for $\forall (s, a) \in \mathcal{S} \times \mathcal{A}$, we have:

$$\begin{aligned} & Q_h^*(s, a) - \check{Q}_{k,h}(s, a) \\ &= r_h(s, a) + \mathbb{P}_h V_{h+1}^*(s, a) - \left[r_h(s, a) + \langle \hat{\mu}_{k,h} \check{V}_{k,h+1}, \phi(s, a) \rangle - \check{\beta} \|\phi(s, a)\|_{\hat{\Lambda}_{k,h}^{-1}} \right] \\ &= \langle \mu_h \check{V}_{k,h+1}, \phi(s, a) \rangle - \langle \hat{\mu}_{k,h} \check{V}_{k,h+1}, \phi(s, a) \rangle + \check{\beta} \|\phi(s, a)\|_{\hat{\Lambda}_{k,h}^{-1}} + \mathbb{P}_h V_{h+1}^*(s, a) - \mathbb{P}_h \check{V}_{k,h+1}(s, a) \\ &\geq - \|\hat{\mu}_{k,h} - \mu_h\|_{\hat{\Lambda}_{k,h}} \|\check{V}_{k,h+1}\|_{\hat{\Lambda}_{k,h}^{-1}} + \check{\beta} \|\phi(s, a)\|_{\hat{\Lambda}_{k,h}^{-1}} + \mathbb{P}_h V_{h+1}^*(s, a) - \mathbb{P}_h \check{V}_{k,h+1}(s, a) \\ &\geq \mathbb{P}_h V_{h+1}^*(s, a) - \mathbb{P}_h \check{V}_{k,h+1}(s, a) \\ &\geq 0, \end{aligned}$$

where the first inequality holds due to Cauchy-Schwarz inequality, the second inequality holds by the assumption that $\mu_h \in \check{\mathcal{C}}_{k,h}$ under $\check{\Psi}_{k,h}$, the third inequality holds since the induction assumption $\check{V}_{k,h+1}(\cdot) \leq V_{h+1}^*(\cdot)$ under $\check{\Psi}_{k,h+1}$

and \mathbb{P}_h is a valid distribution. Therefore, we have $\check{Q}_{k,h}(s, a) \leq Q_h^*(s, a)$, for all $(s, a) \in \mathcal{S} \times \mathcal{A}$. Since $\check{V}_{k,h}(\cdot) = \max \{ \max_{a \in \mathcal{A}} \check{Q}_{k,h}(\cdot, a), 0 \}$, for any $s \in \mathcal{S}$, we have the following two cases:

- If $\max_{a \in \mathcal{A}} \check{Q}_{k,h}(s, a) \leq 0$, we have $0 = \check{V}_{k,h}(s) \leq V_h^*(s)$.
- Otherwise, $\check{V}_{k,h}(s) = \max_{a \in \mathcal{A}} \check{Q}_{k,h}(s, a) \leq \max_{a \in \mathcal{A}} Q_h^*(s, a) = V_h^*(s)$.

Therefore, we have $\check{V}_{k,h}(\cdot) \leq V_h^*(\cdot)$ under $\check{\Psi}_{k,h}$ for any $(k, h) \in [K] \times [H]$.

(b): Optimism: For any $k \in [K]$, the statement holds for $h = H + 1$ since $\hat{V}_{k,H+1}(\cdot) = V_{H+1}^*(\cdot) = 0$.

Assume the statement holds for $h + 1$, which means $\hat{V}_{k,h+1}(\cdot) \geq V_{h+1}^*(\cdot)$ under $\hat{\Psi}_{k,h+1}$. Since $\hat{Q}_{k,h}(\cdot, \cdot) = r_h(\cdot, \cdot) + \langle \hat{\mu}_{k,h} \hat{V}_{k,h+1}, \phi(\cdot, \cdot) \rangle + \hat{\beta} \|\phi(\cdot, \cdot)\|_{\hat{\Lambda}_{k,h}^{-1}}$, for $\forall (s, a) \in \mathcal{S} \times \mathcal{A}$, we have

$$\begin{aligned}
 & \hat{Q}_{k,h}(s, a) - Q_h^*(s, a) \\
 &= r_h(s, a) + \langle \hat{\mu}_{k,h} \hat{V}_{k,h+1}, \phi(s, a) \rangle + \hat{\beta} \|\phi(s, a)\|_{\hat{\Lambda}_{k,h}^{-1}} - [r_h(s, a) + \mathbb{P}_h V_{h+1}^*(s, a)] \\
 &= \langle \hat{\mu}_{k,h} \hat{V}_{k,h+1}, \phi(s, a) \rangle - \langle \mu_h \hat{V}_{k,h+1}, \phi(s, a) \rangle + \hat{\beta} \|\phi(s, a)\|_{\hat{\Lambda}_{k,h}^{-1}} + \mathbb{P}_h \hat{V}_{k,h+1}(s, a) - \mathbb{P}_h V_{h+1}^*(s, a) \\
 &\geq - \|(\hat{\mu}_{k,h} - \mu_h) \hat{V}_{k,h+1}\|_{\hat{\Lambda}_{k,h}} \|\phi(s, a)\|_{\hat{\Lambda}_{k,h}^{-1}} + \hat{\beta} \|\phi(s, a)\|_{\hat{\Lambda}_{k,h}^{-1}} + \mathbb{P}_h \hat{V}_{k,h+1}(s, a) - \mathbb{P}_h V_{h+1}^*(s, a) \\
 &\geq \mathbb{P}_h \hat{V}_{k,h+1}(s, a) - \mathbb{P}_h V_{h+1}^*(s, a) \\
 &\geq 0,
 \end{aligned}$$

where the first inequality holds due to Cauchy-Schwarz inequality, the second inequality holds by the assumption that $\mu_h \in \hat{C}_{k,h}$ under $\hat{\Psi}_{k,h}$, the third inequality holds by the induction assumption $\hat{V}_{k,h+1}(\cdot) \geq V_{h+1}^*(\cdot)$ under $\hat{\Psi}_{k,h+1}$ and \mathbb{P}_h is a valid distribution. Therefore, we have $\hat{Q}_{k,h}(s, a) \geq Q_h^*(s, a)$, for all $(s, a) \in \mathcal{S} \times \mathcal{A}$.

Since $\hat{V}_{k,h}(\cdot) = \min \{ \max_{a \in \mathcal{A}} \hat{Q}_{k,h}(\cdot, a), H \}$, we have the following two cases:

- If $\max_{a \in \mathcal{A}} \hat{Q}_{k,h}(s, a) \leq H$, we have $H = \hat{V}_{k,h}(s) \geq V_h^*(s)$.
- Otherwise, $\hat{V}_{k,h}(s) = \max_{a \in \mathcal{A}} \hat{Q}_{k,h}(s, a) \geq \max_{a \in \mathcal{A}} \hat{Q}_{k,h}^*(s, a) = V_h^*(s)$.

Therefore, we have $\hat{V}_{k,h}(\cdot) \geq V_h^*(\cdot)$ under $\hat{\Psi}_{k,h}$ for any $(k, h) \in [K] \times [H]$. \square

Lemma D.2 (Over-optimism). *In Algorithm 1, for any $i, j \in [K]$ with $i \leq j$ and any $h \in [H]$, if $\hat{\Psi}_{i,h} \cap \hat{\Psi}_{j,h}$ holds, we have*

$$\hat{V}_{i,h}(s) \geq \hat{V}_{j,h}(s), \quad \forall s \in \mathcal{S}. \quad (38)$$

Proof. We prove the state by induction. Initially, the statement holds for $h = H + 1$ since $\hat{V}_{i,H+1}(\cdot) = \hat{V}_{j,H+1}(\cdot) = 0$.

Assume the statement holds for $h + 1$, which means $\hat{V}_{i,h+1}(\cdot) \geq \hat{V}_{j,h+1}(\cdot)$ under $\hat{\Psi}_{i,h+1} \cap \hat{\Psi}_{j,h+1}$. Since $\hat{Q}_{i,h}(\cdot, \cdot) = r_h(\cdot, \cdot) + \langle \hat{\mu}_{i,h} \hat{V}_{i,h+1}, \phi(\cdot, \cdot) \rangle + 3\hat{\beta} \|\phi(\cdot, \cdot)\|_{\hat{\Lambda}_{i,h}^{-1}}$ and $\hat{Q}_{j,h}(\cdot, \cdot) = r_h(\cdot, \cdot) + \langle \hat{\mu}_{j,h} \hat{V}_{j,h+1}, \phi(\cdot, \cdot) \rangle + \hat{\beta} \|\phi(\cdot, \cdot)\|_{\hat{\Lambda}_{j,h}^{-1}}$, for $\forall (s, a) \in \mathcal{S} \times \mathcal{A}$, we have

$$\begin{aligned}
 & \hat{Q}_{i,h}(s, a) - \hat{Q}_{j,h}(s, a) \\
 &= r_h(s, a) + \langle \hat{\mu}_{i,h} \hat{V}_{i,h+1}, \phi(s, a) \rangle + 3\hat{\beta} \|\phi(s, a)\|_{\hat{\Lambda}_{i,h}^{-1}} - [r_h(s, a) + \langle \hat{\mu}_{j,h} \hat{V}_{j,h+1}, \phi(s, a) \rangle + \hat{\beta} \|\phi(s, a)\|_{\hat{\Lambda}_{j,h}^{-1}}] \\
 &= \langle \hat{\mu}_{i,h} \hat{V}_{i,h+1}, \phi(s, a) \rangle - \langle \mu_h \hat{V}_{i,h+1}, \phi(s, a) \rangle + \langle \mu_h \hat{V}_{j,h+1}, \phi(s, a) \rangle - \langle \hat{\mu}_{j,h} \hat{V}_{j,h+1}, \phi(s, a) \rangle \\
 &\quad + 3\hat{\beta} \|\phi(s, a)\|_{\hat{\Lambda}_{i,h}^{-1}} - \hat{\beta} \|\phi(s, a)\|_{\hat{\Lambda}_{j,h}^{-1}} + \mathbb{P}_h \hat{V}_{i,h+1}(s, a) - \mathbb{P}_h \hat{V}_{j,h+1}(s, a)
 \end{aligned}$$

$$\begin{aligned}
 &\geq -\|(\hat{\boldsymbol{\mu}}_{i,h} - \boldsymbol{\mu}_h)\hat{\mathbf{V}}_{i,h+1}\|_{\hat{\boldsymbol{\Lambda}}_{i,h}}\|\boldsymbol{\phi}(s,a)\|_{\hat{\boldsymbol{\Lambda}}_{i,h}^{-1}} - \|(\hat{\boldsymbol{\mu}}_{j,h} - \boldsymbol{\mu}_h)\hat{\mathbf{V}}_{j,h+1}\|_{\hat{\boldsymbol{\Lambda}}_{j,h}}\|\boldsymbol{\phi}(s,a)\|_{\hat{\boldsymbol{\Lambda}}_{j,h}^{-1}} \\
 &\quad + 3\hat{\beta}\|\boldsymbol{\phi}(s,a)\|_{\hat{\boldsymbol{\Lambda}}_{i,h}^{-1}} - \hat{\beta}\|\boldsymbol{\phi}(s,a)\|_{\hat{\boldsymbol{\Lambda}}_{j,h}^{-1}} + \mathbb{P}_h\hat{\mathbf{V}}_{i,h+1}(s,a) - \mathbb{P}_h\hat{\mathbf{V}}_{j,h+1}(s,a) \\
 &\geq 2\hat{\beta}\|\boldsymbol{\phi}(s,a)\|_{\hat{\boldsymbol{\Lambda}}_{i,h}^{-1}} - 2\hat{\beta}\|\boldsymbol{\phi}(s,a)\|_{\hat{\boldsymbol{\Lambda}}_{j,h}^{-1}} + \mathbb{P}_h\hat{\mathbf{V}}_{i,h+1}(s,a) - \mathbb{P}_h\hat{\mathbf{V}}_{j,h+1}(s,a) \\
 &\geq \mathbb{P}_h\hat{\mathbf{V}}_{i,h+1}(s,a) - \mathbb{P}_h\hat{\mathbf{V}}_{j,h+1}(s,a) \\
 &\geq 0,
 \end{aligned}$$

where the first inequality holds due to Cauchy-Schwarz inequality, the second inequality holds by the assumption that $\boldsymbol{\mu}_h \in \hat{\mathcal{C}}_{i,h} \cap \hat{\mathcal{C}}_{j,h}$ under $\hat{\Psi}_{i,h} \cap \hat{\Psi}_{j,h}$, the third inequality holds since $\hat{\boldsymbol{\Lambda}}_{i,h} \leq \hat{\boldsymbol{\Lambda}}_{j,h}$, and the last inequality by the induction assumption $\hat{\mathbf{V}}_{i,h+1}(\cdot) \geq \hat{\mathbf{V}}_{j,h+1}(\cdot)$ under $\hat{\Psi}_{i,h+1} \cap \hat{\Psi}_{j,h+1}$ and \mathbb{P}_h is a valid distribution. Therefore, we have $\hat{Q}_{i,h}(s,a) \geq \hat{Q}_{j,h}(s,a)$, for all $(s,a) \in \mathcal{S} \times \mathcal{A}$.

