

HORIZON-FREE REGRET FOR LINEAR MARKOV DECISION PROCESSES

Zihan Zhang*

Jason D. Lee*

Yuxin Chen†

Simon S. Du‡

ABSTRACT

A recent line of works showed regret bounds in reinforcement learning (RL) can be (nearly) independent of planning horizon, a.k.a. the horizon-free bounds. However, these regret bounds only apply to settings where a polynomial dependency on the size of transition model is allowed, such as tabular Markov Decision Process (MDP) and linear mixture MDP. We give the first horizon-free bound for the popular linear MDP setting where the size of the transition model can be exponentially large or even uncountable. In contrast to prior works which explicitly estimate the transition model and compute the inhomogeneous value functions at different time steps, we directly estimate the value functions and confidence sets. We obtain the horizon-free bound by: (1) maintaining *multiple* weighted least square estimators for the value functions; and (2) a structural lemma which shows the maximal total variation of the *inhomogeneous* value functions is bounded by a polynomial factor of the feature dimension.

1 INTRODUCTION

In reinforcement learning (RL), an agent learns to interact with an unknown environment by observing the current states and taking actions sequentially. The goal of the agent is to maximize the accumulative reward. In RL, the sample complexity describes the number of samples needed to learn a near-optimal policy.

It has been shown that the sample complexity needs to scale with the state-action space for the tabular Markov Decision Process (MDP) (Domingues et al., 2021). When the state-action space is large, function approximation is often used to generalize across states and actions. One popular model is linear MDP where the transition model is assumed to be low-rank (Yang & Wang, 2019; Jin et al., 2020b). We denote by d the rank (or feature dimension) and the sample complexity will depend on d instead of the size of the state-action space.

However, there remains a gap between our theoretical understanding of tabular MDPs and linear MDPs. For tabular MDPs, a line of works give horizon-free bounds, i.e., the sample complexity can be (nearly) independent of the planning horizon (Wang et al., 2020; Zhang et al., 2021a). More recently, the horizon-free bounds were also obtained in linear mixture MDP where the underlying MDP could be presented by a linear combination of d known MDPs (Ayoub et al., 2020; Modi et al., 2020; Jia et al., 2020). A natural theoretical question is:

Can we obtain horizon-free regret bounds for linear MDPs?

One major technical challenge in obtaining horizon-free bounds is that the value function is inhomogeneous, i.e., for an MDP with planning horizon H , the optimal value functions for each time step $\{V_h^*\}_{h=1}^H$ can vary across $h = 1, \dots, H$. For tabular MDPs and linear mixture MDPs, one can resolve this challenge by first estimating the transition kernel and then computing the value function based on the learned model. The sample complexity will then scale with the size of the transition

*Department of Electrical and Computer Engineering, Princeton University; email: {zz5478, jasonlee}@princeton.edu.

†Department of Statistics and Data Science, University of Pennsylvania; email: yuxinc@wharton.upenn.edu.

‡Paul G. Allen School of Computer Science and Engineering, University of Washington; email: ssdu@cs.washington.edu.

kernel, which is homogeneous and does not depend on H . For tabular MDPs and linear mixture MDPs, the dependence on the model size is allowed.¹ Unfortunately, this approach cannot be readily applied to handle linear MDPs because the model size of a linear MDP scales with the size of the state space, which can be exponentially large or even uncountable.

Contributions. In this paper, we answer the above question affirmatively by establishing a regret bound of $\tilde{O}(\text{poly}(d)\sqrt{K})$ for linear MDPs. Formally we have the result below.

Theorem 1. *Choose Reward – Confidence as VOFUL (see Algorithm 3). For any MDP satisfying the total-bounded reward assumption (Assumption 1) and linear MDP assumption (see Assumption 2), then with probability $1 - \delta$, the regret of Algorithm 1 is bounded by $\tilde{O}(d^{5.5}\sqrt{K} + d^{6.5})$, where logarithmic factors of $(d, K, H, 1/\delta)$ are hidden by the $\tilde{O}(\cdot)$ parameter.*

By virtue of Theorem 1, we show that linear MDP has a sample-complexity with only poly-logarithmic dependence on H . Although the proposed algorithm (Algorithm 1) is inefficient in computation, we believe our method provides intuitions to remove the horizon dependence in the view of statistical efficiency.

In terms of technical innovations, we design a novel method to share the samples to solve different linear bandit problems. Since the optimal value functions are not homogeneous, we need to learn H different linear bandit problems using the same dataset. We first show that it suffices to bound the regret for each single bandit problem, and then bound the maximal total variance by bounding the variation of the optimal value function. See Section 3 for more details. Due to space limitation, we postpone the full proof of Theorem 1 to Appendix B.

1.1 RELATED WORKS

Tabular MDPs. There has been a long list of algorithms proposed for episodic tabular MDPs (e.g., Kearns & Singh (2002); Brafman & Tennenholtz (2003); Kakade (2003); Agrawal & Jia (2017); Azar et al. (2017); Jin et al. (2018); Zhang et al. (2020); Wang et al. (2020); Jin et al. (2020a); Zhang et al. (2021a); Li et al. (2021b;c; 2023); Zhang et al. (2022)). For finite-horizon *inhomogeneous* MDPs with the immediate reward at each step bounded by $1/H$, Azar et al. (2017); Zhang et al. (2020); Zanette & Brunskill (2019); Li et al. (2021b) achieved asymptotically minimax-optimal regret $\tilde{\Theta}(\sqrt{SAHK})$ (ignoring lower order terms), where SA is the size of state-action space, H is the planning horizon, and K is the number of episodes. Motivated by a conjecture raised by Jiang & Agarwal (2018), Wang et al. (2020) developed—for *time-homogeneous* MDPs with total rewards in any episode bounded above by 1—the first sample complexity upper bound that exhibits only logarithmic dependence on the horizon H , which was later on improved by Zhang et al. (2021a) to yield a near-optimal regret bound of $\tilde{O}(\sqrt{SAK} + S^2A)$. Subsequently, Li et al. (2021c); Zhang et al. (2022) proved that even the poly-logarithmic horizon dependency in the sample complexity can be removed, albeit at the price of suboptimal scaling with SA .

Contextual linear bandits. The linear bandit problem has been extensively studied in past decades (Auer, 2002; Dani et al., 2008; Chu et al., 2011; Abbasi-Yadkori et al., 2011). For linear bandits with infinite arms, the minimax regret bound of $\tilde{\Theta}(\sqrt{dK})$ is achieved by OFUL (Abbasi-Yadkori et al., 2011), where d is the feature dimension and K the number of rounds. With regards to variance-aware algorithms, Zhang et al. (2021b); Kim et al. (2022) proposed VOFUL (VOFUL+) to achieve a variance-dependent regret bound of $\tilde{O}(\text{poly}(d)\sqrt{\sum_{k=1}^K \sigma_k^2} + \text{poly}(d))$, in the absence of the knowledge of $\{\sigma_k\}_{k=1}^K$ (with σ_k the conditional variance of the noise in the k -th round). By assuming prior knowledge of $\{\sigma_k\}_{k=1}^K$, Zhou et al. (2021); Zhou & Gu (2022) obtained similar regret bounds with improved dependence on d . Another work (Faury et al., 2020) proposed a Bernstein-style confidence set for the logistic bandit problem, also assuming availability of the noise variances.

RL with linear function approximation. It has been an important problem in the RL community to determine the generalization capability of linear function approximation (Jiang et al., 2017; Dann et al., 2018; Yang & Wang, 2019; Jin et al., 2020b; Wang et al., 2019; Sun et al., 2019; Zanette

¹For linear mixture MDP, the model size scales linearly with the feature dimension.

et al., 2020; Weisz et al., 2020; Li et al., 2021a; Ayoub et al., 2020; Zhang et al., 2021b; Kim et al., 2022; Zhou et al., 2021; Zhou & Gu, 2022; He et al., 2022). Several model assumptions have been proposed and exploited to capture the underlying dynamics via linear functions. For example, Jiang et al. (2017) investigated low Bellman-rank, which described the algebraic dimension between the decision process and value-function approximator. Another setting proposed and studied by Jia et al. (2020); Ayoub et al. (2020); Modi et al. (2020) is that of linear mixture MDPs, which postulates that the underlying dynamics is a linear combination of d known environments. Focusing on linear mixture MDPs, Zhang et al. (2021b) proposed the first sample-efficient algorithm to achieve horizon-free $\tilde{O}(\text{poly}(d)\sqrt{K})$ regret, and later on Kim et al. (2022) obtained better d -dependency in the regret bound; further, a recent study Zhou & Gu (2022) designed a variance- & uncertainty-aware exploration bonus with weighted least-square regression, achieving near-optimal regret bounds with computation efficiency. Another recent strand of research Yang & Wang (2019); Jin et al. (2020b); He et al. (2022); Agarwal et al. (2022) studied the setting of linear MDPs, where the transition kernel and reward function are assumed to be linear functions of several known low-dimensional feature vectors. Take episodic inhomogeneous linear MDPs for example: when the feature dimension is d and the immediate reward in each step is bounded above by $1/H$, (Jin et al., 2020b) established the regret bound of $\tilde{O}(\sqrt{d^3 H^2 K})$, whereas the follow-up works He et al. (2022); Agarwal et al. (2022) improved the regret to $\tilde{O}(d\sqrt{HK})$. It remained unclear whether and when horizon-free solutions are plausible in linear MDPs, in the hope of accommodating scenarios with exceedingly large H .

2 PRELIMINARIES

In this section, we present the basics of MDPs and the learning process, and introduce our key assumptions. Throughout the paper, $\Delta(X)$ denotes the set of probability distributions over the set X .

Episodic MDPs. A finite-horizon episodic MDP can be represented by a tuple $(\mathcal{S}, \mathcal{A}, R, P, K, H)$, where \mathcal{S} denotes the state space containing S states, \mathcal{A} is the action space containing A different actions, $R : \mathcal{S} \times \mathcal{A} \rightarrow \Delta([0, 1])$ indicates the reward distribution, $P : \mathcal{S} \times \mathcal{A} \rightarrow \Delta(\mathcal{S})$ represents the probability transition kernel, K stands for the total number of sample episodes that can be collected, and H is the planning horizon. In particular, P is assumed throughout to be *time-homogeneous*, which is necessary to enable nearly horizon-free regret bounds; in light of this assumption, we denote by $P_{s,a} := P(\cdot | s, a) \in \Delta(\mathcal{S})$ the transition probability from state s to state s' while taking action a . The reward distribution R is also assumed to be time-homogeneous, so that the immediate reward at a state-action pair (s, a) at any step h is drawn from $R(s, a)$ with mean $\mathbb{E}_{r' \sim R(s,a)}[r'] = r(s, a)$. Moreover, a deterministic and possibly non-stationary policy $\pi = \{\pi_h : \mathcal{S} \rightarrow \mathcal{A}\}_{h=1}^H$ describes an action selection strategy, with $\pi_h(s)$ specifying the action chosen in state s at step h .

At each sample episode, the learner starts from an initial state s_1 ; for each step $h = 1, \dots, H$, the learner observes the current state s_h , takes action a_h accordingly, receives an immediate reward $r_h \sim R(s_h, a_h)$, and then the environment transits to the next state s_{h+1} in accordance with $P(\cdot | s_h, a_h)$. When the actions are selected based on policy π , we can define the Q -function and the value function at step h respectively as follows:

$$Q_h^\pi(s, a) := \mathbb{E}_\pi \left[\sum_{h'=h}^H r_{h'} \mid (s_h, a_h) = (s, a) \right] \quad \text{and} \quad V_h^\pi(s) := \mathbb{E}_\pi \left[\sum_{h'=h}^H r_{h'} \mid s_h = s \right]$$

for any $(s, a) \in \mathcal{S} \times \mathcal{A}$, where $\mathbb{E}_\pi[\cdot]$ denotes the expectation following π , i.e., we execute $a_{h'} = \pi_{h'}(s_{h'})$ for all $h < h' \leq H$ (resp. $h \leq h' \leq H$) in the definition of Q_h^π (resp. V_h^π). The optimal Q -function and value function at step h can then be defined respectively as

$$Q_h^*(s, a) = \max_{\pi} Q_h^\pi(s, a) \quad \text{and} \quad V_h^*(s) = \max_{\pi} V_h^\pi(s), \quad \forall (s, a) \in \mathcal{S} \times \mathcal{A}.$$

These functions satisfy the Bellman optimality condition in the sense that $V_h^*(s) = \max_a Q_h^*(s, a)$, $\forall s \in \mathcal{S}$, and $Q_h^*(s, a) = r(s, a) + \mathbb{E}_{s' \sim P(\cdot | s, a)}[V_{h+1}^*(s')]$, $\forall (s, a) \in \mathcal{S} \times \mathcal{A}$.

The learning process. The learning process entails collection of K sample episodes. At each episode $k = 1, 2, \dots, K$, a policy π^k is selected carefully based on the samples collected in the previous $k - 1$ episodes; the learner then starts from a given initial state s_1^k and executes π^k to collect

the k -th episode $\{(s_h^k, a_h^k, r_h^k)\}_{1 \leq h \leq H}$, where s_h^k, a_h^k and r_h^k denote respectively the state, action and immediate reward at step h of this episode. The learning performance is measured by the total regret

$$\text{Regret}(K) := \sum_{k=1}^K \left(V_1^*(s_1^k) - V_1^{\pi^k}(s_1^k) \right), \quad (1)$$

and our ultimate goal is to design a learning algorithm that minimizes the above regret (1).

Key assumptions. We now introduce two key assumptions imposed throughout this paper, which play a crucial role in determining the minimal regret. The first assumption is imposed upon the rewards, requiring the aggregate reward in any episode to be bounded above by 1 almost surely.

Assumption 1 (Bounded total rewards). *In any episode, we assume that $\sum_{h=1}^H r_h \leq 1$ holds almost surely regardless of the policy in use.*

Compared to the common assumption where the immediate reward at each step is bounded by $1/H$, Assumption 1 is much weaker in that it allows the rewards to be spiky (e.g., we allow the immediate reward at one step to be on the order of 1 with the remaining ones being small). The interested reader is referred to Jiang & Agarwal (2018) for more discussions about the above reward assumption.

The second assumption postulates that the transition kernel and the reward function reside within some known low-dimensional subspace, a scenario commonly referred to as a linear MDP.

Assumption 2 (Linear MDP (Jin et al., 2020b)). *Let \mathcal{B} represent the unit ℓ_2 ball in \mathbb{R}^d , and let $\{\phi(s, a)\}_{(s,a) \in \mathcal{S} \times \mathcal{A}} \subset \mathcal{B}$ be a set of known feature vectors such that $\max_{s,a} \|\phi_{s,a}\|_2 \leq 1$. Assume that there exist a reward parameter $\theta_r \in \mathbb{R}^d$ and a transition kernel parameter $\mu \in \mathbb{R}^{\mathcal{S} \times \mathcal{A}}$ such that*

$$r(s, a) = \langle \phi(s, a), \theta_r \rangle \quad \forall (s, a) \in \mathcal{S} \times \mathcal{A} \quad (2a)$$

$$P(\cdot | s, a) = \mu \phi(s, a), \quad \forall (s, a) \in \mathcal{S} \times \mathcal{A} \quad (2b)$$

$$\|\theta_r\|_2 \leq \sqrt{d}, \quad (2c)$$

$$\|\mu^\top v\|_2 \leq \sqrt{d}, \quad \forall v \in \mathbb{R}^{\mathcal{S}} \text{ obeying } \|v\|_\infty \leq 1. \quad (2d)$$

In words, Assumption 2 requires both the reward function and the transition kernel to be linear combinations of a set of d -dimensional feature vectors, which enables effective dimension reduction as long as $d \ll SA$.

In comparison, another line of works Jia et al. (2020); Ayoub et al. (2020); Modi et al. (2020) focus on the setting of linear mixture MDP below.

Assumption 3 (Linear Mixture MDP). *Let $\{(r_i, P_i)\}_{i=1}^d$ be a group of known reward-transition pairs. Assume that there exists a kernel parameter $\theta \in \mathbb{R}^d$ such that*

$$r(s, a) = \sum_{i=1}^d \theta_i r_i(s, a) \quad \forall (s, a) \in \mathcal{S} \times \mathcal{A} \quad (3a)$$

$$P(\cdot | s, a) = \sum_{i=1}^d \theta_i P_i(\cdot | s, a), \quad \forall (s, a) \in \mathcal{S} \times \mathcal{A} \quad (3b)$$

$$\|\theta\|_1 \leq 1. \quad (3c)$$

Roughly speaking, Assumption 3 requires that the underlying reward-transition pair is a linear combination of d known reward-transition pairs. Recent work Zhou & Gu (2022) achieved a near-tight horizon-regret bound in this setting with a computational efficient algorithm. However, we emphasize that learning a linear MDP is fundamentally harder than learning a linear mixture MDP. The reason is that the only unknown parameter in a linear mixture MDP problem is the hidden kernel θ , which has at most d dimensions. So it is possible to learn θ to fully express the transition model. While in linear MDP, the dimension of unknown parameter μ scales linearly in the number of states, where it is impossible to recover the transition model. To address this problem, previous works on linear MDP try to learn the transition kernel in some certain direction, e.g., $\mu^\top v$ for some certain $v \in \mathbb{R}^{\mathcal{S}}$. This approach faces a fundamental problem in sharing samples among difference layers. We refer to Section 3 for more discussion.

Algorithm 1 Main Algorithm

-
- 1: **Input:** Number of episodes K , horizon H , feature dimension d , confidence parameter δ
 - 2: **Initialization:** $\lambda \leftarrow 1/H^2$, $\epsilon \leftarrow 1/(KH)^4$, $\alpha \leftarrow 150d\sqrt{\log^2((KH)/\delta)}$
 - 3: **for** $k = 1, 2, \dots, K$ **do**
 - 4: $\mathcal{D}^k \leftarrow \{s_{h'}^{k'}, a_{h'}^{k'}, s_{h'+1}^{k'}\}_{h' \in [H], k' \in [k-1]}$;
 - 5: *// Construct the confidence region for the transition kernel.*
 - 6: **for** $v \in \mathcal{W}_\epsilon$ **do**
 - 7: $(\hat{\theta}^k(v), \hat{\theta}^k(v), \Lambda^k(v)) \leftarrow \text{HF - Estimator}(\mathcal{D}^k, v)$;
 - 8: $b^k(v, \phi) \leftarrow \alpha\sqrt{\phi^\top (\Lambda^k(v))^{-1} \phi} + 4\epsilon$;
 - 9: **end for**
 - 10: $\mathcal{U}^k \leftarrow \left\{ \tilde{\mu} \in \mathcal{U} \mid \|\phi^\top \tilde{\mu}^\top v - \phi^\top \hat{\theta}(v)\| \leq b^k(v, \phi), \forall \phi \in \Phi, v \in \mathcal{W}_\epsilon \right\}$
 - 11: *// Construct the confidence region for the reward function.*
 - 12: $\Theta^k \leftarrow \text{Reward - Confidence} \left(\{\phi_h^{k'}/\sqrt{d}\}_{k' \in [k-1], h \in [H]}, \{r_h^{k'}/\sqrt{d}\}_{k' \in [k-1], h \in [H]} \right)$
 - 13: *// Optimistic planning.*
 - 14: $(\mu^k, \theta^k) \leftarrow \arg \max_{\tilde{\mu} \in \mathcal{U}^k, \theta \in \Theta^k} \max_{\pi} \mathbb{E}_{\pi} [\sum_{h=1}^H r_h \mid \tilde{\mu}, \theta]$;
 - 15: π^k be the optimal policy w.r.t. the reward parameter as θ^k and transition parameter as μ^k ;
 - 16: Play π^k in the k -th episode;
 - 17: **end for**
-

Algorithm 2 HF – Estimator

-
- 1: **Input :** A group of samples $\mathcal{D} := \{s_i, a_i, s'_i\}_{i=1}^n$, value function $v \in \mathbb{R}^S$;
 - 2: **Initialization:** $\lambda \leftarrow 1/H^2$, $\alpha \leftarrow 150d\sqrt{\log^2((KH)/\delta)}$, $\phi_i \leftarrow \phi(s_i, a_i)$, $1 \leq i \leq n$, $\epsilon \leftarrow 1/(KH)^4$;
 - 3: $\sigma_1^2 \leftarrow 4$;
 - 4: **for** $i = 2, 3, \dots, n+1$ **do**
 - 5: $\Lambda_{i-1} \leftarrow \lambda \mathbf{I} + \sum_{i'=1}^{i-1} \phi_{i'}^\top \phi_{i'} / \sigma_{i'}^2$;
 - 6: $\tilde{b}_{i-1} \leftarrow \sum_{i'=1}^{i-1} \frac{v^2(s'_{i'})}{\sigma_{i'}^2} \phi_{i'}$, $\tilde{\theta}_{i-1} \leftarrow \Lambda_{i-1}^{-1} \tilde{b}_{i-1}$;
 - 7: $b_{i-1} \leftarrow \sum_{i'=1}^{i-1} \frac{v(s'_{i'})}{\sigma_{i'}^2} \phi_{i'}$, $\theta_{i-1} \leftarrow \Lambda_{i-1}^{-1} b_{i-1}$;
 - 8: $\sigma_i^2 \leftarrow \phi_i^\top \tilde{\theta}_{i-1} - (\phi_i^\top \theta_{i-1})^2 + 16\alpha\sqrt{\phi_i^\top (\Lambda_{i-1})^{-1} \phi_i} + 4\epsilon$;
 - 9: **end for**
 - 10: $\theta \leftarrow \Lambda_n^{-1} b_n$, $\tilde{\theta} \leftarrow \Lambda_n^{-1} \tilde{b}_n$, $\Lambda \leftarrow \Lambda_n$;
 - 11: **Return:** $(\theta, \tilde{\theta}, \Lambda)$;
-

Notation. Let us introduce several notation to be used throughout. First, we use ι to abbreviate $\log(2/\delta)$. For any $x \in \mathbb{R}^d$ and $\Lambda \in \mathbb{R}^{d \times d}$, we define the weighted norm $\|x\|_\Lambda := \sqrt{x^\top \Lambda x}$. Let $[N]$ denote the set $\{1, 2, \dots, N\}$ for a positive integer N . Define $\mathcal{B}(x) := \{\theta \in \mathbb{R}^d \mid \|\theta\|_2 \leq x\}$ and let $\mathcal{B} := \mathcal{B}(1)$ be the unit ball. For two vectors u, v with the same dimension, we say $u \geq v$ (resp. $u \leq v$) iff u is elementwise no smaller (resp. larger) than v . For a random variable X , we use $\text{Var}(X)$ to denote its variance. For any probability vector $p \in \Delta(\mathcal{S})$ and any $v = [v_i]_{1 \leq i \leq S} \in \mathbb{R}^S$, we denote by $\mathbb{V}(p, v) := p^\top (v^2) - (p^\top v)^2$ the associated variance, where $v^2 := [v_i^2]_{1 \leq i \leq S}$ denotes the entrywise square of v . Let ϕ_h^k abbreviate $\phi(s_h^k, a_h^k)$ for any proper (h, k) . Also, we say $(h', k') \leq (h, k)$ iff $h' + k'H \leq h + kH$. Let \mathcal{F}_h^k denote the σ -algebra generated by $\{s_{h'}^{k'}, a_{h'}^{k'}\}_{(h', k') \leq (h, k)}$. We employ $\mathbb{E}[\cdot \mid \tilde{\mu}, \theta]$ to denote the expectation when the underlying linear MDP is generated by the transition kernel parameter $\tilde{\mu}$ and the reward parameter θ (cf. (2)). Moreover, let Φ denote the set of all possible features. Without loss of generality, we assume Φ is a convex set.

3 TECHNIQUE OVERVIEW

In this section, we first discuss the hardness of horizon-free bounds for linear MDP, and then introduce the high-level ideas of our approach. To simplify presentation, we focus on the regret incurred by learning the unknown transition dynamics.

3.1 TECHNICAL CHALLENGE

Least-square regression in the linear MDP problem. Jin et al. (2020c) proposed the first efficient algorithm (LSVI-UCB) for the linear MDP problem. In this method, for each $h \in [H]$, the learner maintains an estimation of V_{h+1} , and constructs optimistic estimators of $Q_h(s, a) := r(s, a) + P_{s,a}^\top V_{h+1}$. Since the reward r is assumed to be known, it suffices to estimate $P_{s,a}^\top V_{h+1} = (\phi(s, a))^\top \mu^\top V_{h+1}$. By defining $\theta_{h+1} := \mu^\top V_{h+1}$, we can estimate $(\phi(s, a))^\top \theta_{h+1}$ with least-square regression because all state-action pairs share the same kernel θ_{h+1} . This task is basically the same as a linear bandit problem, except for that additional factors are needed due to uniform bound over all possible choices V_{h+1} .

To obtain horizon-free regret bound, a common approach is to design estimators for $P_{s,a}^\top V_{h+1}$ with smaller confidence intervals. In this way, we can choose a smaller bonus to keep the optimism, and the regret is also reduced since the leading term in the regret is the sum of bonuses.

Recent work Zhou & Gu (2022) made progress in this direction by designing a variance-aware estimators for the linear regression problem. Roughly speaking, given a groups of samples $\{\phi_i, v_i\}_{i=1}^n$ such that (i) $v_i = \phi_i^\top \theta + \epsilon_i, \forall i \in [n]$; (ii) $\mathbb{E}[\epsilon_i | \{\phi_j\}_{j=1}^i, \{\epsilon_j\}_{j=1}^{i-1}] = 0$ and $\mathbb{E}[\epsilon_i^2 | \{\phi_j\}_{j=1}^i, \{\epsilon_j\}_{j=1}^{i-1}] = \sigma_i^2, \forall i \in [n]$, with the method in Zhou & Gu (2022), the width of the confidence interval of $\phi^\top \theta$ is roughly

$$\tilde{O} \left(\text{poly}(d) \sqrt{\phi^\top \Lambda^{-1} \phi} \right), \quad (4)$$

where $\Lambda = \lambda \mathbf{I} + \sum_{i=1}^n \frac{\phi_i \phi_i^\top}{\sigma_i^2}$ and λ is some proper regularization parameter (See Lemma 1 in Appendix A).

Main technical challenge: Variance-aware estimators coupled with inhomogeneous value functions. While the transition kernel is assumed to be time-homogeneous, the value function and the policy can be *time-inhomogeneous* across different steps. Although the confidence width in (4) seems nice, it poses additional difficulty to bound the sum of bonuses due to *time-inhomogeneous* value functions.

Below we give more technical details to elucidate this technical challenge. To simplify the problem, we assume that the learner is informed of both the reward function and the optimal value function $\{V_h^*\}_{h \in [H]}$. Note that the arguments below can be extended to accommodate unknown $\{V_h^*\}_{h \in [H]}$ as well by means of proper exploitation of the linear MDP structure and a discretization method (i.e., applying a union bound over all possible optimal value functions over a suitably discretized set).

Let $\theta_h^* = \mu^\top V_{h+1}^*$. Then it reduces to learning H contextual bandit problems with hidden parameter as $\{\theta_h^*\}_{h=1}^H$. To remove the polynomial factors of H , it is natural to share samples over different layers. That is, we need to use all the samples along the trajectory $\{s_{h'}, a_{h'}, s_{h'+1}\}_{h'=1}^H$ to estimate the value of $\phi^\top \theta_h^*$.

To solve the h -th linear bandit problem, following (4), we could get a regret bound of $\text{Regret}_h(K) := \tilde{O} \left(\sum_{k=1}^K \sqrt{(\phi_h^k)^\top (\Lambda^k(V_{h+1}^*))^{-1} \phi_h^k} \right)$. Here $\Lambda^k(v) = \lambda \mathbf{I} + \sum_{k'=1}^{k-1} \sum_{h'=1}^H \frac{\phi_{h'}^{k'} (\phi_{h'}^{k'})^\top}{(\sigma_{h'}^{k'}(v))^2}$ with $(\sigma_{h'}^{k'}(v))^2$ as an upper bound for the variance $\mathbb{V}(P_{s_{h'}, a_{h'}}^{k'}, v)$ for $v \in \mathbb{R}^S$. Taking sum over h , the resulting regret bound is roughly

$$\sum_{k=1}^K \min \left\{ \sum_{h=1}^H \sqrt{(\phi_h^k)^\top (\Lambda^k(V_{h+1}^*))^{-1} \phi_h^k}, 1 \right\}. \quad (5)$$

We remark that if V_h^* is homogeneous in h , i.e., there exists V^* such that $V_h^* = V^*$ for any $h \in [H]$, we could use Cauchy's inequality to bound (5) by²

$$\sqrt{\sum_{k=1}^K \min \left\{ \sum_{h=1}^H (\phi_h^k / (\sigma_h^k(V^*)))^\top (\Lambda^k(V^*))^{-1} (\phi_h^k / \sigma_h^k(V^*)), 1 \right\}} \cdot \sqrt{\sum_{k=1}^K \sum_{h=1}^H (\sigma_h^k(V^*))^2}. \quad (6)$$

Noting that

$$\Lambda^{k+1}(V^*) = \Lambda^k(V^*) + \sum_{h'=1}^H \frac{\phi_{h'}^k (\phi_{h'}^k)^\top}{(\sigma_{h'}^k(V^*))^2}, \quad (7)$$

we can further use the elliptical potential lemma (Lemma 6) to bound the first term in (6), and the total variance lemma for MDPs to bound the second term in (6). As a result, we can easily bound (5) by $\tilde{O}(\text{poly}(d)\sqrt{K})$.