Since $\hat{V}_{i,h}(\cdot) = \min\{\max_{a \in \mathcal{A}} \hat{Q}_{i,h}(\cdot, a), H\}$ and $\hat{V}_{j,h}(\cdot) = \min\{\max_{a \in \mathcal{A}} \hat{Q}_{j,h}(\cdot, a), H\}$, we have $\hat{V}_{i,h}(\cdot) \geq \hat{V}_{j,h}(\cdot)$ under $\hat{\Psi}_{i,h} \cap \hat{\Psi}_{j,h}$ for any $(k,h) \in [K] \times [H]$. \square

D.2. Suboptimality Gap

In this subsection, we establish Lemma D.5 and Lemma D.6 that bound the distance of the over-optimistic value function $\hat{V}_{k,h}$ to the value function $V_{h+1}^{\pi_k}$ associated with policy π_k and the pessimistic value function $\check{V}_{k,h}$, respectively. Before that, we present two high probability events Ξ_1 and Ξ_2 in Lemma D.3 and D.4, respectively.

Lemma D.3. *In Algorithm 1, for any $\delta \in (0, 1)$, with probability at least $1 - \delta$, simultaneously for all $h' \in [H]$, we have*

$$\sum_{k=1}^K \sum_{h=h'}^H \left[\mathbb{P}_h(\hat{V}_{k,h+1} - V_{h+1}^{\pi_k})(s_h^k, a_h^k) - [\hat{V}_{k,h+1} - V_{h+1}^{\pi_k}](s_{h+1}^k) \right] \leq 2H\sqrt{2T \log\left(\frac{H}{\delta}\right)}$$

Proof. Denote $\Delta_{k,h} = [\mathbb{P}_h(\hat{V}_{k,h+1} - V_{h+1}^{\pi_k})(s_h^k, a_h^k) - [\hat{V}_{k,h+1} - V_{h+1}^{\pi_k}](s_{h+1}^k)]$. Since s_{h+1}^k is $\mathcal{F}_{k,h+1}$ -measurable, $\Delta_{k,h}$ is $\mathcal{F}_{k,h+1}$ -measurable and $\mathbb{E}[\Delta_{k,h} | \mathcal{F}_{k,h}] = 0$. Thus, for some $h' \in [H]$, $\{\Delta_{k,h'}, \Delta_{k,h'+1}, \dots, \Delta_{k,H}\}_{k \in [K]}$ is a martingale difference sequence. Since $|\Delta_{k,h}| \leq 2H$ by $-H \leq \hat{V}_{k,h+1}(\cdot) - V_{h+1}^{\pi_k}(\cdot) \leq H$, we can apply Azuma-Hoeffding inequality (Lemma F.1) to this martingale difference sequence and obtain

$$\sum_{k=1}^K \sum_{h=h'}^H \Delta_{k,h} \leq 2H\sqrt{2K(H-h'+1)\log(1/\delta)} \leq 2H\sqrt{2T\log(1/\delta)}, \quad (39)$$

for some $h' \in [H]$ with probability at least $1 - \delta$. Taking a union bound over all $h' \in [H]$ gives the final conclusion. \square

Lemma D.4. *In Algorithm 1, for any $\delta \in (0, 1)$, with probability at least $1 - \delta$, simultaneously for all $h' \in [H]$, we have*

$$\sum_{k=1}^K \sum_{h=h'}^H \left[\mathbb{P}_h(\hat{V}_{k,h+1} - \check{V}_{k,h+1})(s_h^k, a_h^k) - [\hat{V}_{k,h+1} - \check{V}_{k,h+1}](s_{h+1}^k) \right] \leq 2H\sqrt{2T \log\left(\frac{H}{\delta}\right)}$$

Proof. The proof is almost the same as that of Lemma D.3, except for replacing $V_{k,h+1}^{\pi_k}(\cdot)$ by $\check{V}_{k,h+1}(\cdot)$. \square

Lemma D.5. *In Algorithm 1, under $\hat{\Psi}_1 \cap \Xi_1$, we have*

$$\begin{aligned}
 \sum_{k=1}^K \hat{V}_{k,1}(s_1^k) - V_1^{\pi_k}(s_1^k) &\leq 4\hat{\beta}\sqrt{\sum_{k=1}^K \sum_{h=1}^H \hat{\sigma}_{k,h}^2} \sqrt{H \cdot 2d \log\left(1 + \frac{K}{Hd\lambda}\right)} \\
 &\quad + \frac{3H^2d}{\log(2)} \log\left(1 + \frac{1}{\lambda H \log(2)}\right) + 2H\sqrt{2T \log\left(\frac{H}{\delta}\right)} \quad (40)
 \end{aligned}$$

$$\begin{aligned} \sum_{k=1}^K \sum_{h=1}^H [\mathbb{P}_h(\dot{\hat{V}}_{k,h+1} - V_{h+1}^{\pi^k})](s_h^k, a_h^k) &\leq 4\hat{\beta}H \sqrt{\sum_{k=1}^K \sum_{h=1}^H \hat{\sigma}_{k,h}^2} \sqrt{H \cdot 2d \log\left(1 + \frac{K}{Hd\lambda}\right)} \\ &\quad + \frac{3H^3d}{\log(2)} \log\left(1 + \frac{1}{\lambda \log(2)}\right) + 4H^2 \sqrt{2T \log\left(\frac{H}{\delta}\right)} \end{aligned} \quad (41)$$

Proof. By Algorithm 1, for any $k \in [K], h \in [H]$, we have

$$\begin{aligned} \dot{Q}_{k,1}(s, a) &\leq r_1(s, a) + \langle \hat{\boldsymbol{\mu}}_{k,1} \hat{\mathbf{V}}_{k,2}, \boldsymbol{\phi}(s, a) \rangle + 3\hat{\beta} \|\boldsymbol{\phi}(s, a)\|_{\hat{\boldsymbol{\Lambda}}_{k,1}^{-1}} \\ &= r_1(s, a) + \hat{\mathbb{P}}_{k,1} \hat{\mathbf{V}}_{k,2}(s, a) + 3\hat{\beta} \|\boldsymbol{\phi}(s, a)\|_{\hat{\boldsymbol{\Lambda}}_{k,1}^{-1}}, \\ Q_1^{\pi^k}(s, a) &= r_1(s, a) + \mathbb{P}_1 V_2^{\pi^k}(s, a). \end{aligned} \quad (42)$$

Then,

$$\begin{aligned} \dot{\hat{V}}_{k,1}(s_1^k) - V_h^{\pi^k}(s_1^k) &= \dot{Q}_{k,1}(s_1^k, a_1^k) - Q_1^{\pi^k}(s_1^k, a_1^k) \\ &\leq r_1(s_1^k, a_1^k) + \hat{\mathbb{P}}_{k,1} \hat{\mathbf{V}}_{k,2}(s_1^k, a_1^k) + 3\hat{\beta} \|\boldsymbol{\phi}(s_1^k, a_1^k)\|_{\hat{\boldsymbol{\Lambda}}_{k,1}^{-1}} - r_1(s_1^k, a_1^k) - \mathbb{P}_1 V_2^{\pi^k}(s_1^k, a_1^k) \\ &= 3\hat{\beta} \|\boldsymbol{\phi}(s_1^k, a_1^k)\|_{\hat{\boldsymbol{\Lambda}}_{k,1}^{-1}} + [\hat{\mathbb{P}}_{k,1} \hat{\mathbf{V}}_{k,2}(s_1^k, a_1^k) - \mathbb{P}_1 \hat{\mathbf{V}}_{k,2}(s_1^k, a_1^k)] + [\mathbb{P}_1(\hat{\mathbf{V}}_{k,2} - V_2^{\pi^k})](s_1^k, a_1^k) \\ &= 3\hat{\beta} \|\boldsymbol{\phi}(s_1^k, a_1^k)\|_{\hat{\boldsymbol{\Lambda}}_{k,1}^{-1}} + \langle (\hat{\boldsymbol{\mu}}_{k,1} - \boldsymbol{\mu}_1) \hat{\mathbf{V}}_{k,2}, \boldsymbol{\phi}(s_1^k, a_1^k) \rangle + [\mathbb{P}_1(\hat{\mathbf{V}}_{k,2} - V_2^{\pi^k})](s_1^k, a_1^k) \\ &\leq 3\hat{\beta} \|\boldsymbol{\phi}(s_1^k, a_1^k)\|_{\hat{\boldsymbol{\Lambda}}_{k,1}^{-1}} + \|(\hat{\boldsymbol{\mu}}_{k,1} - \boldsymbol{\mu}_1) \hat{\mathbf{V}}_{k,2}\|_{\hat{\boldsymbol{\Lambda}}_{k,1}} \|\boldsymbol{\phi}(s_1^k, a_1^k)\|_{\hat{\boldsymbol{\Lambda}}_{k,1}^{-1}} + [\mathbb{P}_1(\hat{\mathbf{V}}_{k,2} - V_2^{\pi^k})](s_1^k, a_1^k) \\ &\leq 4\hat{\beta} \|\boldsymbol{\phi}(s_1^k, a_1^k)\|_{\hat{\boldsymbol{\Lambda}}_{k,1}^{-1}} + [\mathbb{P}_1(\hat{\mathbf{V}}_{k,2} - V_2^{\pi^k})](s_1^k, a_1^k) \\ &\leq 4\hat{\beta} \|\boldsymbol{\phi}(s_1^k, a_1^k)\|_{\hat{\boldsymbol{\Lambda}}_{k,1}^{-1}} + [\mathbb{P}_1(\dot{\hat{V}}_{k,2} - V_2^{\pi^k})](s_1^k, a_1^k) \\ &= 4\hat{\beta} \|\boldsymbol{\phi}(s_1^k, a_1^k)\|_{\hat{\boldsymbol{\Lambda}}_{k,1}^{-1}} + [\dot{\hat{V}}_{k,2} - V_2^{\pi^k}](s_2^k) + [\mathbb{P}_1(\dot{\hat{V}}_{k,2} - V_2^{\pi^k})](s_1^k, a_1^k) - [\dot{\hat{V}}_{k,2} - V_2^{\pi^k}](s_2^k) \\ &\quad \dots \\ &\leq \sum_{h=1}^H 4\hat{\beta} \|\boldsymbol{\phi}(s_h^k, a_h^k)\|_{\hat{\boldsymbol{\Lambda}}_{k,h}^{-1}} + \sum_{h=1}^H \left[[\mathbb{P}_h(\dot{\hat{V}}_{k,h+1} - V_{h+1}^{\pi^k})](s_h^k, a_h^k) - [\dot{\hat{V}}_{k,h+1} - V_{h+1}^{\pi^k}](s_{h+1}^k) \right] \end{aligned} \quad (43)$$

where the first inequality is due to Eq. (42), the second inequality holds by the Cauchy-Schwarz inequality, third inequality holds since $\boldsymbol{\mu}_1 \in \hat{\mathcal{C}}_{k,1}$ under $\hat{\Psi}_1$, the fourth inequality holds since $\hat{\mathbf{V}}_{k,2}(\cdot) \leq \dot{\hat{V}}_{k,2}(\cdot)$ by definitions, and the last inequality holds since we can expand $\dot{\hat{V}}_{k,h+1}(s_h^k) - \dot{V}_{k,h+1}(s_h^k)$ in a recursive way until stage H .

Summing up Eq.(43) for $k \in [K]$ gives

$$\begin{aligned} &\sum_{k=1}^K \dot{\hat{V}}_{k,1}(s_1^k) - V_1^{\pi^k}(s_1^k) \\ &\leq \sum_{k=1}^K \sum_{h=1}^H 4\hat{\beta} \|\boldsymbol{\phi}(s_h^k, a_h^k)\|_{\hat{\boldsymbol{\Lambda}}_{k,1}^{-1}} + \sum_{k=1}^K \sum_{h=1}^H \left[[\mathbb{P}_h(\dot{\hat{V}}_{k,h+1} - V_{h+1}^{\pi^k})](s_h^k, a_h^k) - [\dot{\hat{V}}_{k,h+1} - V_{h+1}^{\pi^k}](s_{h+1}^k) \right] \\ &\leq 4\hat{\beta} \underbrace{\sum_{k=1}^K \sum_{h=1}^H \hat{\sigma}_{k,h} \|\hat{\sigma}_{k,h}^{-1} \boldsymbol{\phi}(s_h^k, a_h^k)\|_{\hat{\boldsymbol{\Lambda}}_{k,h}^{-1}}}_I + 2H \sqrt{2T \log\left(\frac{H}{\delta}\right)}, \end{aligned}$$

where the second inequality holds under Ξ_1 by Lemma D.3. Subsequently, we try to bound I . For fixed $h \in [H]$, denote $\mathbf{x}_k = \hat{\sigma}_{k,h}^{-1} \boldsymbol{\phi}(s_h^k, a_h^k)$ in Lemma D.8. Then, there are at most $3d \log[1 + d/(\lambda H \log(2))]/\log(2)$ episodes that $\|\hat{\sigma}_{k,h}^{-1} \boldsymbol{\phi}(s_h^k, a_h^k)\|_{\hat{\boldsymbol{\Lambda}}_{k,h}^{-1}} \geq 1$, which further implies that there are at most $H \cdot 3d \log[1 + d/(\lambda H \log(2))]/\log(2)$ episodes that

there exists $h' \in [H]$ such that $\|\hat{\sigma}_{k,h'}^{-1}\phi(s_{h'}, a_{h'})\|_{\hat{\Lambda}_{k,h'}} \geq 1$. Moreover, we can bound $\hat{V}_{k,1}(s_1^k) - V_1^{\pi^k}(s_1^k)$ in these episodes by H since $\hat{V}_{k,1}(\cdot) - V_1^{\pi^k}(\cdot) \leq H$ for any $k \in [K]$. Thus, we have