However, the arguments above cannot work when V_h^* depends on h . In such cases, the first term in (5) would be

$$\sqrt{\sum_{k=1}^K \min \left\{ \sum_{h=1}^H (\phi_h^k / \sigma_h^k(V_h^*))^\top (\Lambda^k(V_{h+1}^*))^{-1} (\phi_h^k / \sigma_h^k(V_h^*)), 1 \right\}}. \quad (8)$$

To invoke elliptical potential lemma, we need $\Lambda^{k+1}(V_{h+1}^*) - \Lambda^k(V_{h+1}^*) = \sum_{h'=1}^H \frac{\phi_{h'}^k (\phi_{h'}^k)^\top}{(\sigma_{h'}^k(V_{h+1}^*))^2}$, which does not hold since $\sigma_{h'}^k(V_{h+1}^*) \neq \sigma_{h'}^k(V_h^*)$.

In comparison, for tabular MDP, the variance aware bonus has a simple form of $\sqrt{\frac{\mathbb{V}(P_{s,a}, V_{h+1}^*)}{N}}$, so that one can invoke Cauchy's inequality to bound the sum of bonuses; for linear mixture MDP, because there is only one kernel parameter θ and one information matrix, it suffices to analyze like (6) and (7).

3.2 OUR METHODS

In high-level idea, by noting that the main obstacle is the *time-inhomogeneous* value function, we aim to prove that the value function $\{V_h^*\}_{h=1}^H$ could be divided into several groups such that in each group, the value functions are similar to each other measured by the variance.

Technique 1: a uniform upper bound for the variances. We consider using a uniform upper bound $(\bar{\sigma}_{h'}^{k'})^2 := \max_{h \in [H]} (\sigma_{h'}^{k'}(V_{h+1}^*))^2$ to replace $(\sigma_{h'}^{k'}(V_{h+1}^*))^2$ when computing $\Lambda^k(V_{h+1}^*)$. That is, by setting $\bar{\Lambda}^k = \lambda \mathbf{I} + \sum_{k'=1}^{k-1} \sum_{h'=1}^H \frac{\phi_{h'}^{k'} (\phi_{h'}^{k'})^\top}{(\bar{\sigma}_{h'}^{k'})^2} \preceq \Lambda^k(V_{h+1}^*)$ for any $h \in [H]$, we can bound (5) as below:

$$\begin{aligned} & \sum_{k=1}^K \min \left\{ \sum_{h=1}^H \sqrt{(\phi_h^k)^\top (\Lambda^k(V_{h+1}^*))^{-1} \phi_h^k}, 1 \right\} \\ & \leq \sum_{k=1}^K \min \left\{ \sum_{h=1}^H \sqrt{(\phi_h^k)^\top (\bar{\Lambda}^k)^{-1} \phi_h^k}, 1 \right\} \\ & \approx \sqrt{\sum_{k=1}^K \min \left\{ \sum_{h=1}^H (\phi_h^k / \bar{\sigma}_h^k)^\top (\bar{\Lambda}^k)^{-1} (\phi_h^k / \bar{\sigma}_h^k), 1 \right\}} \cdot \sqrt{\sum_{k=1}^K \sum_{h=1}^H (\bar{\sigma}_h^k)^2}. \end{aligned} \quad (9)$$

With the elliptical potential lemma (Lemma 6), we have that

$$\sum_{k=1}^K \min \left\{ \sum_{h=1}^H (\phi_h^k / \bar{\sigma}_h^k)^\top (\bar{\Lambda}^k)^{-1} (\phi_h^k / \bar{\sigma}_h^k), 1 \right\} = \tilde{O}(\sqrt{d}).$$

²Here we omit a lower order term.

So it suffices to deal with $\sum_{k=1}^K \sum_{h=1}^H (\bar{\sigma}_h^k)^2$. For simplicity, we assume that $(\sigma_h^k(v))^2$ is exactly $\mathbb{V}(P_{s_h^k, a_h^k}, v)$ and consider to bound $\sum_{k=1}^K \sum_{h=1}^H \max_{h' \in [H]} \mathbb{V}(P_{s_h^k, a_h^k}, V_{h'+1}^*)$.

Noting that $\mathbb{V}(P_{s,a}, v)$ could be written as $\phi(s, a)^\top (\theta(v^2)) - (\phi(s, a)^\top \theta(v))^2$, which is a linear function of the matrix $\begin{bmatrix} \phi(s, a) \phi^\top(s, a) & \phi(s, a) \\ \phi^\top(s, a) & 1 \end{bmatrix}$, we can bound $\sum_{k=1}^K \sum_{h=1}^H \max_{h' \in [H]} \mathbb{V}(P_{s_h^k, a_h^k}, V_{h'+1}^*)$ by $2(d+1)^2 \max_{h' \in [H]} \sum_{k=1}^K \sum_{h=1}^H \mathbb{V}(P_{s_h^k, a_h^k}, V_{h'+1}^*)$ with a useful technical lemma (See Lemma 5).

As a result, it suffices to bound $\sum_{k=1}^K \sum_{h=1}^H \mathbb{V}(P_{s_h^k, a_h^k}, V_{h'+1}^*)$ for each $h' \in [H]$. However, because $V_{h'+1}^*$ can vary significantly when h' is closed to H , $\sum_{k=1}^K \sum_{h=1}^H \mathbb{V}(P_{s_h^k, a_h^k}, V_{h'+1}^*)$ might be large in the worst case. We consider the toy example below.

Example 1. Fix some $\epsilon > 0$. Let $\mathcal{S} := \{s_1, s_2, s_3, z\}$, $\mathcal{A} = \{a_1, a_2\}$. Let $P_{s_1, a_1} = P_{s_2, a_1} = [\frac{1}{2} - \epsilon, \frac{1}{2} - \epsilon, \epsilon, 0]^\top$, $r(s_1, a_1) = r(s_2, a_1) = 0$, $P_{s_1, a_2} = P_{s_2, a_2} = [0, 0, 0, 1]^\top$, $r(s_1, a_2) = \frac{1}{2}$, $r(s_2, a_2) = 0$, $P_{s_3, a_1} = P_{s_3, a_2} = [0, 0, 0, 1]^\top$, $r(s_3, a_1) = r(s_3, a_2) = 1$, $P_{z, a_1} = P_{z, a_2} = [0, 0, 0, 1]^\top$, and $r(z, a_1) = r(z, a_2) = 0$.

In this toy example, we have two frequent states $\{s_1, s_2\}$, one transient state $\{s_3\}$ with reward 1 and one ending state z with no reward. The transition dynamics at $\{s_1, s_2\}$ is the same, but one can get reward $\frac{1}{2}$ in one step by taking action a_2 at s_1 . Suppose $H \gg \frac{1}{\epsilon}$ and $h \leq \frac{H}{2}$, then the optimal action for $\{s_1, s_2\}$ at the h -th step should be a_1 , and $V_h^*(s_1) \approx V_h^*(s_2) \approx 1$. On the other hand, it is easy to observe $V_H^*(s_1) = \frac{1}{2}$ and $V_H^*(s_2) = 0$. Let the initial state be s_1 . Following the optimal policy, we have $\mathbb{E} \left[\sum_{h=1}^{\frac{H}{2}} \mathbb{V}(P_{s_h^k, a_h^k}, V_H^*) \right] = \Omega(\frac{1}{\epsilon}) \gg 1$ when choosing ϵ small enough.

Technique 2: bounding the total variation. Direct computation shows that for $1 \leq h_1 < h_2 \leq [H]$,

$$\sum_{k=1}^K \sum_{h=h_1}^{h_2} \mathbb{V}(P_{s_h^k, a_h^k}, V_{h'+1}^*) = \tilde{O}(K + K(h_2 - h_1 + 1) \|V_{h'}^* - V_{h'+1}^*\|_\infty). \quad (10)$$

Let $l_h = \|V_h^* - V_{h+1}^*\|_\infty$. It is easy to observe that $l_h \leq l_{h+1}$ for $1 \leq h \leq H-1$ since the Bellman operator Γ is a contraction, i.e., $\|\Gamma(v_1) - \Gamma(v_2)\|_\infty \leq \|v_1 - v_2\|_\infty$ for any $v_1, v_2 \in \mathbb{R}^S$. So we can obtain $l_h \leq \frac{\sum_{h'=1}^{H-1} l_{h'}}{H-h+1}$. For tabular MDP, it is easy to bound $\sum_{h=1}^H l_h \leq S$ since $\|V_h^* - V_{h+1}^*\|_\infty \leq \sum_s (V_h(s) - V_{h+1}(s))$. As a generalization to linear MDP, by Lemma 5 we have that

$$\sum_{h=1}^{H-1} l_h^* \leq \sum_{h=1}^{H-1} \max_{\phi \in \Phi} \phi^\top \mu^\top (V_{h+1} - V_{h+2}) \leq \max_{\phi \in \Phi} 2d \phi^\top \sum_{h=1}^{H-1} \mu^\top (V_{h+1} - V_{h+2}) \leq 2d. \quad (11)$$

As a result, $l_h^* \leq \frac{2d}{H-h+1}$.

Technique 3: doubling segments. By choosing $h_1 = \frac{H}{2} + 1$ and $h_2 = H$ in (10), for $h' \in [h_1, h_2]$,

$$\sum_{k=1}^K \sum_{h=h_1}^{h_2} \mathbb{V}(P_{s_h^k, a_h^k}, V_{h'+1}^*) = \tilde{O}(K + K(h_2 - h_1 + 1) \|V_{h'}^* - V_{h'+1}^*\|_\infty) = \tilde{O}(Kd).$$

This inspires us to divide $[H]$ several segments $[H] = \cup_i \mathcal{H}_i$ with $\mathcal{H}_i = \{h | H - \frac{H}{2^{i-1}} + 1 \leq h \leq H - \frac{H}{2^i}\}$ and $\mathcal{H}_{\log_2(H)+1} = \{H\}$ ³. Consequently, for any i and $h' \in \mathcal{H}_i$, using (10) and the fact that $l_{h'}^* \leq \frac{2d}{H-h'+1} \leq \frac{2^{i+1}d}{H}$, $\sum_{k=1}^K \sum_{h \in \mathcal{H}_i} \mathbb{V}(P_{s_h^k, a_h^k}, V_{h'+1}^*) = \tilde{O}(Kd)$.

Note that we only bound $\sum_{k=1}^K \sum_{i=1}^{\log_2(H)+1} \max_{h' \in \mathcal{H}_i} \sum_{h \in \mathcal{H}_i} \mathbb{V}(P_{s_h^k, a_h^k}, V_{h'+1}^*)$, which does not imply any bound for $\max_{h' \in [H]} \sum_{k=1}^K \sum_{h=1}^H \mathbb{V}(P_{s_h^k, a_h^k}, V_{h'+1}^*)$. Recall that our initial target is to

³We assume $\log_2(H)$ is an integer without loss of generality.

bound $\sum_{k=1}^K \min \left\{ \sum_{h=1}^H \sqrt{(\phi_h^k)^\top (\Lambda^k(V_{h+1}^*))^{-1} \phi_h^k}, 1 \right\}$. A natural idea is to group $h \in \mathcal{H}_i$ for each i to avoid the term $\max_{h' \in [H]} \sum_{k=1}^K \sum_{h=1}^H \mathbb{V}(P_{s_h^k, a_h^k}, V_{h'+1}^*)$. In other words, we turn to bound $\sum_{k=1}^K \min \left\{ \sum_{h \in \mathcal{H}_i} \sqrt{(\phi_h^k)^\top (\Lambda^k(V_{h+1}^*))^{-1} \phi_h^k}, 1 \right\}$ for each i separately. More precisely, for fixed i , we let $(\bar{\sigma}_{h'}^{k'})^2 = \max_{h \in \mathcal{H}_i} (\sigma_{h'}^{k'}(V_{h+1}^*))^2$, and $\bar{\Lambda}^k = \lambda \mathbf{I} + \sum_{k'=1}^{k-1} \sum_{h' \in \mathcal{H}_i} \frac{\phi_{h'}^{k'} (\phi_{h'}^{k'})^\top}{(\bar{\sigma}_{h'}^{k'})^2}$. With the arguments above, we have that

$$\begin{aligned} & \sum_{k=1}^K \min \left\{ \sum_{h \in \mathcal{H}_i} \sqrt{(\phi_h^k)^\top (\Lambda^k(V_{h+1}^*))^{-1} \phi_h^k}, 1 \right\} \\ & \leq \sqrt{\sum_{k=1}^K \min \left\{ \sum_{h \in \mathcal{H}_i} (\phi_h^k / \bar{\sigma}_h^k) (\bar{\Lambda}^k)^{-1} (\phi_h^k / \bar{\sigma}_h^k), 1 \right\}} \cdot \sqrt{\sum_{k=1}^K \sum_{h \in \mathcal{H}_i} (\bar{\sigma}_h^k)^2} + \tilde{O}(d) \\ & = \tilde{O}(\sqrt{Kd^4}). \end{aligned} \tag{12}$$

4 ALGORITHM

In this section, we introduce Algorithm 1. The algorithm is based on model-elimination. At each episode $k = 1, 2, \dots, K$, we maintain \mathcal{U}^k as confidence region of μ and Θ^k as confidence region for θ_r . Then we select the optimistic transition model and reward function from $\mathcal{U}^k \times \Theta^k$ and then execute the corresponding optimal policy. The key step is how to construct \mathcal{U}^k . Inspired by recent work Zhou & Gu (2022), we consider the weighted least square regression to estimate the value function and corresponding variance, which is presented in Algorithm 2. We also borrow VOFUL in Zhang et al. (2021b) to construct the confidence region for θ_r .