$$\begin{aligned}
 & \sum_{k=1}^K \hat{V}_{k,1}(s_1^k) - V_1^{\pi^k}(s_1^k) \\
 & \leq 4\hat{\beta} \sum_{k=1}^K \sum_{h=1}^H \hat{\sigma}_{k,h} \min \left\{ \|\hat{\sigma}_{k,h}^{-1}\phi(s_h^k, a_h^k)\|_{\hat{\Lambda}_{k,h}}, 1 \right\} + H \cdot H \cdot \frac{3d}{\log(2)} \log \left(1 + \frac{1}{\lambda H \log(2)} \right) + 2H \sqrt{2T \log \left(\frac{H}{\delta} \right)} \\
 & \leq 4\hat{\beta} \sqrt{\sum_{k=1}^K \sum_{h=1}^H \hat{\sigma}_{k,h}^2} \sqrt{\sum_{k=1}^K \sum_{h=1}^H \min \left\{ \|\hat{\sigma}_{k,h}^{-1}\phi(s_h^k, a_h^k)\|_{\hat{\Lambda}_{k,h}}^2, 1 \right\}} + \frac{3H^2d}{\log(2)} \log \left(1 + \frac{1}{\lambda H \log(2)} \right) + 2H \sqrt{2T \log \left(\frac{H}{\delta} \right)} \\
 & \leq 4\hat{\beta} \sqrt{\sum_{k=1}^K \sum_{h=1}^H \hat{\sigma}_{k,h}^2} \sqrt{H \cdot 2d \log \left(1 + \frac{K}{Hd\lambda} \right)} + \frac{3H^2d}{\log(2)} \log \left(1 + \frac{L^2}{\lambda H \log(2)} \right) + 2H \sqrt{2T \log \left(\frac{H}{\delta} \right)},
 \end{aligned} \tag{44}$$

where the second inequality holds by Cauchy-Schwarz inequality, and the last inequality holds by Lemma F.5 with the fact that $\|\hat{\sigma}_{k,h}^{-1}\phi(s_h^k, a_h^k)\|_2 \leq 1/\sqrt{H}$. Thus, Eq. (40) is obtained. Besides, by similar argument in Eq. (43), we obtain

$$\hat{V}_{k,h'}(s_{h'}^k) - V_{h'}^{\pi^k}(s_{h'}^k) \leq \sum_{h=h'}^H 4\hat{\beta} \|\phi(s_h^k, a_h^k)\|_{\hat{\Lambda}_{k,h}^{-1}} + \sum_{h=h'}^H \left[[\mathbb{P}_h(\hat{V}_{k,h+1} - V_{h+1}^{\pi^k})](s_h^k, a_h^k) - [\hat{V}_{k,h+1} - V_{h+1}^{\pi^k}](s_{h+1}^k) \right]$$

which further gives

$$\begin{aligned}
 & \sum_{k=1}^K \hat{V}_{k,h'}(s_{h'}^k) - V_{h'}^{\pi^k}(s_{h'}^k) \\
 & \leq \sum_{k=1}^K \sum_{h=h'}^H 4\hat{\beta} \|\phi(s_h^k, a_h^k)\|_{\hat{\Lambda}_{k,h}^{-1}} + \sum_{k=1}^K \sum_{h=h'}^H \left[[\mathbb{P}_h(\hat{V}_{k,h+1} - V_{h+1}^{\pi^k})](s_h^k, a_h^k) - [\hat{V}_{k,h+1} - V_{h+1}^{\pi^k}](s_{h+1}^k) \right] \\
 & \leq 4\hat{\beta} \sum_{k=1}^K \sum_{h=1}^H \hat{\sigma}_{k,h} \|\hat{\sigma}_{k,h}^{-1}\phi(s_h^k, a_h^k)\|_{\hat{\Lambda}_{k,h}^{-1}} + \sum_{k=1}^K \sum_{h=h'}^H \left[[\mathbb{P}_h(\hat{V}_{k,h+1} - V_{h+1}^{\pi^k})](s_h^k, a_h^k) - [\hat{V}_{k,h+1} - V_{h+1}^{\pi^k}](s_{h+1}^k) \right] \\
 & \leq 4\hat{\beta} \sqrt{\sum_{k=1}^K \sum_{h=1}^H \hat{\sigma}_{k,h}^2} \sqrt{H \cdot 2d \log \left(1 + \frac{K}{Hd\lambda} \right)} + \frac{3H^2d}{\log(2)} \log \left(1 + \frac{1}{\lambda \log(2)} \right) + 2H \sqrt{2T \log \left(\frac{H}{\delta} \right)},
 \end{aligned} \tag{45}$$

where the last inequality holds by similar argument of bounding I in Eq. (44) and Lemma D.3 under Ξ_1 .

Note that

$$\begin{aligned}
 & \sum_{k=1}^K \sum_{h=1}^H [\mathbb{P}_h(\hat{V}_{k,h+1} - V_{h+1}^{\pi^k})](s_h^k, a_h^k) \\
 & = \sum_{k=1}^K \sum_{h=2}^H [\hat{V}_{k,h+1} - V_{h+1}^{\pi^k}](s_{h+1}^k) + \sum_{k=1}^K \sum_{h=1}^H \left[[\mathbb{P}_h(\hat{V}_{k,h+1} - V_{h+1}^{\pi^k})](s_h^k, a_h^k) - [\hat{V}_{k,h+1} - V_{h+1}^{\pi^k}](s_{h+1}^k) \right] \\
 & \leq 4\hat{\beta}H \sqrt{\sum_{k=1}^K \sum_{h=1}^H \hat{\sigma}_{k,h}^2} \sqrt{H \cdot 2d \log \left(1 + \frac{K}{Hd\lambda} \right)} + \frac{3H^3d}{\log(2)} \log \left(1 + \frac{1}{\lambda \log(2)} \right) \\
 & \quad + 2H^2 \sqrt{2T \log \left(\frac{H}{\delta} \right)} + 2H \sqrt{2T \log \left(\frac{H}{\delta} \right)}
 \end{aligned}$$

$$\leq 4\hat{\beta}H\sqrt{\sum_{k=1}^K\sum_{h=1}^H\hat{\sigma}_{k,h}^2}\sqrt{H\cdot 2d\log\left(1+\frac{K}{Hd\lambda}\right)+\frac{3H^3d}{\log(2)}\log\left(1+\frac{1}{\lambda\log(2)}\right)+4H^2\sqrt{2T\log\left(\frac{H}{\delta}\right)},$$

where the first inequality holds by sum up Eq. (45) for $h' = 2, \dots, H$ and Lemma D.3 under Ξ_1 . Thus, Eq. (41) is also obtained. \square

Lemma D.6 (Gap between Optimism and Pessimism). *In Algorithm 1, under $\hat{\Psi}_1 \cap \check{\Psi}_1 \cap \Xi_2$, we have*

$$\sum_{k=1}^K\sum_{h=1}^H[\mathbb{P}_h(\hat{V}_{k,h+1}-\check{V}_{k,h+1})](s_h^k, a_h^k) \leq 2(2\hat{\beta}+\check{\beta})H\sqrt{\sum_{k=1}^K\sum_{h=1}^H\hat{\sigma}_{k,h}^2}\sqrt{H\cdot 2d\log\left(1+\frac{K}{Hd\lambda}\right)+4H^2\sqrt{2T\log\left(\frac{H}{\delta}\right)}}$$

Proof. By Algorithm 1, for any $k \in [K], h \in [H]$, we have

$$\begin{aligned}\hat{Q}_{k,h}(s, a) &\leq r_h(s, a) + \langle \hat{\mu}_{k,h} \hat{V}_{k,h+1}, \phi(s, a) \rangle + 3\hat{\beta} \|\phi(s, a)\|_{\hat{\Lambda}_{k,h}^{-1}} \\ &= r_h(s, a) + \hat{\mathbb{P}}_{k,h} \hat{V}_{k,h+1}(s, a) + 3\hat{\beta} \|\phi(s, a)\|_{\hat{\Lambda}_{k,h}^{-1}}, \\ \check{Q}_{k,h}(s, a) &\geq r_h(s, a) + \langle \hat{\mu}_{k,h} \check{V}_{k,h+1}, \phi(s, a) \rangle - \check{\beta} \|\phi(s, a)\|_{\hat{\Lambda}_{k,h}^{-1}} \\ &= r_h(s, a) + \hat{\mathbb{P}}_{k,h} \check{V}_{k,h+1}(s, a) - \check{\beta} \|\phi(s, a)\|_{\hat{\Lambda}_{k,h}^{-1}},\end{aligned}\tag{46}$$

Then,

$$\begin{aligned}\hat{V}_{k,h}(s_h^k) - \check{V}_{k,h}(s_h^k) &\leq \hat{Q}_{k,h}(s_h^k, a_h^k) - \check{Q}_{k,h}(s_h^k, a_h^k) \\ &\leq r_h(s_h^k, a_h^k) + \hat{\mathbb{P}}_{k,h} \hat{V}_{k,h+1}(s_h^k, a_h^k) + 3\hat{\beta} \|\phi(s_h^k, a_h^k)\|_{\hat{\Lambda}_{k,h}^{-1}} - [r_h(s_h^k, a_h^k) + \hat{\mathbb{P}}_{k,h} \check{V}_{k,h+1}(s_h^k, a_h^k) - \check{\beta} \|\phi(s_h^k, a_h^k)\|_{\hat{\Lambda}_{k,h}^{-1}}] \\ &= 3\hat{\beta} \|\phi(s_h^k, a_h^k)\|_{\hat{\Lambda}_{k,h}^{-1}} + \check{\beta} \|\phi(s_h^k, a_h^k)\|_{\hat{\Lambda}_{k,h}^{-1}} + [\hat{\mathbb{P}}_{k,h} \hat{V}_{k,h+1}(s_h^k, a_h^k) - \mathbb{P}_h \hat{V}_{k,h+1}(s_h^k, a_h^k)] \\ &\quad + [\mathbb{P}_h \check{V}_{k,h+1}(s_h^k, a_h^k) - \hat{\mathbb{P}}_{k,h} \check{V}_{k,h+1}(s_h^k, a_h^k)] + [\mathbb{P}_h(\hat{V}_{k,h+1} - \check{V}_{k,h+1})](s_h^k, a_h^k) \\ &= 3\hat{\beta} \|\phi(s_h^k, a_h^k)\|_{\hat{\Lambda}_{k,h}^{-1}} + \check{\beta} \|\phi(s_h^k, a_h^k)\|_{\hat{\Lambda}_{k,h}^{-1}} + \langle (\hat{\mu}_{k,h+1} - \mu_h) \hat{V}_{k,h+1}, \phi(s_h^k, a_h^k) \rangle \\ &\quad + \langle (\hat{\mu}_{k,h+1} - \mu_h) \check{V}_{k,h+1}, \phi(s_h^k, a_h^k) \rangle + [\mathbb{P}_h(\hat{V}_{k,h+1} - \check{V}_{k,h+1})](s_h^k, a_h^k) \\ &\leq 3\hat{\beta} \|\phi(s_h^k, a_h^k)\|_{\hat{\Lambda}_{k,h}^{-1}} + \check{\beta} \|\phi(s_h^k, a_h^k)\|_{\hat{\Lambda}_{k,h}^{-1}} + \|(\hat{\mu}_{k,h} - \mu_h) \hat{V}_{k,h+1}\|_{\hat{\Lambda}_{k,h}} \|\phi(s_h^k, a_h^k)\|_{\hat{\Lambda}_{k,1}^{-1}} \\ &\quad + \|(\hat{\mu}_{k,h} - \mu_h) \check{V}_{k,h+1}\|_{\hat{\Lambda}_{k,h}} \|\phi(s_h^k, a_h^k)\|_{\hat{\Lambda}_{k,1}^{-1}} + [\mathbb{P}_h(\hat{V}_{k,h+1} - \check{V}_{k,h+1})](s_h^k, a_h^k) \\ &\leq 4\hat{\beta} \|\phi(s_h^k, a_h^k)\|_{\hat{\Lambda}_{k,h}^{-1}} + 2\check{\beta} \|\phi(s_h^k, a_h^k)\|_{\hat{\Lambda}_{k,h}^{-1}} + [\mathbb{P}_h(\hat{V}_{k,h+1} - \check{V}_{k,h+1})](s_h^k, a_h^k) \\ &\leq 4\hat{\beta} \|\phi(s_h^k, a_h^k)\|_{\hat{\Lambda}_{k,h}^{-1}} + 2\check{\beta} \|\phi(s_h^k, a_h^k)\|_{\hat{\Lambda}_{k,h}^{-1}} + [\mathbb{P}_h(\hat{V}_{k,h+1} - \check{V}_{k,h+1})](s_h^k, a_h^k) \\ &= (4\hat{\beta} + 2\check{\beta}) \|\phi(s_h^k, a_h^k)\|_{\hat{\Lambda}_{k,h}^{-1}} + [\hat{V}_{k,h+1} - \check{V}_{k,h+1}](s_{h+1}^k) + [\mathbb{P}_h(\hat{V}_{k,h+1} - \check{V}_{k,h+1})](s_h^k, a_h^k) - [\hat{V}_{k,h+1} - \check{V}_{k,h+1}](s_{h+1}^k)\end{aligned}\tag{47}$$

where the first inequality holds since $\check{V}_{k,h+1}(s_h^k) \geq \check{Q}_{k,h}(s_h^k, a_h^k)$, the second inequality holds due to Eq. (46), the third inequality holds by the Cauchy-Schwarz inequality, the fourth inequality holds since $\mu_h \in \hat{\mathcal{C}}_{k,h} \cap \check{\mathcal{C}}_{k,h}$ under $\hat{\Psi}_1 \cap \check{\Psi}_1$, and the last inequality holds since $\hat{V}_{k,h+1}(\cdot) \leq \hat{V}_{k,h+1}(\cdot)$ by definitions.

Since $\hat{V}_{k,h'}(s_{h'}^k) - \check{V}_{k,h'}(s_{h'}^k) \leq H$, we further obtains

$$\begin{aligned}
 & \hat{V}_{k,h}(s_h^k) - \check{V}_{k,h}(s_h^k) \leq \min \left\{ \hat{Q}_{k,h}(s_h^k, a_h^k) - \check{Q}_{k,h}(s_h^k, a_h^k), H \right\} \\
 & \leq \min \left\{ (4\hat{\beta} + 2\check{\beta}) \|\phi(s_h^k, a_h^k)\|_{\hat{\Lambda}_{k,h}^{-1}}, H \right\} + [\mathbb{P}_h(\hat{V}_{k,h+1} - \check{V}_{k,h+1})](s_h^k, a_h^k) \\
 & \leq (4\hat{\beta} + 2\check{\beta}) \hat{\sigma}_{k,h} \min \left\{ \|\hat{\sigma}_{k,h}^{-1} \phi(s_h^k, a_h^k)\|_{\hat{\Lambda}_{k,h}^{-1}}, 1 \right\} + [\mathbb{P}_h(\hat{V}_{k,h+1} - \check{V}_{k,h+1})](s_h^k, a_h^k) \\
 & = (4\hat{\beta} + 2\check{\beta}) \hat{\sigma}_{k,h} \min \left\{ \|\hat{\sigma}_{k,h}^{-1} \phi(s_h^k, a_h^k)\|_{\hat{\Lambda}_{k,h}^{-1}}, 1 \right\} + [\hat{V}_{k,h+1} - \check{V}_{k,h+1}](s_{h+1}^k) \\
 & \quad + [\mathbb{P}_h(\hat{V}_{k,h+1} - \check{V}_{k,h+1})](s_h^k, a_h^k) - [\hat{V}_{k,h+1} - \check{V}_{k,h+1}](s_{h+1}^k),
 \end{aligned} \tag{48}$$

where the second inequality holds since $\check{\beta} \hat{\sigma}_{k,h} \geq \sqrt{Hd^2} \sqrt{H} \geq H$. Summing up Eq. (48) for $k \in [K]$ and $h = h', \dots, H$ gives

$$\begin{aligned}
 & \sum_{k=1}^K \hat{V}_{k,h'}(s_{h'}^k) - \check{V}_{k,h'}(s_{h'}^k) \\
 & \leq \sum_{k=1}^K \sum_{h=h'}^H (4\hat{\beta} + 2\check{\beta}) \hat{\sigma}_{k,h} \min \left\{ \|\hat{\sigma}_{k,h}^{-1} \phi(s_h^k, a_h^k)\|_{\hat{\Lambda}_{k,h}^{-1}}, 1 \right\} \\
 & \quad + \sum_{k=1}^K \sum_{h=h'}^H [\mathbb{P}_h(\hat{V}_{k,h+1} - \check{V}_{k,h+1})](s_h^k, a_h^k) - [\hat{V}_{k,h+1} - \check{V}_{k,h+1}](s_{h+1}^k) \\
 & \leq \sum_{k=1}^K \sum_{h=1}^H (4\hat{\beta} + 2\check{\beta}) \hat{\sigma}_{k,h} \min \left\{ \|\hat{\sigma}_{k,h}^{-1} \phi(s_h^k, a_h^k)\|_{\hat{\Lambda}_{k,h}^{-1}}, 1 \right\} + 2H \sqrt{2T \log \left(\frac{H}{\delta} \right)} \\
 & \leq (4\hat{\beta} + 2\check{\beta}) \sqrt{\sum_{k=1}^K \sum_{h=1}^H \hat{\sigma}_{k,h}^2} \sqrt{\sum_{k=1}^K \sum_{h=1}^H \min \left\{ \|\hat{\sigma}_{k,h}^{-1} \phi(s_h^k, a_h^k)\|_{\hat{\Lambda}_{k,h}^{-1}}^2, 1 \right\}} + 2H \sqrt{2T \log \left(\frac{H}{\delta} \right)} \\
 & \leq (4\hat{\beta} + 2\check{\beta}) \sqrt{\sum_{k=1}^K \sum_{h=1}^H \hat{\sigma}_{k,h}^2} \sqrt{H \cdot 2d \log \left(1 + \frac{K}{Hd\lambda} \right)} + 2H \sqrt{2T \log \left(\frac{H}{\delta} \right)}
 \end{aligned} \tag{49}$$

where the second inequality holds by Lemma D.4 under Ξ_2 , the third inequality holds due to Cauchy-Schwarz inequality, and the last inequality holds due to Lemma F.5 with the fact that $\|\hat{\sigma}_{k,h}^{-1} \phi(s_h^k, a_h^k)\|_2 \leq 1/\sqrt{H}$.

Note that

$$\begin{aligned}
 & \sum_{k=1}^K \sum_{h=1}^H [\mathbb{P}_h(\hat{V}_{k,h+1} - \check{V}_{k,h+1})](s_h^k, a_h^k) \\
 & = \sum_{k=1}^K \sum_{h=2}^H [\hat{V}_{k,h+1} - \check{V}_{k,h+1}](s_{h+1}^k) + \sum_{k=1}^K \sum_{h=1}^H \left[[\mathbb{P}_h(\hat{V}_{k,h+1} - \check{V}_{k,h+1})](s_h^k, a_h^k) - [\hat{V}_{k,h+1} - \check{V}_{k,h+1}](s_{h+1}^k) \right] \\
 & \leq (4\hat{\beta} + 2\check{\beta}) H \sqrt{\sum_{k=1}^K \sum_{h=1}^H \hat{\sigma}_{k,h}^2} \sqrt{H \cdot 2d \log \left(1 + \frac{K}{Hd\lambda} \right)} + 2H^2 \sqrt{2T \log \left(\frac{H}{\delta} \right)} + 2H \sqrt{2T \log \left(\frac{H}{\delta} \right)} \\
 & \leq 2(2\hat{\beta} + \check{\beta}) H \sqrt{\sum_{k=1}^K \sum_{h=1}^H \hat{\sigma}_{k,h}^2} \sqrt{H \cdot 2d \log \left(1 + \frac{K}{Hd\lambda} \right)} + 4H^2 \sqrt{2T \log \left(\frac{H}{\delta} \right)},
 \end{aligned}$$

where the first inequality holds by sum up Eq. (49) for $h' = 2, \dots, H$ and Lemma D.4 under Ξ_2 . \square

D.3. Summation of Estimated Variances

In this subsection, we try to bound the summation of estimated variance in Lemma D.10. As shown in Lemma D.10, the summation of estimated variance $\sum_{k=1}^K \sum_{h=1}^H \hat{\sigma}_{k,h}^2 = \tilde{O}(HT)$, which utilizes the Law of Total Variance in (Lattimore & Hutter, 2012), detailed in Lemma D.7.

Lemma D.7 (Total variance lemma, Lemma C.5 in (Jin et al., 2018)). *With probability at least $1 - \delta$, we have*

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

By the definition of $\hat{\sigma}_{k,h}$, the summation of $\varsigma_{k,h}^2$ will influence the summation of $\hat{\sigma}_{k,h}^2$. We need to make $\sum_{k=1}^K \sum_{h=1}^H \varsigma_{k,h}^2$ small such that it will not become dominant term in the upper bound. However, enlarging $\varsigma_{k,h}$ is required in some stages of some episodes, as stated in Remark 7.4 in the main paper. To address this dilemma, we build the following critical lemma which characterizes the conservatism of the elliptical potential, i.e., $\|\mathbf{x}_t\|_{\mathbf{Z}_{t-1}^{-1}}$ is small in most episodes, as detailed in Lemma D.8.

Lemma D.8 (Elliptical Potentials: You cannot have many big intervals). *Given $\lambda > 0$ and sequence $\{\mathbf{x}_t\}_{t=1}^T \subset \mathbb{R}^d$ with $\|\mathbf{x}_t\|_2 \leq L$ for all $t \in [T]$, define $\mathbf{Z}_t = \lambda \mathbf{I} + \sum_{i=1}^t \mathbf{x}_i \mathbf{x}_i^\top$ for $t \geq 1$ and $\mathbf{Z}_0 = \lambda \mathbf{I}$. During $[T]$, the number of times $\|\mathbf{x}_t\|_{\mathbf{Z}_{t-1}^{-1}} \geq c$ is at most*

$$\frac{3d}{\log(1+c^2)} \log \left(1 + \frac{L^2}{\lambda \log(1+c^2)} \right),$$

where $c > 0$ is a constant.

Proof. The proof of this lemma is firstly proposed at Exercise 19.3 in (Lattimore & Szepesvári, 2020) for the case of $C = 1$, i.e., Lemma F.6, we generalize it to the case with any positive constant C .

Let \mathcal{T} be the set of rounds t when $\|\mathbf{x}_t\|_{\mathbf{Z}_{t-1}^{-1}} \geq C$ for $t \in [T]$ and $\mathbf{Y}_t = \mathbf{Z}_0 + \sum_{i=1}^t \mathbb{1}\{i \in \mathcal{T}\} \mathbf{x}_i \mathbf{x}_i^\top$. Then

$$\begin{aligned} \left(\frac{d\lambda + |\mathcal{T}|L^2}{d} \right)^d &\geq \left(\frac{\text{trace}(\mathbf{Y}_T)}{d} \right)^d \\ &\geq \det(\mathbf{Y}_T) \\ &= \det(\mathbf{Z}_0) \prod_{t \in \mathcal{T}} \left(1 + \|\mathbf{x}_t\|_{\mathbf{Z}_{t-1}^{-1}}^2 \right) \\ &\geq \det(\mathbf{Z}_0) \prod_{t \in \mathcal{T}} \left(1 + \|\mathbf{x}_t\|_{\mathbf{Z}_{t-1}^{-1}}^2 \right) \\ &\geq \lambda^d (1+c^2)^{|\mathcal{T}|} \end{aligned}$$

Rearranging and taking the logarithm show that

$$|\mathcal{T}| \leq \frac{d}{\log(1+c^2)} \log \left(1 + \frac{|\mathcal{T}|L^2}{d\lambda} \right).$$

Abbreviate $x = d/\log(1+c^2)$ and $y = L^2/d\lambda$, which are both positive. Then

$$x \log(1 + y(3x \log(1 + xy))) \leq x \log(1 + 3x^2 y^2) \leq x \log(1 + xy)^3 = 3x \log(1 + xy).$$

Define $f(z) = z - x \log(1 + yz)$ for $z \geq 0$, we have $f'(z) = [1 + y(z - x)]/(1 + yz)$, which implies $f(z)$ is increasing if $1 - xy \geq 0$, or $f(z)$ is first decreasing then increasing, otherwise. Since $f(0) = 0$ and $f(3x \log(1 + xy)) \geq 0$, if $f(z) \leq 0$, we must have $z \leq 3x \log(1 + xy)$. In other words, $f(z)$ is increasing for $z \geq 3x \log(1 + xy)$. It then follows that

$$|\mathcal{T}| \leq 3x \log(1 + xy) = \frac{3d}{\log(1+c^2)} \log \left(1 + \frac{L^2}{\lambda \log(1+c^2)} \right).$$

□

The following Lemma is required to upper bounds the summation of offset term $U_{k,h}$.