Recall that $\mathcal{B}(2\sqrt{d}) = \{\theta \in \mathbb{R}^d \mid \|\theta\|_2 \leq 2\sqrt{d}\}$. For fixed $\epsilon > 0$, there exists an ϵ -net $\mathcal{B}_\epsilon(2\sqrt{d})$ w.r.t. L_∞ for $\mathcal{B}(2\sqrt{d})$ such that $|\mathcal{B}_\epsilon(2\sqrt{d})| \leq O((4\sqrt{d}/\epsilon)^d)$. By Assumption 2, for any $v \in \mathbb{R}^S$ such that $\|v\|_\infty \leq 1$, it holds that $\|\mu^\top v\|_2 \leq \sqrt{d}$. Therefore, for any MDP such that Assumption 1 and 2 holds, its optimal value function is in the set

$$\mathcal{W} := \{v \in \mathbb{R}^S \mid \exists \theta \in \mathcal{B}(2\sqrt{d}), v(s) = \max_a \{\min \{\max_a \phi^\top(s, a)\theta, 1\}, 0\}, \forall s \in \mathcal{S}\}.$$

Define $\mathcal{W}_\epsilon = \left\{ v \in \mathbb{R}^S \mid \exists \theta \in \mathcal{B}_\epsilon(2\sqrt{d}), v(s) = \max_a \{\min \{\max_a \phi^\top(s, a)\theta, 1\}, 0\}, \forall s \in \mathcal{S} \right\}$. For fixed $\theta \in \mathcal{B}(2\sqrt{d})$ and $s \in \mathcal{S}$, the function $\max \left\{ \min \left\{ \max_a \phi(s, a)^\top \theta, 1 \right\}, 0 \right\}$ is $O(1)$ -Lipschitz continuous w.r.t L_∞ norm. As a result, \mathcal{W}_ϵ is an ϵ -net w.r.t. L_∞ norm of \mathcal{W} . Besides, the size of \mathcal{W}_ϵ is bounded by $|\mathcal{W}_\epsilon| = O((4\sqrt{d}/\epsilon)^d)$.

Confidence region for the transition kernel. Fix a group of sequential samples $\{\phi_i\}_{i=1}^n$ and a value function $v \in \mathcal{W}_\epsilon$. Fix $\phi \in \Phi$ and let $\theta(v) = \mu^\top v$. We aim to construct a confidence interval from $\phi^\top \mu^\top v$, and then eliminate all the transition kernels $\tilde{\mu}$ which fails to satisfy the confidence interval for some v and ϕ . To obtain variance-aware confidence interval, we need to compute the variance to feed the weight least-square estimator in Zhou & Gu (2022). For this purpose, for the i -th variance $\mathbb{V}(\mu\phi_i, v)$, we construct σ_i^2 such that $\sigma_i^2 \geq \mathbb{V}(\mu\phi_i, v)$ and the error $\sigma_i^2 - \mathbb{V}(\mu\phi_i, v)$ is well controlled. To compute σ_i^2 , we need to estimate $\phi_i^\top \theta(v^2)$ and $\phi_i^\top \theta(v)$ using the first $i-1$ -th samples, which requires the knowledge of $\mathbb{V}(\mu\phi_{i'}, v^2)$ for $i' \leq i-1$. To address this problem, Zhou & Gu (2022) recursively estimated the 2^m -th order momentum for $m = 1, 2, \dots, \log_2(H)$. In comparison, by the fact that $\mathbb{V}(\mu\phi_{i'}, v^2) \leq 4\mathbb{V}(\mu\phi_{i'}, v)$ (see Lemma 2), we can use $4\sigma_{i'}^2$ as an upper bound for $\mathbb{V}(\mu\phi_{i'}, v^2)$.

Confidence region for the reward parameter. To estimate the reward parameter θ_r , we invoke VOFUL in Zhang et al. (2021b). We remark that the randomness in reward is independent of the randomness in transition dynamics, so that learning the transition dynamic does not help to estimate the variance of reward. More precisely, the variance of $R(s, a)$ could be a non-linear function in $\phi(s, a)$, while the variance of $V_h^*(s')$ with $s' \sim P(\cdot | s, a)$ must be a linear function in $\phi(s, a)$. In Appendix C, we present VOFUL and summarize some useful properties to bound the error due to uncertainty of reward parameter.

ACKNOWLEDGEMENT

Y. Chen is supported in part by the Alfred P. Sloan Research Fellowship, the Google Research Scholar Award, the AFOSR grants FA9550-19-1-0030 and FA9550-22-1-0198, the ONR grant N00014-22-1-2354, and the NSF grants CCF-2221009 and CCF-1907661. JDL acknowledges support of the ARO under MURI Award W911NF-11-1-0304, the Sloan Research Fellowship, NSF CCF 2002272, NSF IIS 2107304, NSF CIF 2212262, ONR Young Investigator Award, and NSF CAREER Award 2144994. SSD acknowledges the support of NSF IIS 2110170, NSFDMS 2134106, NSF CCF 2212261, NSF IIS 2143493, NSF CCF 2019844, NSF IIS 2229881.

REFERENCES

- Yasin Abbasi-Yadkori, Dávid Pál, and Csaba Szepesvári. Improved algorithms for linear stochastic bandits. *Advances in neural information processing systems*, 24, 2011.
- Alekh Agarwal, Yujia Jin, and Tong Zhang. Vo q l: Towards optimal regret in model-free rl with nonlinear function approximation. *arXiv preprint arXiv:2212.06069*, 2022.
- Shipra Agrawal and Randy Jia. Optimistic posterior sampling for reinforcement learning: worst-case regret bounds. In *Advances in Neural Information Processing Systems*, pp. 1184–1194, 2017.
- Peter Auer. Using confidence bounds for exploitation-exploration trade-offs. *Journal of Machine Learning Research*, 3(Nov):397–422, 2002.
- Alex Ayoub, Zeyu Jia, Csaba Szepesvari, Mengdi Wang, and Lin F Yang. Model-based reinforcement learning with value-targeted regression. In *Proceedings of the 37th International Conference on Machine Learning*, 2020.
- Mohammad Gheshlaghi Azar, Ian Osband, and Rémi Munos. Minimax regret bounds for reinforcement learning. In *Proceedings of the 34th International Conference on Machine Learning*, pp. 263–272, 2017.
- Ronen I. Brafman and Moshe Tennenholtz. R-max - a general polynomial time algorithm for near-optimal reinforcement learning. *J. Mach. Learn. Res.*, 3(Oct):213–231, March 2003. ISSN 1532-4435.
- Liyu Chen, Mehdi Jafarnia-Jahromi, Rahul Jain, and Haipeng Luo. Implicit finite-horizon approximation and efficient optimal algorithms for stochastic shortest path. *Advances in Neural Information Processing Systems*, 34, 2021.
- Wei Chu, Lihong Li, Lev Reyzin, and Robert Schapire. Contextual bandits with linear payoff functions. In *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics*, pp. 208–214. JMLR Workshop and Conference Proceedings, 2011.
- Varsha Dani, Thomas P Hayes, and Sham M Kakade. Stochastic linear optimization under bandit feedback. 2008.
- Christoph Dann, Nan Jiang, Akshay Krishnamurthy, Alekh Agarwal, John Langford, and Robert E. Schapire. On oracle-efficient PAC-RL with rich observations. In *Advances in Neural Information Processing Systems*, 2018.
- Omar Darwiche Domingues, Pierre Ménard, Emilie Kaufmann, and Michal Valko. Episodic reinforcement learning in finite mdps: Minimax lower bounds revisited. In *Algorithmic Learning Theory*, pp. 578–598. PMLR, 2021.
- Louis Faury, Marc Abeille, Clément Calauzènes, and Olivier Fercoq. Improved optimistic algorithms for logistic bandits. In *International Conference on Machine Learning*, pp. 3052–3060. PMLR, 2020.
- Jiafan He, Heyang Zhao, Dongruo Zhou, and Quanquan Gu. Nearly minimax optimal reinforcement learning for linear markov decision processes. *arXiv preprint arXiv:2212.06132*, 2022.

- Zeyu Jia, Lin Yang, Csaba Szepesvari, and Mengdi Wang. Model-based reinforcement learning with value-targeted regression. In *Learning for Dynamics and Control*, pp. 666–686. PMLR, 2020.
- Nan Jiang and Alekh Agarwal. Open problem: The dependence of sample complexity lower bounds on planning horizon. In *Conference On Learning Theory*, pp. 3395–3398, 2018.
- Nan Jiang, Akshay Krishnamurthy, Alekh Agarwal, John Langford, and Robert E Schapire. Contextual decision processes with low Bellman rank are PAC-learnable. In *Proceedings of the 34th International Conference on Machine Learning*, pp. 1704–1713, 2017.
- Chi Jin, Zeyuan Allen-Zhu, Sebastien Bubeck, and Michael I Jordan. Is Q-learning provably efficient? In *Advances in Neural Information Processing Systems*, pp. 4863–4873, 2018.
- Chi Jin, Akshay Krishnamurthy, Max Simchowitz, and Tiancheng Yu. Reward-free exploration for reinforcement learning. *International Conference on Machine Learning*, 2020a.
- Chi Jin, Zhuoran Yang, Zhaoran Wang, and Michael I Jordan. Provably efficient reinforcement learning with linear function approximation. In *Conference on Learning Theory*, pp. 2137–2143, 2020b.
- Chi Jin, Zhuoran Yang, Zhaoran Wang, and Michael I Jordan. Provably efficient reinforcement learning with linear function approximation. In *Conference on Learning Theory*, pp. 2137–2143. PMLR, 2020c.
- Sham M Kakade. *On the sample complexity of reinforcement learning*. PhD thesis, University of London London, England, 2003.
- Michael Kearns and Satinder Singh. Near-optimal reinforcement learning in polynomial time. *Machine learning*, 49(2-3):209–232, 2002.
- Yeoneung Kim, Insoon Yang, and Kwang-Sung Jun. Improved regret analysis for variance-adaptive linear bandits and horizon-free linear mixture mdps. *Advances in Neural Information Processing Systems*, 35:1060–1072, 2022.
- Gen Li, Yuxin Chen, Yuejie Chi, Yuantao Gu, and Yuting Wei. Sample-efficient reinforcement learning is feasible for linearly realizable MDPs with limited revisiting. *Advances in Neural Information Processing Systems*, 34:16671–16685, 2021a.
- Gen Li, Laixi Shi, Yuxin Chen, Yuantao Gu, and Yuejie Chi. Breaking the sample complexity barrier to regret-optimal model-free reinforcement learning. *Advances in Neural Information Processing Systems*, 34, 2021b.
- Gen Li, Yuling Yan, Yuxin Chen, and Jianqing Fan. Minimax-optimal reward-agnostic exploration in reinforcement learning. *arXiv preprint arXiv:2304.07278*, 2023.
- Yuanzhi Li, Ruosong Wang, and Lin F Yang. Settling the horizon-dependence of sample complexity in reinforcement learning. In *IEEE Symposium on Foundations of Computer Science*, 2021c.
- Aditya Modi, Nan Jiang, Ambuj Tewari, and Satinder Singh. Sample complexity of reinforcement learning using linearly combined model ensembles. In *International Conference on Artificial Intelligence and Statistics*, pp. 2010–2020. PMLR, 2020.
- Wen Sun, Nan Jiang, Akshay Krishnamurthy, Alekh Agarwal, and John Langford. Model-based RL in contextual decision processes: PAC bounds and exponential improvements over model-free approaches. In *Conference on Learning Theory*, pp. 2898–2933, 2019.
- Ruosong Wang, Simon S Du, Lin F Yang, and Sham M Kakade. Is long horizon reinforcement learning more difficult than short horizon reinforcement learning? In *Advances in Neural Information Processing Systems*, 2020.
- Yining Wang, Ruosong Wang, Simon S Du, and Akshay Krishnamurthy. Optimism in reinforcement learning with generalized linear function approximation. *arXiv preprint arXiv:1912.04136*, 2019.

- Gellert Weisz, Philip Amortila, and Csaba Szepesvári. Exponential lower bounds for planning in mdps with linearly-realizable optimal action-value functions. *arXiv preprint arXiv:2010.01374*, 2020.
- Lin Yang and Mengdi Wang. Sample-optimal parametric Q-learning using linearly additive features. In *International Conference on Machine Learning*, pp. 6995–7004, 2019.
- Andrea Zanette and Emma Brunskill. Tighter problem-dependent regret bounds in reinforcement learning without domain knowledge using value function bounds. In *International Conference on Machine Learning*, pp. 7304–7312, 2019.
- Andrea Zanette, Alessandro Lazaric, Mykel Kochenderfer, and Emma Brunskill. Learning near optimal policies with low inherent Bellman error. In *International Conference on Machine Learning*, 2020.
- Zihan Zhang, Yuan Zhou, and Xiangyang Ji. Almost optimal model-free reinforcement learning via reference-advantage decomposition. In *Advances in Neural Information Processing Systems*, 2020.
- Zihan Zhang, Xiangyang Ji, and Simon Du. Is reinforcement learning more difficult than bandits? a near-optimal algorithm escaping the curse of horizon. In *Conference on Learning Theory*, pp. 4528–4531. PMLR, 2021a.
- Zihan Zhang, Jiaqi Yang, Xiangyang Ji, and Simon S Du. Variance-aware confidence set: Variance-dependent bound for linear bandits and horizon-free bound for linear mixture mdp. In *Advances in Neural Information Processing Systems*, 2021b.
- Zihan Zhang, Xiangyang Ji, and Simon Du. Horizon-free reinforcement learning in polynomial time: the power of stationary policies. In *Conference on Learning Theory*, pp. 3858–3904. PMLR, 2022.
- Dongruo Zhou and Quanquan Gu. Computationally efficient horizon-free reinforcement learning for linear mixture mdps. *arXiv preprint arXiv:2205.11507*, 2022.
- Dongruo Zhou, Quanquan Gu, and Csaba Szepesvari. Nearly minimax optimal reinforcement learning for linear mixture markov decision processes. In *Conference on Learning Theory*, pp. 4532–4576. PMLR, 2021.

A TECHNICAL LEMMAS

Lemma 1 (Theorem 4.3 in Zhou & Gu (2022)). *Let $\{\mathcal{F}_i\}_{i=1}^\infty$ be a filtration, and $\{\psi_i, \zeta_i\}_{i \geq 1}$ be a stochastic process such that $\psi_i \in \mathbb{R}^d$ is \mathcal{F}_i -measurable and $\epsilon_i \in \mathbb{R}$ is \mathcal{F}_{i+1} measurable. Let $L, \sigma, R, \lambda, c > 0, \theta^* \in \mathbb{R}^d$. For $i \geq 1$, let $y_i = \langle \psi_i, \theta^* \rangle + \zeta_i$ and suppose that*

$$\mathbb{E}[\zeta_i | \mathcal{F}_i] = 0, \quad \mathbb{E}[\zeta_i^2 | \mathcal{F}_i] \leq \sigma^2, |\zeta_i| \leq R, \|\psi_i\|_2 \leq L. \quad (13)$$

For $i \geq 1$, let $\Lambda_i = \lambda \mathbf{I} + \sum_{i'=1}^i \psi_{i'} \psi_{i'}^\top$ and $b_i = \sum_{i'=1}^i y_{i'} \psi_{i'}$, $\theta_i = \Lambda_i^{-1} b_i$, and

$$\begin{aligned} \kappa_i &= 12\sqrt{\sigma^2 d \log(1 + iL^2/(d\lambda)) \log(32(\log(R/c) + 1)i^2/\delta)} \\ &\quad + 24 \log(32(\log(R/c) + 1)k^2/\delta) \max_{1 \leq i' \leq i} \{|\zeta_{i'}| \min\{1, \|\psi_{i'}\|_{\Lambda_{i'-1}^{-1}}\}\} + 6 \log(32(\log(R/c) + 1)k^2/\delta)c. \end{aligned}$$

Then, for any $0 < \delta < 1$, with probability $1 - \delta$.

$$\forall i \geq 1, \quad \left\| \sum_{i'=1}^i \psi_{i'} \zeta_{i'} \right\|_{\Lambda_i^{-1}} \leq \kappa_i, \quad \|\theta_i - \theta^*\|_{\Lambda_i} \leq \kappa_i + \sqrt{\lambda} \|\theta^*\|_2.$$

Lemma 2 (Chen et al. (2021)). *Let X be a random variable taking value in $[-C, C]$ for some $C \geq 0$. Let $\text{var}(Y)$ denote the variance of a random variable Y . It then holds that $\text{var}(X^2) \leq 4C^2 \text{var}(X)$.*

Lemma 3. [Lemma 10 in Zhang et al. (2020)] *Let $(M_n)_{n \geq 0}$ be a martingale such that $M_0 = 0$ and $|M_n - M_{n-1}| \leq c$ for some $c > 0$ and any $n \geq 1$. Let $\text{Var}_n = \sum_{k=1}^n \mathbb{E}[(M_k - M_{k-1})^2 | \mathcal{F}_{k-1}]$ for $n \geq 0$, where $\mathcal{F}_k = \sigma(M_1, M_2, \dots, M_k)$. Then for any positive integer n , and any $\epsilon, p > 0$, we have that*

$$\mathbb{P} \left[|M_n| \geq 2\sqrt{\text{Var}_n \log\left(\frac{1}{p}\right)} + 2\sqrt{\epsilon \log\left(\frac{1}{p}\right)} + 2c \log\left(\frac{1}{p}\right) \right] \leq \left(\frac{2nc^2}{\epsilon} + 2 \right) p. \quad (14)$$

Lemma 4. *Let Φ be a bounded convex closed subset of \mathbb{R}^d . Let $\Phi^* = \{\psi \in \mathbb{R}^d | \phi^\top \psi \geq 0, \forall \phi \in \Phi\}$. Then there exists $\bar{\phi} \in \Phi$ satisfying that $\bar{\phi}^\top \psi \geq \frac{1}{2d} \max_{\phi \in \Phi} \phi^\top \psi$ for any $\psi \in \Phi^*$.*

Proof. Without loss of generality, we assume that $\|\psi\|_2 = 1$. Let u be the standard measure of \mathbb{R}^d . Define $\bar{\phi} = \frac{\int_{\Phi} \phi du(\phi)}{\int_{\Phi} du(\phi)}$ be the geometry center of Φ . We will show that $\bar{\phi}^\top \psi \geq \frac{1}{2d} \max_{\phi \in \Phi} \phi^\top \psi$ for any $\psi \in \Phi^*$. Fix $\psi \in \Phi^*$ and define $l = \max_{\phi \in \Phi} \phi^\top \psi$. For $0 \leq x \leq l$, we define $f(x) = \lim_{\epsilon \rightarrow 0} \frac{\int_{\phi \in \Phi, \phi^\top \psi \in [x, x+\epsilon]} du(\phi)}{\epsilon \cdot \int_{\Phi} du(\phi)}$. Because Φ is a convex bounded set, $f(x)$ is well defined and continuous. By definition, it holds that $\int_0^l f(x) dx = 1$.

We claim that

$$\frac{f(l-y)}{y^{d-1}} \leq \frac{f(l-x)}{x^{d-1}} \quad (15)$$

for $0 < x \leq y \leq l$.

Proof of (15). Let $\mathcal{V}(w) := \{\phi \in \Phi : \phi^\top \psi = w\}$ for $0 \leq w \leq l$. $\mathcal{V}(l)$ is non-empty by definition. Fix $\tilde{\phi} \in \mathcal{V}(l)$. By convexity of Φ , for any $\phi \in \mathcal{V}(l-y)$, $\frac{y-x}{y} \tilde{\phi} + \frac{x}{y} \phi \in \mathcal{V}(l-x)$ for any $0 < x \leq y \leq l$. As a result,

$$\frac{y-x}{y} \tilde{\phi} + \frac{x}{y} \mathcal{V}(l-y) := \left\{ \frac{y-x}{y} \tilde{\phi} + \frac{x}{y} \phi \mid \phi \in \mathcal{V}(l-y) \right\} \subset \mathcal{V}(l-x).$$

Let \tilde{u} be the standard measure of \mathbb{R}^{d-1} . It then holds that

$$\tilde{u}(\mathcal{V}(l-x)) \geq \tilde{u}(\mathcal{V}(l-y)) \cdot \left(\frac{x}{y} \right)^{d-1} \quad (16)$$

for $0 < x \leq y \leq l$. By definition of f , we have that

$$\begin{aligned} & \lim_{\epsilon \rightarrow 0} \frac{\epsilon \min_{w' \in [w, w+\epsilon]} \tilde{u}(\mathcal{V}(w'))}{\epsilon \int_{\Phi} du(\phi)} \\ & \leq f(w) := \lim_{\epsilon \rightarrow 0} \frac{\|\psi\|_2 \int_{\phi \in \Phi, \phi^\top \psi \in [w, w+\epsilon]} du(\phi)}{\epsilon \int_{\Phi} du(\phi)} \\ & = \lim_{\epsilon \rightarrow 0} \frac{\epsilon \max_{w' \in [w, w+\epsilon]} \tilde{u}(\mathcal{V}(w'))}{\epsilon \int_{\Phi} du(\phi)}. \end{aligned}$$

It then follows that

$$\lim_{\epsilon \rightarrow 0} \frac{\min_{w' \in [w, w+\epsilon]} \tilde{u}(\mathcal{V}(w'))}{\int_{\Phi} du(\phi)} \leq f(w) \leq \lim_{\epsilon \rightarrow 0} \frac{\max_{w' \in [w, w+\epsilon]} \tilde{u}(\mathcal{V}(w'))}{\int_{\Phi} du(\phi)}.$$

As a result, we have that

$$\begin{aligned} \frac{f(l-y)}{y^{d-1}} & \leq \lim_{\epsilon \rightarrow 0} \max_{v \in [0, \epsilon]} \frac{\tilde{u}(\mathcal{V}(l-y+v))}{y^{d-1} \int_{\Phi} du(\phi)} \\ & \leq \lim_{\epsilon \rightarrow 0} \max_{v \in [0, \epsilon]} \min_{v' \in [0, \epsilon]} \frac{\tilde{u}(\mathcal{V}(l-x+v'))}{x^{d-1} \int_{\Phi} du(\phi)} \cdot \left(\frac{(y-v)x}{y(x-v')} \right)^{d-1} \\ & \leq \frac{f(l-x)}{x^{d-1}}. \end{aligned}$$

The proof of (15) is finished.

Let $z = \bar{\phi}^\top \psi$ and $f(z) = b$, we then have that $f(x) \geq \frac{b \cdot (l-x)^{d-1}}{(l-z)^{d-1}}$ for any $x \in [z, l]$, and $f(x) \leq \frac{b(l-x)^{d-1}}{(l-z)^{d-1}}$ for any $x \in [z]$.

By definition, we have that

$$\begin{aligned} z & = \frac{\int_{\Phi} \phi^\top \psi du(\phi)}{\int_{\Phi} du(\phi)} \\ & = \lim_{\epsilon \rightarrow 0} \sum_{i=0}^{\lceil l/\epsilon \rceil} \frac{\int_{\phi \in \Phi, \phi^\top \psi \in [i\epsilon, (i+1)\epsilon]} \phi^\top \psi du(\phi)}{\int_{\Phi} du(\phi)} \\ & = \lim_{\epsilon \rightarrow 0} \sum_{i=0}^{\lceil l/\epsilon \rceil} \frac{\int_{\phi \in \Phi, \phi^\top \psi \in [i\epsilon, (i+1)\epsilon]} i\epsilon du(\phi)}{\int_{\Phi} du(\phi)} \\ & = \lim_{\epsilon \rightarrow 0} \sum_{i=0}^{\lceil l/\epsilon \rceil} f(i\epsilon) i\epsilon^2 \\ & = \int_0^l f(x) x dx. \end{aligned}$$

As a result, we have that

$$\int_{0 \leq x \leq z} f(x) \cdot (z-x) dx = \int_{z \leq x \leq l} f(x) (x-z) dx,$$

which implies

$$\int_{z \leq x \leq l} \frac{b(l-x)^{d-1}(x-z)}{(l-z)^{d-1}} \leq \int_{z \leq x \leq l} f(x)(x-z) dx = \int_{0 \leq x \leq z} f(x) \cdot (z-x) dx \leq \int_{0 \leq x \leq z} \frac{b(l-x)^{d-1}(z-x)}{(l-z)^{d-1}}.$$

Therefore, $\int_{z \leq x \leq l} (l-x)^{d-1}(x-z) dx \leq \int_{0 \leq x \leq z} (l-x)^{d-1}(z-x) dx$, which means

$$\frac{(l-z)^{d+1}}{d(d+1)} \leq \frac{l^{d+1}}{d+1} - \frac{(l-z)l^d}{d} + \frac{(l-z)^{d+1}}{d(d+1)}. \quad (17)$$

Then it holds that $z \geq \frac{l}{d+1} \geq \frac{l}{2d}$. The proof is completed. \square

Lemma 5. Let l be a positive integer. Let $\Phi = \{\phi_i\}_{i=1}^n$ and $\Psi = \{\psi_j\}_{j=1}^m$ be two group of vectors in \mathbb{R}^l satisfying that $\phi_i^\top \psi_j \geq 0$ for any $1 \leq i \leq n$ and $1 \leq j \leq m$. It then holds that

$$\sum_{i=1}^n \max_j \phi_i^\top \psi_j \leq 2l \max_j \sum_{i=1}^n \phi_i^\top \psi_j.$$

Proof. By Lemma 4, there exists $\psi^* \in \text{Conv}(\Psi)$ such that $\phi_i^\top \psi^* \geq \frac{1}{2l} \max_j \phi_i^\top \psi_j$ for any $1 \leq i \leq n$. As a result, we have that

$$\sum_{i=1}^n \max_j \phi_i^\top \psi_j \leq 2l \sum_{i=1}^n \phi_i^\top \psi^* \leq 2l \max_j \sum_{i=1}^n \phi_i^\top \psi_j.$$

The proof is completed. \square

Lemma 6. Let $\{\phi_i\}_{i=1}^n$ be a group of vectors in \mathbb{R}^d such that $\|\phi_i\|_2 \leq L$. Fix $\lambda > 0$ and let $\Lambda_i = \lambda \mathbf{I} + \sum_{i'=1}^i \phi_{i'} \phi_{i'}^\top$. For any sequence $0 = i_1 < i_2 < \dots < i_k = n$,

$$\sum_{j=1}^k \min \left\{ \sum_{i=i_j+1}^{i_{j+1}} \phi_i^\top \Lambda_{i_j}^{-1} \phi_i, 1 \right\} \leq 6d \log(nL/\lambda). \quad (18)$$

Proof. Let $\mathcal{J} \subset [k-1]$ be the indices j such that $\Lambda_{i_{j+1}} \preceq 2\Lambda_{i_j}$ does not hold. Then we have

$$2^{|\mathcal{J}|} \leq \prod_{j \in \mathcal{J}} \frac{\det(\Lambda_{i_{j+1}})}{\det(\Lambda_{i_j})} \leq (nL^2)^d, \quad (19)$$

which implies $|\mathcal{J}| \leq 2d \log_2(nL/\lambda)$.

Continue the computation,

$$\begin{aligned} \sum_{j=1}^k \min \left\{ \sum_{i=i_j+1}^{i_{j+1}} \phi_i^\top \Lambda_{i_j}^{-1} \phi_i, 1 \right\} &\leq |\mathcal{J}| + \sum_{j \notin \mathcal{J}} \sum_{i=i_j+1}^{i_{j+1}} \phi_i^\top \Lambda_{i_j}^{-1} \phi_i \\ &\leq 2d \log_2(nL/\lambda) + 2 \sum_{i=1}^n \phi_i^\top \Lambda_{i+1}^{-1} \phi_i \\ &\leq 6d \log_2(nL/\lambda). \end{aligned} \quad (20)$$

The proof is completed. \square

B REGRET ANALYSIS (PROOF OF THEOREM 1)

In this section, we present regret analysis for Algorithm 1, i.e., the proof of Theorem 1.

B.1 THE SUCCESSFUL EVENT

We first introduce the successful event \mathcal{G} . Fix $k \in [K]$. For any $v \in \mathcal{W}_\epsilon$, let $(\hat{\theta}^k(v), \tilde{\theta}^k(v), \Lambda^k(v))$ be the output of Algorithm 2 with input as $\{s_h^{k'}, a_h^{k'}, s_{h+1}^{k'}\}_{k' \in [k-1], h \in [H]}$ and v . Recall that $\theta(v) = \mu^\top v$. Let $\kappa = 13\sqrt{6d^2 \log^2(KH/\delta) + 72 \log(KH/\delta)}$. Define $\mathcal{G}^k(v)$ to be the event where

$$\|\theta(v) - \hat{\theta}(v)\|_{\Lambda^k(v)} \leq \kappa, \quad \|\theta(v^2) - \tilde{\theta}(v)\|_{\Lambda^k(v)} \leq 4\kappa. \quad (21)$$

With Lemma 10, we have that $\Pr(\mathcal{G}^k(v)) \geq 1 - 10KH\delta/|\mathcal{W}_\epsilon|$.