Lemma D.9. *In Algorithm 1, under Υ , for any $k \in [K]$ and any $h \in [H]$, we have*

$$\left| \langle \hat{\boldsymbol{\mu}}_{k,h}(\hat{\mathbf{V}}_{k,h+1} - \check{\mathbf{V}}_{k,h+1}), \boldsymbol{\phi}(s_h^k, a_h^k) \rangle \right| \leq \bar{\beta} \|\boldsymbol{\phi}(s_h^k, a_h^k)\|_{\hat{\Lambda}_{k,h}^{-1}} + \mathbb{P}_h(\hat{\mathbf{V}}_{k,h+1} - \check{\mathbf{V}}_{k,h+1})(s_h^k, a_h^k) + \check{\beta} \|\boldsymbol{\phi}(s_h^k, a_h^k)\|_{\hat{\Lambda}_{k,h}^{-1}}$$

Proof.

$$\begin{aligned} & \left| \langle \hat{\boldsymbol{\mu}}_{k,h}(\hat{\mathbf{V}}_{k,h+1} - \check{\mathbf{V}}_{k,h+1}), \boldsymbol{\phi}(s_h^k, a_h^k) \rangle \right| = \left| \hat{\mathbb{P}}_{k,h} \hat{\mathbf{V}}_{k,h+1}(s_h^k, a_h^k) - \hat{\mathbb{P}}_{k,h} \check{\mathbf{V}}_{k,h+1}(s_h^k, a_h^k) \right| \\ & \leq \left| \hat{\mathbb{P}}_{k,h} \hat{\mathbf{V}}_{k,h+1}(s_h^k, a_h^k) - \mathbb{P}_h \hat{\mathbf{V}}_{k,h+1}(s_h^k, a_h^k) \right| + \left| \mathbb{P}_h \hat{\mathbf{V}}_{k,h+1}(s_h^k, a_h^k) - \mathbb{P}_h \check{\mathbf{V}}_{k,h+1}(s_h^k, a_h^k) \right| \\ & \quad + \left| \mathbb{P}_h \check{\mathbf{V}}_{k,h+1}(s_h^k, a_h^k) - \hat{\mathbb{P}}_{k,h} \check{\mathbf{V}}_{k,h+1}(s_h^k, a_h^k) \right| \\ & = \left| \langle (\boldsymbol{\mu}_h - \hat{\boldsymbol{\mu}}_{k,h}) \hat{\mathbf{V}}_{k,h+1}, \boldsymbol{\phi}(s_h^k, a_h^k) \rangle \right| + \left| \mathbb{P}_h(\hat{\mathbf{V}}_{k,h+1} - \check{\mathbf{V}}_{k,h+1})(s_h^k, a_h^k) \right| + \left| \langle (\boldsymbol{\mu}_h - \hat{\boldsymbol{\mu}}_{k,h}) \check{\mathbf{V}}_{k,h+1}, \boldsymbol{\phi}(s_h^k, a_h^k) \rangle \right| \\ & = \left| \langle (\boldsymbol{\mu}_h - \hat{\boldsymbol{\mu}}_{k,h}) \hat{\mathbf{V}}_{k,h+1}, \boldsymbol{\phi}(s_h^k, a_h^k) \rangle \right| + \mathbb{P}_h(\hat{\mathbf{V}}_{k,h+1} - \check{\mathbf{V}}_{k,h+1})(s_h^k, a_h^k) + \left| \langle (\boldsymbol{\mu}_h - \hat{\boldsymbol{\mu}}_{k,h}) \check{\mathbf{V}}_{k,h+1}, \boldsymbol{\phi}(s_h^k, a_h^k) \rangle \right| \\ & \leq \left\| (\boldsymbol{\mu}_h - \hat{\boldsymbol{\mu}}_{k,h}) \hat{\mathbf{V}}_{k,h+1} \right\|_{\hat{\Lambda}_{k,h}} \|\boldsymbol{\phi}(s_h^k, a_h^k)\|_{\hat{\Lambda}_{k,h}^{-1}} + \mathbb{P}_h(\hat{\mathbf{V}}_{k,h+1} - \check{\mathbf{V}}_{k,h+1})(s_h^k, a_h^k) \\ & \quad + \left\| (\boldsymbol{\mu}_h - \hat{\boldsymbol{\mu}}_{k,h}) \check{\mathbf{V}}_{k,h+1} \right\|_{\hat{\Lambda}_{k,h}} \|\boldsymbol{\phi}(s_h^k, a_h^k)\|_{\hat{\Lambda}_{k,h}^{-1}} \\ & \leq \bar{\beta} \|\boldsymbol{\phi}(s_h^k, a_h^k)\|_{\hat{\Lambda}_{k,h}^{-1}} + \mathbb{P}_h(\hat{\mathbf{V}}_{k,h+1} - \check{\mathbf{V}}_{k,h+1})(s_h^k, a_h^k) + \check{\beta} \|\boldsymbol{\phi}(s_h^k, a_h^k)\|_{\hat{\Lambda}_{k,h}^{-1}} \\ & \leq \bar{\beta} \|\boldsymbol{\phi}(s_h^k, a_h^k)\|_{\hat{\Lambda}_{k,h}^{-1}} + \mathbb{P}_h(\hat{\mathbf{V}}_{k,h+1} - \check{\mathbf{V}}_{k,h+1})(s_h^k, a_h^k) + \check{\beta} \|\boldsymbol{\phi}(s_h^k, a_h^k)\|_{\hat{\Lambda}_{k,h}^{-1}}, \end{aligned}$$

where the first inequality holds due to triangle inequality, the second equality holds since $\hat{\mathbf{V}}_{k,h+1}(\cdot) \geq \check{\mathbf{V}}_{k,h+1}(\cdot)$ under Υ by Lemma D.1, the second inequality holds due to Cauchy-Schwarz inequality, the fourth inequality holds since under Υ we have $\boldsymbol{\mu}_h \in \bar{\mathcal{C}}_{k,h} \cap \check{\mathcal{C}}_{k,h}$, and the last inequality holds due to Lemma D.2 under Υ . \square

Now we are ready to upper bounds $\sum_{k=1}^K \sum_{h=1}^H \hat{\sigma}_{k,h}^2$ in Lemma D.10 under the high probability event $\Upsilon \cap \Xi_2 \cap \Xi_3$.

Lemma D.10. *In Algorithm 1, under event $\Upsilon \cap \Xi_2 \cap \Xi_3$, we have*

$$\begin{aligned} \sum_{k=1}^K \sum_{h=1}^H \hat{\sigma}_{k,h}^2 & \leq 8HT + \frac{6H^3 d^4}{\log(1+c^2)} \log \left(1 + \frac{d}{\lambda H \log(1+c^2)} \right) + 6H^3 \log \left(\frac{1}{\delta} \right) \\ & \quad + 8H^3(6+d) \sqrt{2T \log \left(\frac{H}{\delta} \right)} + \frac{12H^4 d}{\log(2)} \log \left(1 + \frac{1}{\lambda \log(2)} \right) \\ & \quad + 8Hd \left[H \left(2H(d+6)\hat{\beta} + d\check{\beta} + 4\bar{\beta} + (H+1)(d+4)\check{\beta} \right) + \check{\beta} \right]^2 \log \left(1 + \frac{K}{Hd\lambda} \right), \end{aligned}$$

where $c = 1/(Hd^3)$.

Proof. Initially, by definition of $\hat{\sigma}_{k,h}$ in Algorithm 1, we have

$$\sum_{k=1}^K \sum_{h=1}^H \hat{\sigma}_{k,h}^2 \leq \underbrace{\sum_{k=1}^K \sum_{h=1}^H \varsigma_k^2}_{I_1} + \underbrace{\sum_{k=1}^K \sum_{h=1}^H dE_{k,h}}_{I_2} + \underbrace{\sum_{k=1}^K \sum_{h=1}^H \left[\hat{\mathbb{V}}_{k,h} \hat{\mathbf{V}}_{k,h+1}(s_h^k, a_h^k) + U_{k,h} \right]}_{I_3}, \quad (50)$$

Bounding I_1 : Denote $c = 1/(Hd^3)$. For fixed $h \in [H]$, set \mathbf{x}_k as $\tilde{\sigma}_{k,h}^{-1} \boldsymbol{\phi}(s_h^k, a_h^k)$ in Lemma D.8. Then, there are at most $3d \log[1 + d/(\lambda H \log(1+c^2))]/\log(1+c^2)$ episodes that $\|\tilde{\sigma}_{k,h}^{-1} \boldsymbol{\phi}(s_h^k, a_h^k)\|_{\hat{\Lambda}_{k,h}^{-1}} \geq c$ such that $\varsigma_k = H\sqrt{d^3}$. Thus, we obtain

$$I_1 \leq \sum_{k=1}^K \sum_{h=1}^H H + H \cdot H^2 d^3 \frac{3d}{\log(1+c^2)} \log \left(1 + \frac{d}{\lambda H \log(1+c^2)} \right). \quad (51)$$

Bounding I_2 : Since $\mu_h \in \hat{\mathcal{C}}_{k,h} \cap \check{\mathcal{C}}_{k,h}$ under Υ , we have

$$\begin{aligned} \langle \hat{\mu}_{k,h} \hat{V}_{k,h+1}, \phi(s_h^k, a_h^k) \rangle &= \hat{\mathbb{P}}_{k,h} \hat{V}_{k,h+1}(s_h^k, a_h^k) \leq \mathbb{P}_h \hat{V}_{k,h+1}(s_h^k, a_h^k) + \hat{\beta} \|\phi(s_h^k, a_h^k)\|_{\hat{\Lambda}_{k,h}^{-1}} \\ \langle \hat{\mu}_{k,h} \check{V}_{k,h+1}, \phi(s_h^k, a_h^k) \rangle &= \hat{\mathbb{P}}_{k,h} \check{V}_{k,h+1}(s_h^k, a_h^k) \geq \mathbb{P}_h \check{V}_{k,h+1}(s_h^k, a_h^k) - \check{\beta} \|\phi(s_h^k, a_h^k)\|_{\hat{\Lambda}_{k,h}^{-1}} \end{aligned} \quad (52)$$

Combing Eq. (52) and the definition of $E_{k,h}$ in Eq. (36) gives

$$\begin{aligned} I_2 &\leq Hd \sum_{k=1}^K \sum_{h=1}^H [\mathbb{P}_h(\hat{V}_{k,h+1} - \check{V}_{k,h+1})](s_h^k, a_h^k) + Hd \sum_{k=1}^K \sum_{h=1}^H \min \left\{ 2(\hat{\beta} + \check{\beta}) \|\phi(s_h^k, a_h^k)\|_{\hat{\Lambda}_{k,h}^{-1}}, H \right\} \\ &\leq Hd \sum_{k=1}^K \sum_{h=1}^H [\mathbb{P}_h(\hat{V}_{k,h+1} - \check{V}_{k,h+1})](s_h^k, a_h^k) + 2(\hat{\beta} + \check{\beta}) Hd \sum_{k=1}^K \sum_{h=1}^H \hat{\sigma}_{k,h} \min \left\{ \|\hat{\sigma}_{k,h}^{-1} \phi(s_h^k, a_h^k)\|_{\hat{\Lambda}_{k,h}^{-1}}, 1 \right\}, \end{aligned} \quad (53)$$

where the second inequality holds since $\check{\beta} \hat{\sigma}_{k,h} \geq \sqrt{Hd^2} \sqrt{H} \geq H$.

Bounding I_3 : I_3 can be bounded by

$$\begin{aligned} I_3 &\leq \underbrace{\sum_{k=1}^K \sum_{h=1}^H [\mathbb{V}_h V_{h+1}^*](s_h^k, a_h^k) - [\mathbb{V}_h V_{h+1}^{\pi^k}](s_h^k, a_h^k)}_{J_1} + 2 \underbrace{\sum_{k=1}^K \sum_{h=1}^H U_{k,h}}_{J_2} + \underbrace{\sum_{k=1}^K \sum_{h=1}^H [\mathbb{V}_h V_{h+1}^{\pi^k}](s_h^k, a_h^k)}_{J_3} \\ &\quad + \underbrace{\sum_{k=1}^K \sum_{h=1}^H \left[[\hat{\mathbb{V}}_{k,h} \hat{V}_{k,h+1}](s_h^k, a_h^k) - [\mathbb{V}_h V_{h+1}^*](s_h^k, a_h^k) - U_{k,h} \right]}_{J_4} \end{aligned}$$

Bounding J_1

$$\begin{aligned} J_1 &= \sum_{k=1}^K \sum_{h=1}^H \left[\mathbb{P}_h V_{h+1}^{*2}(s_h^k, a_h^k) - [\mathbb{P}_h V_{h+1}^*(s_h^k, a_h^k)]^2 \right] - \left[\mathbb{P}_h V_{h+1}^{\pi^k 2}(s_h^k, a_h^k) - [\mathbb{P}_h V_{h+1}^{\pi^k}(s_h^k, a_h^k)]^2 \right] \\ &\leq \sum_{k=1}^K \sum_{h=1}^H \mathbb{P}_h V_{h+1}^{*2}(s_h^k, a_h^k) - \mathbb{P}_h V_{h+1}^{\pi^k 2}(s_h^k, a_h^k) \\ &\leq 2H \sum_{k=1}^K \sum_{h=1}^H [\mathbb{P}_h (V_{h+1}^* - V_{h+1}^{\pi^k})](s_h^k, a_h^k) \\ &\leq 2H \sum_{k=1}^K \sum_{h=1}^H [\mathbb{P}_h (\hat{V}_{k,h+1} - V_{h+1}^{\pi^k})](s_h^k, a_h^k), \end{aligned} \quad (54)$$

where the first inequality holds since $V_{h+1}^*(\cdot) \geq V_{h+1}^{\pi^k}$, the second inequality holds since $V_{h+1}^{\pi^k} \leq V_{h+1}^*(\cdot) \leq H$, and the last inequality holds since $\hat{V}_{k,h+1}(\cdot) \geq \hat{V}_{k,h+1}(\cdot) \geq V_{h+1}^*(\cdot)$ under Υ by Lemma D.1 and Lemma D.2.

Bounding J_2 By the definition of $U_{k,h}$ in Eq. (35) and Lemma D.9, we have

$$\begin{aligned} J_2 &\leq \sum_{k=1}^K \sum_{h=1}^H \left\{ 2\tilde{\beta} \|\phi(s_h^k, a_h^k)\|_{\hat{\Lambda}_{k,h}^{-1}}, 2H^2 \right\} + 4H \sum_{k=1}^K \sum_{h=1}^H [\mathbb{P}_h (\hat{V}_{k,h+1} - \check{V}_{k,h+1})](s_h^k, a_h^k) \\ &\quad + \sum_{k=1}^K \sum_{h=1}^H \min \left\{ 8H(\tilde{\beta} + \check{\beta}) \|\phi(s_h^k, a_h^k)\|_{\hat{\Lambda}_{k,h}^{-1}}, 2H^2 \right\} \\ &\leq [2\tilde{\beta} + 8H(\tilde{\beta} + \check{\beta})] \sum_{k=1}^K \sum_{h=1}^H \hat{\sigma}_{k,h} \left\{ \|\hat{\sigma}_{k,h}^{-1} \phi(s_h^k, a_h^k)\|_{\hat{\Lambda}_{k,h}^{-1}}, 1 \right\} + 4H \sum_{k=1}^K \sum_{h=1}^H [\mathbb{P}_h (\hat{V}_{k,h+1} - \check{V}_{k,h+1})](s_h^k, a_h^k), \end{aligned} \quad (55)$$

where the second inequality holds since $\tilde{\beta} \hat{\sigma}_{k,h} \geq \sqrt{H^3} \cdot \sqrt{H} \geq H^2$, $8H(\tilde{\beta} + \check{\beta}) \hat{\sigma}_{k,h} \geq 8H\sqrt{H} \cdot \sqrt{H} \geq 2H^2$.

Bounding J_3 Since Ξ_3 holds, we have

$$J_3 \leq 3 \left[HT + H^3 \log \left(\frac{1}{\delta} \right) \right] \quad (56)$$

Bounding J_4 Due to Lemma 7.5, we have

$$J_4 \leq 0 \quad (57)$$

Putting Together Initially, we have

$$\begin{aligned} & \sum_{k=1}^K \sum_{h=h'}^H \hat{\sigma}_{k,h} \min \left\{ \|\hat{\sigma}_{k,h}^{-1} \phi(s_h^k, a_h^k)\|_{\hat{\Lambda}_{k,h}^{-1}}, 1 \right\} \\ & \leq \sqrt{\sum_{k=1}^K \sum_{h=1}^H \hat{\sigma}_{k,h}^2} \sqrt{\sum_{k=1}^K \sum_{h=1}^H \min \left\{ \|\hat{\sigma}_{k,h}^{-1} \phi(s_h^k, a_h^k)\|_{\hat{\Lambda}_{k,h}^{-1}}^2, 1 \right\}} \\ & \leq \sqrt{\sum_{k=1}^K \sum_{h=1}^H \hat{\sigma}_{k,h}^2} \sqrt{H \cdot 2d \log \left(1 + \frac{K}{Hd\lambda} \right)}, \end{aligned} \quad (58)$$

where the first inequality holds due to Cauchy-Schwarz inequality, and the second inequality holds due to Lemma F.5 with the fact that $\|\hat{\sigma}_{k,h}^{-1} \phi(s_h^k, a_h^k)\|_2 \leq 1/\sqrt{H}$.

Subsequently, combining Eq. (53), (54), (55) gives

$$\begin{aligned} I_2 + J_1 + J_2 & \leq 2H \sum_{k=1}^K \sum_{h=1}^H [\mathbb{P}_h(\hat{V}_{k,h+1} - V_{h+1}^{\pi^k})](s_h^k, a_h^k) + H(4+d) \sum_{k=1}^K \sum_{h=1}^H [\mathbb{P}_h(\hat{V}_{k,h+1} - \check{V}_{k,h+1})](s_h^k, a_h^k) \\ & \quad + 2 \left[H(d\hat{\beta} + 4\bar{\beta} + (4+d)\check{\beta}) + \tilde{\beta} \right] \sum_{k=1}^K \sum_{h=1}^H \hat{\sigma}_{k,h} \left\{ \|\hat{\sigma}_{k,h}^{-1} \phi(s_h^k, a_h^k)\|_{\hat{\Lambda}_{k,h}^{-1}}, 1 \right\} \\ & \leq 4H^3(6+d) \sqrt{2T \log \left(\frac{H}{\delta} \right)} + \frac{6H^4d}{\log(2)} \log \left(1 + \frac{1}{\lambda \log(2)} \right) \\ & \quad + 2 \left[H \left(2H(d+6)\hat{\beta} + d\hat{\beta} + 4\bar{\beta} + (H+1)(d+4)\check{\beta} \right) + \tilde{\beta} \right] \sum_{k=1}^K \sum_{h=1}^H \hat{\sigma}_{k,h} \left\{ \|\hat{\sigma}_{k,h}^{-1} \phi(s_h^k, a_h^k)\|_{\hat{\Lambda}_{k,h}^{-1}}, 1 \right\} \\ & \leq 4H^3(6+d) \sqrt{2T \log \left(\frac{H}{\delta} \right)} + \frac{6H^4d}{\log(2)} \log \left(1 + \frac{1}{\lambda \log(2)} \right) \\ & \quad + 2 \left[H \left(2H(d+6)\hat{\beta} + d\hat{\beta} + 4\bar{\beta} + (H+1)(d+4)\check{\beta} \right) + \tilde{\beta} \right] \sqrt{\sum_{k=1}^K \sum_{h=1}^H \hat{\sigma}_{k,h}^2} \sqrt{2Hd \log \left(1 + \frac{K}{Hd\lambda} \right)} \end{aligned}$$

where the first inequality holds by Lemma D.5 under $\Upsilon \cap \Xi_1$ and Lemma D.6 under $\Upsilon \cap \Xi_2$, and the second inequality holds due to Eq. (58). Further considering Eq. (51), (56), (57) gives

$$\begin{aligned} \sum_{k=1}^K \sum_{h=1}^H \hat{\sigma}_{k,h}^2 & \leq 4HT + \frac{3H^3d^4}{\log(1+c^2)} \log \left(1 + \frac{d}{\lambda H \log(1+c^2)} \right) + 3H^3 \log \left(\frac{1}{\delta} \right) \\ & \quad + 4H^3(6+d) \sqrt{2T \log \left(\frac{H}{\delta} \right)} + \frac{6H^4d}{\log(2)} \log \left(1 + \frac{1}{\lambda \log(2)} \right) \\ & \quad + 2 \left[H \left(2H(d+6)\hat{\beta} + d\hat{\beta} + 4\bar{\beta} + (H+1)(d+4)\check{\beta} \right) + \tilde{\beta} \right] \sqrt{\sum_{k=1}^K \sum_{h=1}^H \hat{\sigma}_{k,h}^2} \sqrt{2Hd \log \left(1 + \frac{K}{Hd\lambda} \right)} \end{aligned} \quad (59)$$

Besides, for any $x, a, b \geq 0$, if $x \leq a\sqrt{x} + b$, then we have $\sqrt{x} \leq \sqrt{b + a^2/4} + \sqrt{a^2/4} \leq \sqrt{2(b + a^2/4 + a^2/4)}$, i.e., $x \leq 2b + a^2$. Thus, Eq. (59) implies the final conclusion. \square

D.4. Proof of Theorem 6.1

Putting these building blocks together, we are finally ready to give high probability upper bound on the regret in this subsection, which is based on high probability event $\Upsilon \cap \Xi_1 \cap \Xi_2 \cap \Xi_3$.

Proof of Theorem 6.1. By construction, taking a union bound, we have that with probability $1 - 10\delta$, $\Upsilon \cap \Xi_1 \cap \Xi_2 \cap \Xi_3$ holds. In the remainder of the proof, assume that we are conditioning on this event. Initially, we have

$$\begin{aligned} \text{Regret}(K) &= \sum_{k=1}^K [V_1^*(s_1^k) - V_1^{\pi^k}(s_1^k)] \leq \sum_{k=1}^K [\hat{V}_{k,1}(s_1^k) - V_1^{\pi^k}(s_1^k)] \\ &\leq 4\hat{\beta} \sqrt{\sum_{k=1}^K \sum_{h=1}^H \hat{\sigma}_{k,h}^2} \sqrt{H \cdot 2d \log\left(1 + \frac{K}{Hd\lambda}\right)} + \frac{3H^2d}{\log(2)} \log\left(1 + \frac{1}{\lambda H \log(2)}\right) + 2H \sqrt{2T \log\left(\frac{H}{\delta}\right)}, \end{aligned} \quad (60)$$

where the first inequality holds since $\hat{V}_{k,1}(\cdot) \geq \hat{V}_{k,1}(\cdot) \geq V_1^*(\cdot)$ under Υ by Lemma D.1 and Lemma D.2, and the second inequality holds due to Lemma D.5.

Then, by Lemma D.10, we have

$$\begin{aligned} \sum_{k=1}^K \sum_{h=1}^H \hat{\sigma}_{k,h}^2 &\leq 8HT + \frac{6H^3d^4}{\log(1+c^2)} \log\left(1 + \frac{d}{\lambda H \log(1+c^2)}\right) + 6H^3 \log\left(\frac{1}{\delta}\right) \\ &\quad + 8H^3(6+d) \sqrt{2T \log\left(\frac{H}{\delta}\right)} + \frac{12H^4d}{\log(2)} \log\left(1 + \frac{1}{\lambda \log(2)}\right) \\ &\quad + 8Hd \left[H \left(2H(d+6)\hat{\beta} + d\dot{\beta} + 4\bar{\beta} + (H+1)(d+4)\check{\beta} \right) + \tilde{\beta} \right]^2 \log\left(1 + \frac{K}{Hd\lambda}\right), \end{aligned}$$

under $\Upsilon \cap \Xi_1 \cap \Xi_2$, where $c = 1/(Hd^3)$.

On the one hand, since $1/\log(1+1/x) \leq 2x$ for any $x \geq 1$, we have

$$\frac{1}{\log(1+c^2)} \leq \frac{1}{\log\left(1 + \frac{1}{H^2d^6}\right)} \leq 2H^2d^6$$

for $H^2d^6 \geq 1$. On the other hand, setting $\lambda = 1/(H^2\sqrt{d})$ gives $\hat{\beta} = O(\sqrt{d}\log(T))$, $\bar{\beta} = O(d\sqrt{H\log(T)})$, $\dot{\beta} = O(d\sqrt{H\log(T)})$, $\check{\beta} = O(d\sqrt{H^3\log(T)})$ and $\tilde{\beta} = O(d\sqrt{H\log(T)})$, where $T = KH$.

Thus, by some basic manipulations, we obtain

$$\begin{aligned} \sum_{k=1}^K \sum_{h=1}^H \hat{\sigma}_{k,h}^2 &\leq O\left(HT + H^5d^{10} + H^3d\sqrt{T\log(T)} + H^6d^5\log^2(T)\right) \\ &\leq O\left(HT + H^5d^{10} + H^5d^2\log(T) + H^6d^5\log^2(T)\right) \\ &\leq O\left(HT + H^5d^{10} + H^6d^5\log^2(T)\right), \end{aligned} \quad (61)$$

where the second inequality holds since $H^3d\sqrt{T\log(T)} \leq [HT + H^5d^2\log(T)]/2$. Substituting Eq. (61) in (60) gives

$$\begin{aligned} &\text{Regret}(K) \\ &\leq O\left(\sqrt{d}\log(T)\sqrt{HT + H^5d^{10} + H^6d^5\log^2(T)}\sqrt{Hd\log(T)} + \frac{3H^2d}{\log(2)} \log\left(1 + \frac{H\sqrt{d}}{\log(2)}\right) + 2H\sqrt{2T \log\left(\frac{H}{\delta}\right)}\right) \end{aligned}$$

$$\begin{aligned} &\leq O\left(Hd\sqrt{T\log^3(T)} + H^3d^6\sqrt{\log^3(T)} + \sqrt{H^7d^7\log^5(T)}\right) \\ &= \tilde{O}\left(Hd\sqrt{T} + H^3d^6 + \sqrt{H^7d^7}\right), \end{aligned}$$

where the second inequality holds by dropping lower order terms and the fact that $\sqrt{a+b+c} \leq \sqrt{a} + \sqrt{b} + \sqrt{c}$ for any $a, b, c > 0$. \square

E. Lower Bound

Remark 23 in Appendix of (Zhou et al., 2021) constructs a hard-to-learn linear MDP instance, which shares the same order of regret lower bound as a hard-to-learn linear mixture MDP instance. According to Theorem 8 in (Zhou et al., 2021), the linear mixture MDP has a regret lower bound of $\Omega(Hd\sqrt{T})$, which means the linear MDP with known reward has the same regret lower bound as well. We present the construction of this hard-to-learn linear MDP from (Zhou et al., 2021) in this section for completeness. This hard-to-learn linear MDP can be regarded as an extension of hard instances in linear bandits literature (Dani et al., 2008; Lattimore & Szepesvári, 2020). We first illustrate the structure of this MDP and then present the specific linear parametrization.

Hard MDP Instance This MDP instance is denoted as $\mathcal{M} = \{\mathcal{S}, \mathcal{A}, H, \{\mathbb{P}_1\}_h, \{r_h\}_h\}$. The state space \mathcal{S} consists of states s_1, \dots, s_{H+2} such that $|\mathcal{S}| = H + 2$. There are 2^{d-1} actions and $\mathcal{A} = \{-1, 1\}^{d-1}$ such that each action $\mathbf{a} \in \mathcal{A}$ is denoted in vector form.

- Reward: For any stage $h \in [H + 2]$, only transitions originating at s_{H+2} incurs a reward.
- Transition: s_{H+1} and s_{H+2} are absorbing regardless of what action is taken. For state s_i with $i \leq H$, the transition probability is given as

$$\mathbb{P}_h(s' | s_i, \mathbf{a}) = \begin{cases} \iota + \langle \boldsymbol{\mu}_h, \mathbf{a} \rangle, & s' = s_{H+2} \\ 1 - (\iota + \langle \boldsymbol{\mu}_h, \mathbf{a} \rangle), & s' = s_{i+1} \\ 0, & \text{Otherwise} \end{cases}$$

where $\iota = 1/H$ and $\boldsymbol{\mu}_h \in \{-\Delta, \Delta\}^{d-1}$ with $\Delta = \sqrt{\iota/K}/(4\sqrt{2})$ to make the probabilities are well-defined. The transition of this MDP is detailed in Figure 1.