Define $\mathcal{G}^k = \cap_{v \in \mathcal{W}_\epsilon} \mathcal{G}^k(v)$ and $\mathcal{G}_\mu = \cap_k \mathcal{G}^k$. Then we have that $\Pr(\mathcal{G}_\mu) \geq 1 - 10K^2H\delta$. On the other hand, we define $\mathcal{G}_r = \{\theta_r \in \Theta^k, \forall 1 \leq k \leq K\}$. By Lemma 13, we have that $\Pr(\mathcal{G}_r) \geq 1 - 10KH\delta$. Let $\mathcal{G} = \mathcal{G}_\mu \cap \mathcal{G}_r$. It then holds that $\Pr(\mathcal{G}) \geq 1 - 20KH\delta$. In the rest of this section, we continue the proof conditioned on \mathcal{G} .

B.2 REGRET DECOMPOSITION

We start with showing that the maintained value function and Q -function are nearly optimistic.

Lemma 7. *Conditioned on \mathcal{G} , it holds that $\theta_r \in \Theta^k$ and $\mu \in \mathcal{U}^k$ for any k .*

Proof. Recall that

$$\mathcal{U}^k = \{\tilde{\mu} \in \mathcal{U} \mid |\phi^\top \tilde{\mu}^\top v - \phi^\top \hat{\theta}^k(v)| \leq b^k(v, \phi), \forall \phi \in \Phi(\epsilon), v \in \mathcal{W}_\epsilon\}$$

and $b^k(v, \phi) = \alpha \|\phi\|_{(\Lambda^k(v))^{-1}} + 4\epsilon$. By the definition of \mathcal{G} , we have that for any $\phi \in \Phi$ and $v \in \mathcal{W}_\epsilon$,

$$|\phi^\top \mu^\top v - \phi^\top \hat{\theta}^k(v)| \leq \|\phi\|_{(\Lambda^k(v))^{-1}} \cdot \|\theta(v) - \hat{\theta}^k(v)\|_{\Lambda^k(v)} \leq \kappa \|\phi\|_{(\Lambda^k(v))^{-1}} \leq b^k(v, \phi).$$

As a result, $\mu \in \mathcal{U}^k$. On the other hand, $\theta_r \in \Theta^k$ by the definition of \mathcal{G} . \square

For any proper s, h, k , we define $V_h^k(s) := \mathbb{E}_{\pi^k}[\sum_{h'=h}^H r_{h'} \mid s_h = s, \tilde{\mu}^k, \theta^k]$ to be the value function w.r.t. the model (μ^k, θ^k) .

By definition of regret, we have that

$$\begin{aligned} & \text{Regret}(K) \\ &= \sum_{k=1}^K \min \left\{ \left(V_1^*(s_1^k) - V_1^{\pi^k}(s_1^k) \right), 1 \right\} \\ &\leq \sum_{k=1}^K \min \left\{ \left(V_1^k(s_1^k) - V_1^{\pi^k}(s_1^k) \right), 1 \right\} \\ &\leq \sum_{k=1}^K \min \left\{ \underbrace{\sum_{h=1}^H \left((\phi_h^k)^\top \theta^k - (\phi_h^k)^\top \theta_r + (\phi_h^k)^\top (\mu^k)^\top V_{h+1}^k - (\phi_h^k)^\top \mu^\top V_{h+1}^k \right)}_{T_1(k)}, 1 \right\} \\ &\quad + \sum_{k=1}^K \min \left\{ \underbrace{\sum_{h=1}^H \left((\phi_h^k)^\top \mu^\top V_{h+1}^k - V_{h+1}^k(s_{h+1}^k) \right)}_{T_2(k)}, 1 \right\} \\ &\quad + \sum_{k=1}^K \underbrace{\left(\sum_{h=1}^H r_h^k - V_1^{\pi^k}(s_1^k) \right)}_{T_3(k)} \end{aligned} \tag{22}$$

Here the first inequality holds by the optimality of $(\tilde{\mu}^k, \theta^k)$. The right hand side of (22) consists of three terms, where $\sum_k T_1(k)$ is the error due to inaccurate transition and reward model, $\sum_k T_2(k)$ is the martingale difference due to state transition, and $\sum_k T_3(k)$ is the difference between the expected accumulative reward and the empirical accumulative reward. We have the lemma below to bound these terms.

Lemma 8. *Conditioned on \mathcal{G} , with probability $1 - 10KH\delta$, it holds that*

$$\sum_{k=1}^K T_1(k) \leq \tilde{O}(d^{5.5} \sqrt{K} + d^{6.5}).$$

Lemma 9. *Conditioned on \mathcal{G} , with probability $1 - 10KH\delta$, it holds that*

$$\sum_{k=1}^K (T_2(k) + T_3(k)) \leq 8\sqrt{K}\iota + 21\iota. \tag{23}$$

The proofs of Lemma 8 and Lemma 9 are presented in Appendix B.4.1 and Appendix B.4.2 respectively.

B.3 PUTTING ALL PIECES TOGETHER

By (22), Lemma 8 and 9, we conclude that, with probability $1 - 50K^2H^2\delta$, it holds that $\text{Regret}(K) = \tilde{O}(d^{5.5}\sqrt{K} + d^{6.5})$. The proof is completed by replacing δ with $\frac{\delta}{50K^2H^2}$.

B.4 MISSING PROOFS

B.4.1 BOUND OF TERM $T_1(k)$ (PROOF OF LEMMA 8)

Direct computation gives that

$$\begin{aligned} \sum_{k=1}^K T_1(k) &= \sum_{k=1}^K \min \left\{ \sum_{h=1}^H ((\phi_h^k)^\top \theta^k - (\phi_h^k)^\top \theta_r + (\phi_h^k)^\top (\tilde{\mu}^k)^\top V_{h+1}^k - (\phi_h^k)^\top \mu^\top V_{h+1}^k), 1 \right\} \\ &\leq \sum_{k=1}^K \min \left\{ \sum_{h=1}^H ((\phi_h^k)^\top \theta^k - (\phi_h^k)^\top \theta_r), 1 \right\} + \sum_{k=1}^K \min \left\{ \sum_{h=1}^H ((\phi_h^k)^\top (\tilde{\mu}^k)^\top V_{h+1}^k - (\phi_h^k)^\top \mu^\top V_{h+1}^k), 1 \right\}. \end{aligned} \quad (24)$$

By Lemma 13, with probability $1 - \delta$,

$$\begin{aligned} \sum_{k=1}^K \min \left\{ \sum_{h=1}^H ((\phi_h^k)^\top \theta^k - (\phi_h^k)^\top \theta_r), 1 \right\} &= \tilde{O} \left(d^6 \sqrt{\sum_{k=1}^K \sum_{h=1}^H \text{Var}(R(s_h^k, a_h^k)/\sqrt{d})} + d^{6.5} \right) \\ &= \tilde{O}(d^{5.5}\sqrt{K} + d^{6.5}). \end{aligned} \quad (25)$$

Here we use the fact that $\sum_{h=1}^H \text{Var}(R(s_h^k, a_h^k)) \leq \sum_{h=1}^H \bar{R}_h^k \leq 1$ with \bar{R}_h^k as the maximal possible value of $R(s_h^k, a_h^k)$.

As for the second term in (24), noting that $V_{h+1}^k \in \mathcal{W}$ for any proper k, h , letting $\bar{V}_{h+1}^k \in \mathcal{W}_\epsilon$ be such that $\|\bar{V}_{h+1}^k - V_{h+1}^k\|_\infty \leq \epsilon$, we have that

$$\begin{aligned} &\sum_{k=1}^K \min \left\{ \sum_{h=1}^H ((\phi_h^k)^\top (\tilde{\mu}^k)^\top V_{h+1}^k - (\phi_h^k)^\top \mu^\top V_{h+1}^k), 1 \right\} \\ &\leq \sum_{k=1}^K \min \left\{ \sum_{h=1}^H ((\phi_h^k)^\top (\tilde{\mu}^k)^\top \bar{V}_{h+1}^k - (\phi_h^k)^\top \mu^\top \bar{V}_{h+1}^k), 1 \right\} + 2KH\epsilon \\ &\leq \sum_{k=1}^K \min \left\{ \sum_{h=1}^H b^k(\bar{V}_{h+1}^k, \phi_h^k), 1 \right\} + 2KH\epsilon \\ &\leq \alpha \sum_{k=1}^K \min \left\{ \sum_{h=1}^H \sqrt{(\phi_h^k)^\top (\Lambda^k(\bar{V}_{h+1}^k))^{-1} \phi_h^k}, 1 \right\} + 6KH\epsilon. \end{aligned} \quad (26)$$

Define $\beta_h^k = \sqrt{(\phi_h^k)^\top (\Lambda^k(\bar{V}_{h+1}^k))^{-1} \phi_h^k}$.

For $v \in \mathbb{R}^S$, let $(\hat{\theta}(v), \tilde{\theta}(v), \Lambda^k(v))$ be the output of Algorithm 2 with input as $\{s_{h'}^{k'}, a_{h'}^{k'}, s_{h'+1}^k, \}_{k' \in [k-1], h' \in [H]}$ and v . Define

$$(\sigma_h^k(v))^2 = (\phi_h^k)^\top \tilde{\theta}(v) - \left((\phi_h^k)^\top \hat{\theta}(v) \right)^2 + 16\alpha \sqrt{(\phi_h^k)^\top (\Lambda^k(v))^{-1} \phi_h^k} + 4\epsilon. \quad (27)$$

In words, $(\sigma_h^k(v))^2$ is the estimator for the variance $\mathbb{V}(P_{s_h^k, a_h^k}, v)$. By the definition of \mathcal{G} and the fact that $\bar{V}_{h'}^{k'}, (\bar{V}_{h'}^{k'})^2 \in \mathcal{W}_\epsilon$, we have that $\sigma_h^k(\bar{V}_{h'}^{k'}) \geq \mathbb{V}(P_{s_h^k, a_h^k}, \bar{V}_{h'}^{k'})$. Let $i_{\max} = \log_2(H) + 1$. Recall $\mathcal{H}_i = \{h | H - \frac{H}{2^{i-1}} + 1 \leq h \leq H - \frac{H}{2^i}\}$ for $i = 1, 2, \dots, i_{\max} - 1$ and $\mathcal{H}_{i_{\max}} = \{H\}$. Then we let $\mathcal{V}_i = \{\bar{V}_{h+1}^k | 1 \leq k \leq K, h \in \mathcal{H}_i\}$ for $i = 1, 2, \dots, \log_2(H)$ and $\mathcal{V}_{i_{\max}} = \{\bar{V}_{H+1}^k | 1 \leq k \leq K\}$.

Fix i . For $h \in \mathcal{H}_i$, we define $\bar{\Lambda}_{(i)}^k$ as

$$\bar{\Lambda}_{(i)}^k = \lambda \mathbf{I} + \sum_{k'=1}^{k-1} \sum_{h \in \mathcal{H}_i} \frac{\phi_{h'}^{k'} (\phi_{h'}^{k'})^\top}{\max_{v \in \mathcal{V}_i} (\sigma_{h'}^{k'}(v))^2}. \quad (28)$$

By definition, it holds that $\bar{\Lambda}_{(i)}^k \preceq \Lambda^k(\bar{V}_{h+1}^k)$ for any k and $h' \in \mathcal{H}_i$. It then holds that

$$\begin{aligned} & \sum_{k=1}^K \sum_{i=1}^{i_{\max}} \min \left\{ \sum_{h \in \mathcal{H}_i} \beta_h^k, 1 \right\} \\ &= \sum_{k=1}^K \sum_{i=1}^{i_{\max}} \min \left\{ \sum_{h \in \mathcal{H}_i} \sqrt{(\phi_h^k)^\top (\Lambda^k(\bar{V}_{h+1}^k))^{-1} \phi_h^k}, 1 \right\}. \end{aligned} \quad (29)$$

Noting that $\bar{\Lambda}_{(i)}^k + \sum_{h \in \mathcal{H}_i} \frac{\phi_h^k (\phi_h^k)^\top}{\max_{v \in \mathcal{V}_i} (\sigma_h^k(v))^2} = \bar{\Lambda}_{(i)}^{k+1}$, by Lemma 6, we have that

$$\sum_{k=1}^K \min \left\{ \sum_{h \in \mathcal{H}_i} \frac{(\phi_h^k)^\top (\bar{\Lambda}_{(i)}^k)^{-1} \phi_h^k}{\max_{v \in \mathcal{V}_i} (\sigma_h^k(v))^2}, 1 \right\} \leq 16d \log(KH), \quad (30)$$

$$\sum_{k=1}^K I_i^k \leq 16d \log(KH). \quad (31)$$

where $I_i^k = \mathbb{I} \left[\sum_{h \in \mathcal{H}_i} \frac{(\phi_h^k)^\top (\bar{\Lambda}_{(i)}^k)^{-1} \phi_h^k}{\max_{v \in \mathcal{V}_i} (\sigma_h^k(v))^2} > 1 \right]$.

By Lemma 6, we have

$$\begin{aligned} \sum_{k=1}^K (1 - I_i^k) \sum_{h \in \mathcal{H}_i} \sqrt{(\phi_h^k)^\top (\bar{\Lambda}_{(i)}^k)^{-1} \phi_h^k} &\leq \sqrt{\sum_{k=1}^K \sum_{h \in \mathcal{H}_i} \frac{(\phi_h^k)^\top (\bar{\Lambda}_{(i)}^k)^{-1} \phi_h^k}{\max_{v \in \mathcal{V}_i} (\sigma_h^k(v))^2}} \cdot \sqrt{\sum_{k=1}^K (1 - I_i^k) \sum_{h \in \mathcal{H}_i} \max_{v \in \mathcal{V}_i} (\sigma_h^k(v))^2} \\ &\leq 16d \log(KH) \cdot \sqrt{\sum_{k=1}^K (1 - I_i^k) \sum_{h \in \mathcal{H}_i} \max_{v \in \mathcal{V}_i} (\sigma_h^k(v))^2}. \end{aligned} \quad (32)$$

By the successful event \mathcal{G} , and noting that $v \in \mathcal{W}_\epsilon$, we have that

$$\begin{aligned} (\sigma_h^k(v))^2 &\leq \mathbb{V}(P_{s_h^k, a_h^k}, v) + (\phi_h^k)^\top (\tilde{\theta}(v) - \theta(v^2)) - 2|(\phi_h^k)^\top (\tilde{\theta}(v) - \theta(v))| + 16\alpha \sqrt{(\phi_h^k)^\top (\Lambda^k(v))^{-1} \phi_h^k} + 4\epsilon \\ &\leq \mathbb{V}(P_{s_h^k, a_h^k}, v) + (6\kappa + 16\alpha) \sqrt{(\phi_h^k)^\top (\Lambda^k(v))^{-1} \phi_h^k} + \frac{1}{KH} \\ &\leq \mathbb{V}(P_{s_h^k, a_h^k}, v) + (6\kappa + 16\alpha) \sqrt{(\phi_h^k)^\top (\bar{\Lambda}^k)^{-1} \phi_h^k} + \frac{1}{KH} \\ &\leq \mathbb{V}(P_{s_h^k, a_h^k}, v) + 32\alpha \sqrt{(\phi_h^k)^\top (\bar{\Lambda}^k)^{-1} \phi_h^k} + \frac{1}{KH}. \end{aligned} \quad (33)$$

Therefore,

$$\sum_{k=1}^K (1 - I_i^k) \sum_{h \in \mathcal{H}_i} \max_{v \in \mathcal{V}_i} (\sigma_h^k(v))^2 \leq \sum_{k=1}^K \sum_{h \in \mathcal{H}_i} \max_{v \in \mathcal{V}_i} \mathbb{V}(P_{s_h^k, a_h^k}, v) + 32\alpha \sum_{k=1}^K (1 - I_i^k) \sum_{h \in \mathcal{H}_i} \sqrt{(\phi_h^k)^\top (\bar{\Lambda}_{(i)}^k)^{-1} \phi_h^k} + 1. \quad (34)$$

By (32) and (34), we obtain that

$$\begin{aligned} & \sum_{k=1}^K (1 - I_i^k) \sum_{h \in \mathcal{H}_i} \sqrt{(\phi_h^k)^\top (\bar{\Lambda}_{(i)}^k)^{-1} \phi_h^k} \\ &\leq 16d \log(KH) \cdot \sqrt{\sum_{k=1}^K \sum_{h \in \mathcal{H}_i} \max_{v \in \mathcal{V}_i} \mathbb{V}(P_{s_h^k, a_h^k}, v) + 32\alpha \sum_{k=1}^K (1 - I_i^k) \sum_{h \in \mathcal{H}_i} \sqrt{(\phi_h^k)^\top (\bar{\Lambda}_{(i)}^k)^{-1} \phi_h^k} + 1}, \end{aligned} \quad (35)$$

which implies that

$$\sum_{k=1}^K (1 - I_i^k) \sum_{h \in \mathcal{H}_i} \sqrt{(\phi_h^k)^\top (\bar{\Lambda}_{(i)}^k)^{-1} \phi_h^k} \leq 32 \log(KH) \sqrt{\sum_{k=1}^K \sum_{h \in \mathcal{H}_i} \max_{v \in \mathcal{V}_i} \mathbb{V}(P_{s_h^k, a_h^k}, v) + 1} + 25000 \log^2(KH) \alpha. \quad (36)$$

Using (29), (30), (31) and (36), we learn that

$$\begin{aligned} & \sum_{k=1}^K \sum_{i=1}^{i_{\max}} \min \left\{ \sum_{h \in \mathcal{H}_i} \beta_h^k, 1 \right\} \\ &= \sum_{k=1}^K \sum_{i=1}^{i_{\max}} \min \left\{ \sum_{h \in \mathcal{H}_i} \sqrt{(\phi_h^k)^\top (\Lambda^k (\bar{V}_{h+1}^k))^{-1} \phi_h^k}, 1 \right\} \\ &\leq \sum_{i=1}^{i_{\max}} \sum_{k=1}^K (1 - I_i^k) \sum_{h \in \mathcal{H}_i} \sqrt{(\phi_h^k)^\top (\bar{\Lambda}_{(i)}^k)^{-1} \phi_h^k} + \sum_{i=1}^{i_{\max}} \sum_{k=1}^K I_i^k \end{aligned} \quad (37)$$

$$\leq i_{\max} \gamma + \sum_{i=1}^{i_{\max}} \left(2\gamma \sqrt{\sum_{k=1}^K \sum_{h \in \mathcal{H}_i} \max_{v \in \mathcal{V}_i} \mathbb{V}(P_{s_h^k, a_h^k}, v) + 1} + 100\gamma^2 \alpha + 16KH\epsilon \right) \quad (38)$$

with $\gamma = 16d \log(KH)$.

Below we fix i and bound the total variance term $\sum_{k=1}^K \sum_{h \in \mathcal{H}_i} \max_{v \in \mathcal{V}_i} \mathbb{V}(P_{s_h^k, a_h^k}, v)$. Let $v \in \mathcal{V}_i$ be fixed.

Using Lemma 11, we have that: for any $v \in \mathcal{V}_i$

$$\sum_{k=1}^K \sum_{h \in \mathcal{H}_i} \mathbb{V}(P_{s_h^k, a_h^k}, v) \leq 36Kd^2 \log(1/\epsilon) \iota, \quad (39)$$

Let $\gamma_1 = 36d^2 \log(1/\epsilon) \iota$. Noting that

$$\mathbb{V}(P_{s_h^k, a_h^k}, v) = (\phi_h^k)^\top \mu^\top v^2 - ((\phi_h^k)^\top \mu^\top v)^2 = (\phi_h^k)^\top \theta(v^2) - ((\phi_h^k)^\top \theta(v))^2,$$

we have that for any $v \in \mathcal{V}_i$,

$$\sum_{k=1}^K \sum_{h \in \mathcal{H}_i} ((\phi_h^k)^\top \theta(v^2) - ((\phi_h^k)^\top \theta(v))^2) = \sum_{k=1}^K \sum_{h \in \mathcal{H}_i} \mathbb{V}(P_{s_h^k, a_h^k}, v) \leq K\gamma_1. \quad (40)$$

By regarding $\begin{bmatrix} \phi_h^k (\phi_h^k)^\top & \phi_h^k \\ (\phi_h^k)^\top & 1 \end{bmatrix}_{k \in [K], h \in \mathcal{H}_i}$ and $\begin{bmatrix} -\theta(v) \theta(v)^\top & \frac{1}{2} \theta(v^2) \\ \frac{1}{2} \theta(v^2)^\top & 0 \end{bmatrix}_{v \in \mathcal{V}_i}$ as two groups of vectors with dimension $(d+1)^2$ and applying Lemma 5, we obtain that

$$\begin{aligned} & \sum_{k=1}^K \sum_{h \in \mathcal{H}_i} \max_{v \in \mathcal{V}_i} \mathbb{V}(P_{s_h^k, a_h^k}, v) \\ &= \sum_{k=1}^K \sum_{h \in \mathcal{H}_i} \max_{v \in \mathcal{V}_i} ((\phi_h^k)^\top \theta(v^2) - ((\phi_h^k)^\top \theta(v))^2) \\ &\leq 2(d+1)^2 \max_{v \in \mathcal{V}_i} \sum_{k=1}^K \sum_{h \in \mathcal{H}_i} ((\phi_h^k)^\top \theta(v^2) - ((\phi_h^k)^\top \theta(v))^2) \\ &\leq 2(d+1)^2 K\gamma_1. \end{aligned} \quad (41)$$

By (38) and (41), we obtain that

$$\sum_{k=1}^K \min \left\{ \sum_{h \in \mathcal{H}_i} \beta_h^k, 1 \right\} = \tilde{O}(\gamma \sqrt{d^2 \gamma_1 K}) = \tilde{O}(d^3 \sqrt{K}). \quad (42)$$

Taking sum over i , putting (24), (25), (26) and (42) together, and noting that $\epsilon = \frac{1}{K^4 H^4}$, with probability $1 - 10KH\delta$, it holds that

$$\sum_{k=1}^K T_1(k) \leq \tilde{O}(d^{5.5} \sqrt{K} + d^4 \sqrt{K} + d^{6.5}) = \tilde{O}(d^{5.5} \sqrt{K} + d^{6.5}). \quad (43)$$

B.4.2 BOUND OF TERM $T_2(k) + T_3(k)$ (PROOF OF LEMMA 9)

With a slight abuse of notation, here we use pv as shorthand of $p^\top v$ for $p \in \Delta^S$ and $v \in \mathbb{R}^S$.

Recall that

$$\begin{aligned} T_2(k) &:= \min \left\{ \sum_{h=1}^H ((\phi_h^k)^\top \mu^\top V_{h+1}^k - V_{h+1}^k(s_{h+1}^k)), 1 \right\} \\ &= \min \left\{ \sum_{h=1}^H ((P_{s_h^k, a_h^k} V_{h+1}^k - V_{h+1}^k(s_{h+1}^k)), 1 \right\}. \end{aligned} \quad (44)$$

Using Lemma 3, with probability $1 - 4KH\delta$, it holds that

$$\begin{aligned} \sum_{k=1}^K T_2(k) &\leq 2\sqrt{2} \sqrt{\sum_{k=1}^K \sum_{h=1}^H \mathbb{V}(P_{s_h^k, a_h^k} V_{h+1}^k) \iota} + 3\iota \\ &\leq 2\sqrt{2} \sqrt{\sum_{k=1}^K \sum_{h=1}^H \mathbb{V}(P_{s_h^k, a_h^k} V_{h+1}^k) \iota} + 3\iota. \end{aligned}$$

Using Lemma 3 and Lemma 2, with probability $1 - 4KH\delta$,

$$\begin{aligned} \sum_{k=1}^K \sum_{h=1}^H \mathbb{V}(P_{s_h^k, a_h^k} V_{h+1}^k) &= \sum_{k=1}^K \sum_{h=1}^H (P_{s_h^k, a_h^k} (V_{h+1}^k)^2 - (P_{s_h^k, a_h^k} V_{h+1}^k)^2) \\ &\leq \sum_{k=1}^K \sum_{h=1}^H ((V_h^k)^2 (s_h^k) - (P_{s_h^k, a_h^k} V_{h+1}^k)^2) + \sum_{k=1}^K \sum_{h=1}^H (P_{s_h^k, a_h^k} - \mathbf{1}_{s_{h+1}^k}) (V_{h+1}^k)^2 \\ &\leq 2K + 4\sqrt{2} \sqrt{\sum_{k=1}^K \sum_{h=1}^H \mathbb{V}(P_{s_h^k, a_h^k} V_{h+1}^k) \iota} + 3\iota. \end{aligned} \quad (45)$$

Solving the equation above, we have that $\sum_{k=1}^K T_2(k) \leq \sqrt{16K + 240\iota} + 3\iota$.

Noting that $\mathbb{E}[\sum_{h=1}^H r_h^k] = V_1^{\pi^k}(s_1^k)$ and $0 \leq \sum_{h=1}^H r_h^k \leq 1$ for any $k \in [K]$, with probability $1 - 2\delta$, it holds that $\sum_{k=1}^K T_3(k) \leq 2\sqrt{2K\iota} + 2\iota$. The proof is completed.

B.4.3 ADDITIONAL LEMMAS

Lemma 10. Fix $v \in \mathbb{R}^S$ such that $\|v\|_\infty \leq 1$. Define $\theta(v) = \mu_P^\top v$ and $\theta(v^2) = \mu_P^\top v^2$. Fix $k \in [K]$. Let $(\theta, \tilde{\theta}, \Lambda)$ be the output of Algorithm 2 with input as $\{s_{h'}^{k'}, a_{h'}^{k'}, s_{h'+1}^{k'}\}_{h' \in [H], k' \in [k-1]}$ and v . Recall

$$\kappa = 13\sqrt{6d^2 \log^2(KH/\delta)} + 72 \log(KH/\delta) \leq \alpha$$

For any $0 < \delta < 1$, with probability $1 - 10KH\delta/|\mathcal{W}_\epsilon|$, it holds that

$$\|\theta(v) - \theta\|_\Lambda \leq \kappa \quad \text{and} \quad \|\theta(v^2) - \tilde{\theta}\|_\Lambda \leq 4\kappa. \quad (46)$$

Proof of Lemma 10. For convenience, we regard the sample $\{s_{h'}^{k'}, a_{h'}^{k'}, s_{h'+1}^{k'}\}$ as the $H(k' - 1) + h$ -th sample and rewrite $x_{h'}^{k'}$ as $x_{H(k'-1)+h}$ where x can be any proper notations.

Recall that $\mathbb{V}(P_{s,a}, v)$ denote the variance of $v(s')$ where s' is the reward function and next state by taking (s, a) . Let $\text{var}_i(v)$ be shorthand of $\mathbb{V}(P_{s,a}, v)$. Let $\bar{\text{var}}_i(v)$ denote the variance of $v^2(s')$ by taking state-action (s_i, a_i) . By Lemma 2, we have that $\bar{\text{var}}_i(v) \leq 4\text{var}_i(v)$.

Let $\{\sigma_i, \Lambda_i, \tilde{b}_i, b_i\}_{i \geq 1}$ be the variables computed in Algorithm 2. We then have the following claim.

Claim 1. *If it holds that*

$$\sigma_i^2 \geq \text{var}_i(v) + 2\alpha \|\phi_i\|_{\Lambda_{i-1}^{-1}} \quad (47)$$

for each $i \leq (k-1)H$, then (46) holds with probability $1 - 5\delta/|\mathcal{W}_\epsilon|$.

Proof of Claim 1. Fix $1 \leq i \leq (k-1)H$. Recall the definition of $\Lambda_{i-1}, b_{i-1}, \theta_{i-1}, \tilde{b}_{i-1}$ and $\tilde{\theta}_{i-1}$ in Algorithm 2. Because $\bar{\text{var}}_i(v) \leq 4\text{var}_i(v)$, we have that

$$4\sigma_i^2 \geq \bar{\text{var}}_i(v) + 8\alpha \|\phi_i\|_{\Lambda_{i-1}^{-1}} \quad (48)$$

for each $i \leq (k-1)H$. Let $\epsilon_i = v(s_{i+1}) - \mathbb{E}[v(s_{i+1})|\mathcal{F}_i]$. Using Lemma 1 with $\{\psi_i, \zeta_i\}_{i \geq 1}$ as $\left\{ \frac{\phi_i}{\sigma_i}, \frac{\epsilon_i}{\sigma_i} \right\}_{i=1}^{(k-1)H}$ and parameters as $\sigma^2 = 1, R = H^2, L = H^2, c = 1/(KH)^3, \lambda = 1/H^2$, with probability $1 - 5\delta/|\mathcal{W}_\epsilon|$

$$\begin{aligned} & \|\theta_{i-1} - \theta(v)\|_{\Lambda_{i-1}} \\ & \leq 12\sqrt{6d^2 \log^2(KH/\delta) + 72d \log(KH/\delta)} \max_{1 \leq i' \leq i} |\epsilon_{i'}/\sigma_{i'}| \cdot \min\{1, \|\phi_{i'}/\sigma_{i'}\|_{\Lambda_{i'-1}^{-1}}\} + \frac{1}{H} \\ & \leq 12\sqrt{6d^2 \log^2(KH/\delta) + 72d \log(KH/\delta)} \max_{1 \leq i' \leq i} \frac{|\epsilon_{i'}| \|\phi_{i'}\|_{\Lambda_{i'-1}^{-1}}}{\sigma_{i'}^2} + \frac{1}{H} \\ & \leq 12\sqrt{6d^2 \log^2(KH/\delta) + 72d \frac{\log(KH/\delta)}{\alpha}} + \frac{1}{H} \\ & \leq \kappa. \end{aligned}$$

Here the second last inequality is by (47).