Figure 1. The transition matrix \mathbb{P}_h of the hard-to-learn MDP.

Linear Parametrization Then, we specify the linear parametrization of this MDP. For each $h \in [H]$, the transition probability matrix \mathbb{P}_h and the reward function r_h are defined as $\mathbb{P}_h(s' | s, \mathbf{a}) = \langle \phi(s, \mathbf{a}), \boldsymbol{\mu}_h(s') \rangle$ and $r_h(s, \mathbf{a}) =$

$\langle \phi(s, \mathbf{a}), \boldsymbol{\theta}_h \rangle$, where $\phi(s, \mathbf{a}) \in \mathbb{R}^{d+1}$ is the known feature mapping, $\boldsymbol{\mu}_h = (\mu_h(s_1), \dots, \mu_h(s_{H+2})) \in \mathbb{R}^{(d+1) \times (H+2)}$ and $\boldsymbol{\theta}_h \in \mathbb{R}^{d+1}$ are unknown parameters in linear MDPs. Here, $\phi(s, \mathbf{a}), \boldsymbol{\mu}_h, \boldsymbol{\theta}_h$ are specified as:

$$\begin{aligned} \phi(s, \mathbf{a}) &= \begin{cases} (\alpha, \beta \mathbf{a}^\top, 0)^\top, & s = s_h, h \in [H+1] \\ (0, \mathbf{0}^\top, 1)^\top, & s = s_{H+2} \end{cases} \\ \boldsymbol{\mu}_h(s') &= \begin{cases} ((1-\iota)/\alpha, -\boldsymbol{\mu}_h^\top/\beta, 0)^\top, & s' = s_{H+1} \\ (\iota/\alpha, \boldsymbol{\mu}_h^\top/\beta, 1)^\top, & s' = s_{H+2} \\ \mathbf{0}, & \text{otherwise} \end{cases} \\ \boldsymbol{\theta}_h &= (\mathbf{0}^\top, 1)^\top \end{aligned}$$

where $\alpha = \sqrt{1/(1 + \Delta(d-1))}$, $\beta = \sqrt{\Delta/(1 + \Delta(d-1))}$.

Norm Assumption We check the norm assumption of linear MDPs in the following:

1. For $s = s_h$ where $h \in [H+1]$, $\|\phi(s, \mathbf{a})\|_2 = \sqrt{\alpha^2 + (d-1)\beta^2} = 1$ and $\|\phi(s_{H+2}, \mathbf{a})\|_2 = 1$. Thus, $\|\phi(s, \mathbf{a})\|_2 \leq 1$ for any $(s, \mathbf{a}) \in \mathcal{S} \times \mathcal{A}$.
2. For any $\mathbf{v} = (v_1, \dots, v_{H+2}) \in \mathbb{R}^{H+2}$ such that $\|\mathbf{v}\|_\infty \leq 1$, we have

$$\|\boldsymbol{\mu}_h \mathbf{v}\|_2^2 = \left[\frac{v_1(1-\iota)}{\alpha} + \frac{v_1 \iota}{\alpha} \right]^2 + v_{H+2}^2 \leq \frac{1}{\alpha^2} + 1 = [1 + \Delta(d-1)]^2 + 1 = \left[1 + \frac{\sqrt{\delta/K}}{4\sqrt{2}}(d-1) \right]^2 + 1 \leq d+1,$$

where the last inequality holds by assuming episode number $K \geq (d-1)/(32H(\sqrt{d}-1))$. Thus, $\|\boldsymbol{\mu}_h \mathbf{v}\|_2 \leq \sqrt{d+1}$ for any $h \in [H]$.

3. In addition, $\|\boldsymbol{\theta}_h\|_2 \leq 1 \leq \sqrt{d+1}$ for any $h \in [H]$.

Lower Bound The constructed linear MDP above has the same state space \mathcal{S} , action space \mathcal{A} , episode length H , reward function $\{r_h\}_{h \in [H]}$ and transition probability $\{\mathbb{P}_h\}_{h \in [H]}$ as the constructed hard-to-learn linear mixture MDP in Appendix E. of (Zhou et al., 2021), which shares the same regret lower bound $\Omega(Hd\sqrt{T})$ as shown in Theorem 8 in (Zhou et al., 2021) and is formalized in the following theorem.

Lemma E.1 (Lower bound of linear MDPs). *Let $d > 1$ and suppose $K \geq \max\{(d-1)^2 H/2, (d-1)/(32H(d-1))\}$, $d \geq 4$, $H \geq 3$. Then for any algorithm there exists an episodic linear MDP parameterized by $\{\boldsymbol{\mu}_h\}_{h \in [H]}$, $\{\boldsymbol{\theta}_h\}_{h \in [H]}$ and satisfy the norm assumption given in Definition 3.1, such that the expected regret is lower bounded as follows:*

$$\mathbb{E}[\text{Regret}(K)] \geq \Omega(Hd\sqrt{T}),$$

where $T = KH$ and the expectation is taken over the probability distribution generated by the interconnection of the algorithm and the MDP.

Proof. The proof is the same as that of Theorem 8 in (Zhou et al., 2021), except for changing B to \sqrt{d} to satisfy the norm assumption of linear MDPs. \square

F. Auxiliary Lemmas

In this section, we give some auxiliary lemmas which serve as the preliminary for the proof above. We also include some other lemmas that are unnecessary for our theoretical analysis but can help readers be more familiar with related works. In general, these lemmas are categorized into four subsections:

- Appendix F.1 for some concentration inequalities;
- Appendix F.2 for properties related to elliptical potentials;
- Appendix F.3 presents some useful properties for linear MDPs;
- Appendix F.4 builds the covering number for covering net over some function classes of our interests.

F.1. Concentration Inequality

In this subsection, Lemma F.1 presents the Azuma-Hoeffding inequality, Lemma F.2 presents the Freedman's inequality in (Freedman, 1975), Lemma F.3 a Hoeffding-type self-normalized bound, and Lemma F.4 presents the full version of Theorem 7.1 in main paper.

Lemma F.1 (Azuma-Hoeffding Inequality). *Let $\{x_i\}_{i=1}^n$ be a martingale difference sequence with respect to a filtration $\{\mathcal{G}_i\}_{i=1}^{n+1}$ such that $|x_i| \leq M$ almost surely. That is, x_i is \mathcal{G}_{i+1} -measurable and $\mathbb{E}[x_i | \mathcal{G}_i] = 0$ a.s. Then for any $0 < \delta < 1$, with probability at least $1 - \delta$,*

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

Lemma F.2 (Freedman's Inequality, (Freedman, 1975)). *Let $\{x_i, \mathcal{F}_i\}$ be a martingale difference sequence with $\forall i \geq 1$, $\mathbb{E}(x_i | \mathcal{F}_{i-1}) = 0$, $\mathbb{E}(x_i^2 | \mathcal{F}_{i-1}) = \sigma_i^2$, $V_i^2 = \sum_{j=1}^i \sigma_j^2$. Furthermore, assume that $\mathbb{P}(|x_i| \leq c | \mathcal{F}_{i-1}) = 1$ for any $0 < c < \infty$.*

Then, for fixed $t \geq 1$ and any $\delta > 0$, with probability at least $1 - \delta$, we have:

$$\sum_{i=1}^t d_i \leq \sqrt{2V_t^2 \log(1/\delta)} + \frac{2}{3}c \log(1/\delta).$$

Lemma F.3 (Hoeffding inequality for vector-valued martingales, Theorem 1 in (Abbasi-Yadkori et al., 2011)). *Let $\{\mathcal{G}_t\}_{t=1}^\infty$ be a filtration, $\{\mathbf{x}_t, \eta_t\}_{t \geq 1}$ be a stochastic process so that $\mathbf{x}_t \in \mathbb{R}^d$ is \mathcal{G}_t -measurable and $\eta_t \in \mathbb{R}$ is \mathcal{G}_{t+1} -measurable.*

Denote $\mathbf{Z}_t = \lambda \mathbf{I} + \sum_{i=1}^t \mathbf{x}_i \mathbf{x}_i^\top$ for $t \geq 1$ and $\mathbf{Z}_0 = \lambda \mathbf{I}$. If $\|\mathbf{x}_t\|_2 \leq L$, and η_t satisfies

$$\mathbb{E}[\eta_t | \mathcal{G}_t] = 0, \quad |\eta_t| \leq R$$

for all $t \geq 1$. Then, for any $0 < \delta < 1$, with probability at least $1 - \delta$ we have:

$$\forall t > 0, \left\| \sum_{i=1}^t \mathbf{x}_i \eta_i \right\|_{\mathbf{Z}_t^{-1}} \leq R\sqrt{d \log(1 + tL^2/d\lambda) + \log(1/\delta)}.$$

Lemma F.4 (Bernstein inequality for vector-valued martingales, full version of Theorem 7.1). *Let $\{\mathcal{G}_t\}_{t=1}^\infty$ be a filtration, $\{\mathbf{x}_t, \eta_t\}_{t \geq 1}$ be a stochastic process so that $\mathbf{x}_t \in \mathbb{R}^d$ is \mathcal{G}_t -measurable and $\eta_t \in \mathbb{R}$ is \mathcal{G}_{t+1} -measurable.*

If $\|\mathbf{x}_t\|_2 \leq L$, and η_t satisfies

$$\mathbb{E}[\eta_t | \mathcal{G}_t] = 0, \quad \mathbb{E}[\eta_t^2 | \mathcal{G}_t] \leq \sigma^2, \quad \left| \eta_t \cdot \min \left\{ 1, \|\mathbf{x}_t\|_{\mathbf{Z}_{t-1}^{-1}} \right\} \right| \leq R$$

for all $t \geq 1$. Then, for any $0 < \delta < 1$, with probability at least $1 - \delta$ we have:

$$\forall t > 0, \left\| \sum_{i=1}^t \mathbf{x}_i \eta_i \right\|_{\mathbf{Z}_t^{-1}} \leq 8\sigma\sqrt{d \log(1 + tL^2/(d\lambda)) \log(4t^2/\delta)} + 4R \log(4t^2/\delta)$$

where $\mathbf{Z}_t = \lambda \mathbf{I} + \sum_{i=1}^t \mathbf{x}_i \mathbf{x}_i^\top$ for $t \geq 1$ and $\mathbf{Z}_0 = \lambda \mathbf{I}$.

F.2. Elliptical Potentials

In this subsection, we present Lemma F.5 in (Abbasi-Yadkori et al., 2011), which is an important lemma for building $\sqrt{O}(\sqrt{T})$ worst case regret for many algorithms for linear bandits or RL with linear function approximation. In addition, we also present an important property, i.e., Lemma F.6, about elliptical potentials (Exercise 19.3 in (Lattimore & Szepesvári, 2020)), which states that one cannot have more than $O(d)$ big intervals. Lemma F.6 is further generalized in Lemma D.8 in Appendix D.3.

Lemma F.5 (Lemma 11, (Abbasi-Yadkori et al., 2011)). *Given $\lambda > 0$ and sequence $\{\mathbf{x}_t\}_{t=1}^T \subset \mathbb{R}^d$ with $\|\mathbf{x}_t\|_2 \leq L$ for all $t \in [T]$, define $\mathbf{Z}_t = \lambda \mathbf{I} + \sum_{i=1}^t \mathbf{x}_i \mathbf{x}_i^\top$ for $t \geq 1$ and $\mathbf{Z}_0 = \lambda \mathbf{I}$. We have*

$$\sum_{t=1}^T \min \left\{ 1, \|\mathbf{x}_t\|_{\mathbf{Z}_{t-1}^{-1}}^2 \right\} \leq 2d \log \left(1 + \frac{TL^2}{d\lambda} \right).$$

Lemma F.6 (Elliptical Potentials: You cannot have more than $O(d)$ big intervals. Exercise 19.3 in (Lattimore & Szepesvári, 2020)). *Given $\lambda > 0$ and sequence $\{\mathbf{x}_t\}_{t=1}^T \subset \mathbb{R}^d$ with $\|\mathbf{x}_t\|_2 \leq L$ for all $t \in [T]$, define $\mathbf{Z}_t = \lambda \mathbf{I} + \sum_{i=1}^t \mathbf{x}_i \mathbf{x}_i^\top$ for $t \geq 1$ and $\mathbf{Z}_0 = \lambda \mathbf{I}$. The number of times $\|\mathbf{x}_t\|_{\mathbf{Z}_{t-1}^{-1}}^2 \geq 1$ is at most*

$$\frac{3d}{\log(2)} \log \left(1 + \frac{L^2}{\lambda \log(2)} \right).$$

F.3. Linear MDP Property

This subsection gives some indirect results about the estimated parameter $\hat{\boldsymbol{\mu}}_{k,h}$ in linear MDPs.

Lemma F.7. *In Algorithm 1, for any $k \in [K]$ and any $h \in [H]$, we have:*

$$\hat{\boldsymbol{\mu}}_{k,h} - \boldsymbol{\mu}_h = \hat{\Lambda}_{k,h}^{-1} \left[-\lambda \boldsymbol{\mu}_h + \sum_{i=1}^{k-1} \hat{\sigma}_{i,h}^{-2} \boldsymbol{\phi}(s_h^i, a_h^i) \boldsymbol{\epsilon}_h^i \right] \quad (62)$$

Proof. We start from the closed-form solution of $\hat{\boldsymbol{\mu}}_{k,h}$:

$$\begin{aligned} \hat{\boldsymbol{\mu}}_{k,h} &= \hat{\Lambda}_{k,h}^{-1} \sum_{i=1}^{k-1} \hat{\sigma}_{i,h}^{-2} \boldsymbol{\phi}(s_h^i, a_h^i) \boldsymbol{\delta}(s_{h+1}^i)^\top = \hat{\Lambda}_{k,h}^{-1} \sum_{i=1}^{k-1} \hat{\sigma}_{i,h}^{-2} \boldsymbol{\phi}(s_h^i, a_h^i) \left(\mathbb{P}_h(\cdot | s_h^i, a_h^i)^\top + \boldsymbol{\epsilon}_h^i \right)^\top \\ &= \hat{\Lambda}_{k,h}^{-1} \sum_{i=1}^{k-1} \hat{\sigma}_{i,h}^{-2} \boldsymbol{\phi}(s_h^i, a_h^i) \left(\boldsymbol{\phi}(s_h^i, a_h^i)^\top \boldsymbol{\mu}_h + \boldsymbol{\epsilon}_h^i \right)^\top \\ &= \hat{\Lambda}_{k,h}^{-1} \sum_{i=1}^{k-1} \hat{\sigma}_{i,h}^{-2} \boldsymbol{\phi}(s_h^i, a_h^i) \boldsymbol{\phi}(s_h^i, a_h^i)^\top \boldsymbol{\mu}_h + \hat{\Lambda}_{k,h}^{-1} \sum_{i=1}^{k-1} \hat{\sigma}_{i,h}^{-2} \boldsymbol{\phi}(s_h^i, a_h^i) \boldsymbol{\epsilon}_h^i \top \\ &= \hat{\Lambda}_{k,h}^{-1} \left(\hat{\Lambda}_{k,h} - \lambda \mathbf{I} \right) \boldsymbol{\mu}_h + \hat{\Lambda}_{k,h}^{-1} \sum_{i=1}^{k-1} \hat{\sigma}_{i,h}^{-2} \boldsymbol{\phi}(s_h^i, a_h^i) \boldsymbol{\epsilon}_h^i \top \\ &= \boldsymbol{\mu}_h - \lambda \hat{\Lambda}_{k,h}^{-1} \boldsymbol{\mu}_h + \hat{\Lambda}_{k,h}^{-1} \sum_{i=1}^{k-1} \hat{\sigma}_{i,h}^{-2} \boldsymbol{\phi}(s_h^i, a_h^i) \boldsymbol{\epsilon}_h^i \top. \end{aligned}$$

Rearranging terms gives Eq. (62). □

F.4. Covering Net

This subsection generalizes the covering number argument built for the covering net of the constructed optimistic value function class $\hat{\mathcal{V}}$ in (Jin et al., 2020), to that for the squared optimistic value function class $\hat{\mathcal{V}}^2$ and the pessimistic value function class $\check{\mathcal{V}}$. Before that, a basic fact about the covering number of the Euclidean ball is given in Lemma F.8.

Lemma F.8 (Covering Number of Euclidean Ball, Lemma D.5. in (Jin et al., 2020)). *For any $\varepsilon > 0$, the ε -covering number of the Euclidean ball in \mathbb{R}^d with radius $R > 0$ is upper bounded by $(1 + 2R/\varepsilon)^d$.*

Lemma F.9 (Lemma D.6. in (Jin et al., 2020)). *Let $\hat{\mathcal{N}}_\varepsilon$ be the ε -covering of $\hat{\mathcal{V}}$ with respect to the distance $\text{dist}(V, V') = \sup_x |V(x) - V'(x)|$, where $\hat{\mathcal{V}}$ is defined in Definition C.5. Then*

$$\log |\hat{\mathcal{N}}_\varepsilon| \leq d \log(1 + 4L/\varepsilon) + d^2 \log \left[1 + 8d^{1/2} B^2 / (\lambda \varepsilon^2) \right].$$

Proof. Denote $\mathbf{A} = \beta^2 \Lambda^{-1}$, then for any $\hat{V}(\cdot) \in \hat{\mathcal{V}}$,

$$\hat{V}(\cdot) = \min \left\{ \max_a \mathbf{w}^\top \boldsymbol{\phi}(\cdot, a) + \sqrt{\boldsymbol{\phi}(\cdot, a)^\top \mathbf{A} \boldsymbol{\phi}(\cdot, a)}, H \right\} \quad (63)$$

for $\|\mathbf{w}\| \leq L$ and $\|\mathbf{A}\| \leq B^2\lambda^{-1}$. For any two functions $\widehat{V}_1, \widehat{V}_2 \in \widehat{\mathcal{V}}$, let them take the form in Eq. (63) with parameters $(\mathbf{w}_1, \mathbf{A}_1)$ and $(\mathbf{w}_2, \mathbf{A}_2)$, respectively. Since both $\min\{\cdot, H\}$ and \max_a are contraction mappings, we have

$$\begin{aligned}
 \text{dist}(\widehat{V}_1, \widehat{V}_2) &\leq \sup_{s,a} \left| \left[\mathbf{w}_1^\top \phi(s,a) + \sqrt{\phi(s,a)^\top \mathbf{A}_2 \phi(s,a)} \right] - \left[\mathbf{w}_2^\top \phi(s,a) + \sqrt{\phi(s,a)^\top \mathbf{A}_2 \phi(s,a)} \right] \right| \\
 &\leq \sup_{\phi: \|\phi\|_2 \leq 1} \left| \left[\mathbf{w}_1^\top \phi + \sqrt{\phi^\top \mathbf{A}_2 \phi} \right] - \left[\mathbf{w}_2^\top \phi + \sqrt{\phi^\top \mathbf{A}_2 \phi} \right] \right| \\
 &\leq \sup_{\phi: \|\phi\|_2 \leq 1} \left| (\mathbf{w}_1 - \mathbf{w}_2)^\top \phi \right| + \sup_{\phi: \|\phi\|_2 \leq 1} \sqrt{\left| \phi^\top (\mathbf{A}_1 - \mathbf{A}_2) \phi \right|} \\
 &= \|\mathbf{w}_1 - \mathbf{w}_2\|_2 + \sqrt{\|\mathbf{A}_1 - \mathbf{A}_2\|_2} \leq \|\mathbf{w}_1 - \mathbf{w}_2\|_2 + \sqrt{\|\mathbf{A}_1 - \mathbf{A}_2\|_F}
 \end{aligned} \tag{64}$$

where the second last inequality follows from the fact that $|\sqrt{x} - \sqrt{y}| \leq \sqrt{|x - y|}$ holds for any $x, y \geq 0$. For matrices, $\|\cdot\|_2$ and $\|\cdot\|_F$ denote the matrix operator norm and Frobenius norm respectively.

Let \mathcal{C}_w be an $\varepsilon/2$ -cover of $\{\mathbf{w} \in \mathbb{R}^d \mid \|\mathbf{w}\|_2 \leq L\}$ with respect to the 2-norm, and \mathcal{C}_A be an $\varepsilon^2/4$ -cover of $\{\mathbf{A} \in \mathbb{R}^{d \times d} \mid \|\mathbf{A}\|_F \leq d^{1/2}B^2\lambda^{-1}\}$ with respect to the Frobenius norm. By Lemma F.8, we have:

$$|\mathcal{C}_w| \leq (1 + 4L/\varepsilon)^d, \quad |\mathcal{C}_A| \leq \left[1 + 8d^{1/2}B^2 / (\lambda\varepsilon^2) \right]^{d^2}$$

By Eq. (64), for any $\widehat{V}_1 \in \widehat{\mathcal{V}}$, there exists $\mathbf{w}_2 \in \mathcal{C}_w$ and $\mathbf{A}_2 \in \mathcal{C}_A$ such that \widehat{V}_2 parametrized by $(\mathbf{w}_2, \mathbf{A}_2)$ satisfies $\text{dist}(\widehat{V}_1, \widehat{V}_2) \leq \varepsilon$. Hence, it holds that $|\mathcal{N}_\varepsilon| \leq |\mathcal{C}_w| \cdot |\mathcal{C}_A|$, which gives

$$\log |\mathcal{N}_\varepsilon| \leq \log |\mathcal{C}_w| + \log |\mathcal{C}_A| \leq d \log(1 + 4L/\varepsilon) + d^2 \log \left[1 + 8d^{1/2}B^2 / (\lambda\varepsilon^2) \right].$$

This concludes the proof. \square

Lemma F.10. Let $\widehat{\mathcal{N}}_\varepsilon^2$ be the ε -covering of $\widehat{\mathcal{V}}^2$ with respect to the distance $\text{dist}(V, V') = \sup_x |V(x) - V'(x)|$, where $\widehat{\mathcal{V}}^2$ is defined in Definition C.6. Then

$$\log |\widehat{\mathcal{N}}_\varepsilon^2| \leq d \log(1 + 8LH/\varepsilon) + d^2 \log \left[1 + 32d^{1/2}B^2H^2 / (\lambda\varepsilon^2) \right].$$

Proof. For any $\widehat{V}_1^2(\cdot), \widehat{V}_2^2(\cdot) \in \widehat{\mathcal{V}}^2$, we have

$$\begin{aligned}
 \text{dist}(\widehat{V}_1^2, \widehat{V}_2^2) &= \sup_x \left| \widehat{V}_1^2(x) - \widehat{V}_2^2(x) \right| = \sup_x \left| \widehat{V}_1(x) + \widehat{V}_2(x) \right| \cdot \left| \widehat{V}_1(x) - \widehat{V}_2(x) \right| \\
 &\leq 2H \sup_x \left| \widehat{V}_1(x) - \widehat{V}_2(x) \right|,
 \end{aligned}$$

where $\widehat{V}_1, \widehat{V}_2 \in \widehat{\mathcal{V}}$ and $\widehat{\mathcal{V}}$ is given in Definition C.5. This implies that the $\varepsilon/(2H)$ -cover of $\widehat{\mathcal{V}}$ is a ε -cover of $\widehat{\mathcal{V}}^2$, which combined with Lemma F.9, gives the conclusion. \square

Lemma F.11. Let $\check{\mathcal{N}}_\varepsilon$ be the ε -covering of $\check{\mathcal{V}}$ with respect to the distance $\text{dist}(V, V') = \sup_x |V(x) - V'(x)|$, where $\check{\mathcal{V}}$ is defined in Definition C.7. Then

$$\log |\check{\mathcal{N}}_\varepsilon| \leq d \log(1 + 4L/\varepsilon) + d^2 \log \left[1 + 8d^{1/2}B^2 / (\lambda\varepsilon^2) \right].$$

Proof. Denote $\mathbf{A} = \beta^2\Lambda^{-1}$, then for any $\check{V}(\cdot) \in \check{\mathcal{V}}$,

$$\check{V}(\cdot) = \max \left\{ \max_a \mathbf{w}^\top \phi(\cdot, a) - \sqrt{\phi(\cdot, a)^\top \mathbf{A} \phi(\cdot, a)}, 0 \right\} \tag{65}$$

for $\|\mathbf{w}\| \leq L$ and $\|\mathbf{A}\| \leq B^2\lambda^{-1}$. The proof is almost the same as that for Lemma F.9, since for any two functions $\check{V}_1, \check{V}_2 \in \check{\mathcal{V}}$, let them take the form in Eq. (65) with parameters $(\mathbf{w}_1, \mathbf{A}_1)$ and $(\mathbf{w}_2, \mathbf{A}_2)$, respectively. Since $|\min\{x, 0\} - \min\{y, 0\}| \leq |x - y|$ for any $x, y \in \mathbb{R}$ and \max_a is a contraction mapping, we have

$$\begin{aligned}
 \text{dist}(\hat{V}_1, \hat{V}_2) &\leq \sup_{s,a} \left\| \left[\mathbf{w}_1^\top \phi(s,a) - \sqrt{\phi(s,a)^\top \mathbf{A}_2 \phi(s,a)} \right] - \left[\mathbf{w}_2^\top \phi(s,a) - \sqrt{\phi(s,a)^\top \mathbf{A}_2 \phi(s,a)} \right] \right\| \\
 &\leq \sup_{\phi: \|\phi\|_2 \leq 1} \left\| \left[\mathbf{w}_1^\top \phi - \sqrt{\phi^\top \mathbf{A}_2 \phi} \right] - \left[\mathbf{w}_2^\top \phi - \sqrt{\phi^\top \mathbf{A}_2 \phi} \right] \right\| \\
 &\leq \sup_{\phi: \|\phi\|_2 \leq 1} \left| (\mathbf{w}_1 - \mathbf{w}_2)^\top \phi \right| + \sup_{\phi: \|\phi\|_2 \leq 1} \sqrt{\left| \phi^\top (\mathbf{A}_2 - \mathbf{A}_1) \phi \right|} \\
 &= \|\mathbf{w}_1 - \mathbf{w}_2\|_2 + \sqrt{\|\mathbf{A}_1 - \mathbf{A}_2\|_2} \leq \|\mathbf{w}_1 - \mathbf{w}_2\|_2 + \sqrt{\|\mathbf{A}_1 - \mathbf{A}_2\|_F},
 \end{aligned} \tag{66}$$

which is the same as that of Eq. (64) in Lemma F.9. This implies that the ε -covering of \check{V} can be given similarly to that of \hat{V} , such that the conclusion is given the same as that of Lemma F.9. \square