Let $\bar{\epsilon}_i = v^2(s_{i+1}) - \mathbb{E}[v^2(s_{i+1})|\mathcal{F}_i]$. By setting $\{\psi_i, \zeta_i\}_{i \geq 1}$ as $\left\{ \frac{\phi_i}{2\sigma_i}, \frac{\bar{\epsilon}_i}{2\sigma_i} \right\}_{i=1}^{(k-1)H}$, by (48), with probability $1 - 5\delta/|\epsilon|$,

$$\begin{aligned} & \|\tilde{\theta}_{i-1} - \theta(v^2)\|_{\Lambda_{i-1}/4} \\ & \leq \left(12\sqrt{12d^2 \log^2(KH/\delta) + 72d \log(KH/\delta)} \max_{1 \leq i' \leq i} |\bar{\epsilon}_{i'}/(2\sigma_{i'})| \cdot \min\{1, \|\phi_{i'}/(2\sigma_{i'})\|_{4\Lambda_{i'-1}^{-1}}\} + \frac{1}{H} \right) \\ & \leq \left(12\sqrt{12d^2 \log^2(KH/\delta) + 72d \log(KH/\delta)} \max_{1 \leq i' \leq i} \frac{2\|\phi_{i'}\|_{\Lambda_{i'-1}^{-1}}}{4\sigma_{i'}^2} + \frac{1}{H} \right) \\ & \leq \left(12\sqrt{12d^2 \log^2(KH/\delta) + 72d \log(KH/\delta)/(2\alpha)} + \frac{1}{H} \right) \\ & \leq 2\kappa, \end{aligned} \quad (49)$$

which implies that $\|\tilde{\theta}_{i-1} - \theta(v^2)\|_{\Lambda_{i-1}} \leq 4\kappa$. □

So it suffices to prove that (47) holds for any $1 \leq i \leq (k-1)H$. To prove (47), we use induction on $i = 1, 2, \dots, n = (k-1)H$. By the update rule in Algorithm 2, we have that

$$\sigma_i^2 = \phi_i^\top \tilde{b}_i - (\phi_i^\top b_i)^2 + 16\alpha \sqrt{\phi_i^\top (\Lambda_{i-1})^{-1} \phi_i} + 4\epsilon. \quad (50)$$

By assuming $\sigma_{i'}^2 \geq \text{var}_{i'}(v) + 2\alpha\|\phi_{i'}\|_{\Lambda_{i'-1}^{-1}}$ holds for $1 \leq i' \leq i-1$, using Claim 1 with samples as $\{s_{i'}, a_{i'}, s_{i'+1}, v(s_{i'+1})\}_{i'=1}^{i-1}$, with probability $1 - 5\delta/|\mathcal{W}_\epsilon|$ it holds that

$$|\phi_i^\top \theta_{i-1} - \phi_i^\top \theta(v)| \leq \|\phi_i\|_{\Lambda_{i-1}^{-1}} \cdot \|\theta_{i-1} - \theta(v)\|_{\Lambda_{i-1}} \leq \kappa \sqrt{\phi_i^\top \Lambda_{i-1}^{-1} \phi_i} \leq \alpha \sqrt{\phi_i^\top \Lambda_{i-1}^{-1} \phi_i}. \quad (51)$$

Note that $\theta(v^2) = \mu_P^\top v^2$. Using Claim 1 again, with probability $1 - 5\delta/|\mathcal{W}_\epsilon|$,

$$|\phi_i^\top \tilde{\theta}_{i-1} - \phi_i^\top \theta(v^2)| \leq 4\alpha \sqrt{\phi_i^\top (\Lambda_{i-1})^{-1} \phi_i}. \quad (52)$$

By noting that $\text{var}_i(v) = \phi_i^\top \tilde{\theta}(v) - (\phi_i^\top \theta(v))^2$, we have that

$$\sigma_i^2 \geq \text{var}_i(v) + 10\alpha \sqrt{\phi_i^\top (\Lambda_{i-1})^{-1} \phi_i}.$$

Also recalling that $\bar{\text{var}}_i(v) \leq 4\text{var}_i(v)$, we have that

$$4\sigma_i^2 \geq \bar{\text{var}}_i(v) + 8\alpha\|\phi_i\|_{\Lambda_{i-1}^{-1}}. \quad (53)$$

The proof is completed. \square

Lemma 11. *With probability $1 - 4K^2H^2\delta$, for any $i \in [i_{\max}]$ and $v \in \mathcal{V}_i$, it holds that*

$$\sum_{k=1}^K \sum_{h \in \mathcal{H}_i} \mathbb{V}(P_{s_h^k, a_h^k}, v) \leq K(36\iota + 18d + 10 \log(KH))$$

Proof. With a slight abuse of notation, here we use pv as shorthand of $p^\top v$ for $p \in \Delta^S$ and $v \in \mathbb{R}^S$. Define $\delta' = \delta/|\mathcal{W}_\epsilon|$ and let $\iota' = \log(2/\delta')$. Fix $v \in \mathcal{W}_\epsilon$ and let $\bar{v}(s) = \max\{\max_a P_{s,a}v, v(s)\}$. Fix $1 \leq h_1 \leq h_2 \leq H+1$. With probability $1 - \delta'$, it holds that

$$\begin{aligned} \sum_{h=h_1}^{h_2} \mathbb{V}(P_{s_h^k, a_h^k}, v) &= \sum_{h=h_1}^{h_2} \left(P_{s_h^k, a_h^k} v^2 - (P_{s_h^k, a_h^k} v)^2 \right) \\ &= \sum_{h=h_1}^{h_2} \left(P_{s_h^k, a_h^k} v^2 - v^2(s_{h+1}^k) \right) + \sum_{h=h_1}^{h_2} \left(v^2(s_{h+1}^k) - (P_{s_h^k, a_h^k} v)^2 \right) \\ &\stackrel{(a)}{\leq} 2 \sqrt{\sum_{h=h_1}^{h_2} \mathbb{V}(P_{s_h^k, a_h^k}, v) \iota'} + \sum_{h=h_1}^{h_2} \left(v^2(s_h^k) - (P_{s_h^k, a_h^k} v)^2 \right) + 4\iota' + 2 \\ &\stackrel{(b)}{\leq} 4 \sqrt{\sum_{h=h_1}^{h_2} \mathbb{V}(P_{s_h^k, a_h^k}, v) \iota'} + 2 \sum_{h=h_1}^{h_2} \max\{v(s_h^k) - P_{s_h^k, a_h^k} v, 0\} + 4\iota' + 2 \\ &\stackrel{(c)}{\leq} 4 \sqrt{\sum_{h=h_1}^{h_2} \mathbb{V}(P_{s_h^k, a_h^k}, v) \iota'} + 2 \sum_{h=h_1}^{h_2} \left(\bar{v}(s_h^k) - P_{s_h^k, a_h^k} v \right) + 4\iota' + 2 \\ &\stackrel{(d)}{\leq} 4 \sqrt{\sum_{h=h_1}^{h_2} \mathbb{V}(P_{s_h^k, a_h^k}, v) \iota'} + 2 \sum_{h=h_1}^{h_2} \left(v(s_h^k) - P_{s_h^k, a_h^k} v \right) + 4\iota' + 2(h_2 - h_1 + 1) \|\bar{v} - v\|_\infty + 2 \\ &\stackrel{(e)}{\leq} 4 \sqrt{\sum_{h=h_1}^{h_2} \mathbb{V}(P_{s_h^k, a_h^k}, v) \iota'} + 4 \sqrt{\sum_{h=h_1}^{h_2} \mathbb{V}(P_{s_h^k, a_h^k}, v)} + 12\iota' + 2(h_2 - h_1 + 1) \|\bar{v} - v\|_\infty + 2 \\ &= 8 \sqrt{\sum_{h=h_1}^{h_2} \mathbb{V}(P_{s_h^k, a_h^k}, v) \iota'} + 12\iota' + 2(h_2 - h_1 + 1) \|\bar{v} - v\|_\infty + 2, \end{aligned}$$

which further implies that

$$\sum_{h=1}^H \mathbb{V}(P_{s_h^k, a_h^k}, v) \leq 36\iota' + 6(h_2 - h_1 + 1)\|\bar{v} - v\|_\infty + 6. \quad (54)$$

Here (a) and (e) hold by Lemma 3, (b) holds by Lemma 2, (c) holds by the fact that $\bar{v}(s) \geq P_{s,a}v$ for any proper (s, a) , and (d) holds because $(\bar{v}(s_h^k) - v(s_h^k)) \leq \|\bar{v} - v\|_\infty$ for all proper (h, k) .

Now we bound $\|\bar{v} - v\|_\infty$. By definition of \mathcal{V}_i , if $v \in \mathcal{V}_i$, there then exists some $k \in [K], h \in \mathcal{H}_i$, such that $\|v - V_{h+1}^k\|_\infty \leq \epsilon$. It then follows that $\bar{v}(s) \leq \max_a P_{s,a}V_{h+1}^k + \epsilon \leq V_h(s) + \epsilon$. Therefore, by Lemma 12, we have that $\|\bar{v} - v\|_\infty \leq \epsilon + \|V_h^k - V_{h+1}^k\|_\infty \leq \epsilon + \frac{2^i}{H}$. By choosing $[h_1, h_2] = \mathcal{H}_i$, we have that with probability $1 - \delta'$

$$\mathbb{I}[v \in \mathcal{V}_i] \sum_{h=h_1}^{h_2} \mathbb{V}(P_{s_h^k, a_h^k}, v) \leq 8 \sqrt{\sum_{h=h_1}^{h_2} \mathbb{V}(P_{s_h^k, a_h^k}, v)\iota' + 12\iota' + 2H\epsilon + 4},$$

which implies

$$\mathbb{I}[v \in \mathcal{V}_i] \sum_{h=1}^H \mathbb{V}(P_{s_h^k, a_h^k}, v) \leq 36\iota' + 12 + 6H\epsilon \quad (55)$$

with probability $1 - \delta'$. With a union bound over \mathcal{W}_ϵ and $i \in [i_{\max}]$, we have that: with probability $1 - 4K^2H^2\delta$,

$$\sum_{k=1}^K \sum_{h \in \mathcal{H}_i} \mathbb{V}(P_{s_h^k, a_h^k}, v) \leq K(36\iota + 18d + 10 \log(KH))$$

for any i and $v \in \mathcal{V}_i$. The proof is completed. \square

Lemma 12. For any $k \in [K], h \in [H]$, it holds that $\|V_h^k - V_{h+1}^k\|_\infty \leq \frac{2d}{H-h+1}$.

Proof. Fix k . Let $l_h = \|V_h^k - V_{h+1}^k\|_\infty$. Let Γ denote the Bellman operator under transition kernel μ^k . Since $\|\Gamma(v_1 - v_2)\|_\infty \leq \|v_1 - v_2\|_\infty$ for any $v_1, v_2 \in \mathbb{R}^S$, l_h is non-increasing in h . So it suffices to bound $\sum_{h=1}^H l_h$.

By Lemma 4, for any s ,

$$\begin{aligned} & V_h(s) - V_{h+1}(s) \\ & \leq \max_{\phi \in \Phi} \phi^\top (\mu^k)^\top (V_{h+1} - V_{h+2}) \\ & \leq 2d\bar{\phi}^\top (\mu^k)^\top (V_{h+1} - V_{h+2}). \end{aligned} \quad (56)$$

It then follows that $\sum_{h=1}^H l_h \leq 2d\bar{\phi}^\top (\mu^k)^\top \sum_{h=1}^H (V_{h+1} - V_{h+2}) \leq 2d$. The proof is completed. \square

C DISCUSSION ABOUT VOFUL

We first introduce the VOFUL estimator in Zhang et al. (2021b). The Algorithm is presented in Algorithm 3. Then we have the lemma to bound the error due to uncertainty of reward parameter.

Lemma 13. Let $\{\phi_i\}_{i=1}^n$ be a group of feature vectors in \mathbb{R}^d such that $\|\phi_i\|_2 \leq 1$, and $r_i = \phi_i^\top \theta^* + \epsilon_i \in [-1, 1]$ for some $\theta^* \in \mathbb{R}^d$ with $\|\theta^*\|_2 \leq 1$. Let \mathcal{F}_i be the σ -field of $\{\phi_{i'}, r_{i'}\}_{i' \leq i}$. Assume that $\mathbb{E}[\epsilon_i | \mathcal{F}_{i-1}] = 0$ and $\mathbb{E}[\epsilon_i^2 | \mathcal{F}_{i-1}] = \sigma_i^2$ for any $i \geq 1$. Let $\{\Theta_i\}_{i=1}^n$ to be the confidence region for the true parameter θ^* in Line 5 Algorithm 3 with input as $\{\phi_i\}_{i=1}^n$ and $\{r_i\}_{i=1}^n$. It then holds that with probability $1 - 10n\delta$: (i) $\theta_r \in \Theta_i$ for any $i \in [n]$; (ii) For any sequence $0 = i_1 < i_2 < \dots < i_z = n$, it holds that $\sum_{l=1}^{z-1} \min \left\{ \sum_{i=i_l+1}^{i_{l+1}} (\max_{\theta \in \Theta_{i+1}} \phi_i^\top (\theta - \theta^*), 1) \right\} = \tilde{O}(d^{5.5} \sqrt{\sum_{i=1}^n \sigma_i^2} + d^{6.5})$.

Proof. Now we proof Lemma 13. Firstly, (i) holds by the lemma below.

Lemma 14 (Lemma 18 in Zhang et al. (2021b)). *With probability $1 - 3 \log(n)\delta$, $\theta^* \in \Theta_i$ for any $i \in [n]$.*

To verify (ii), we define

$$\mathcal{L} := \left\{ l \in [z] \mid \exists j \in [L_2], u \in \mathcal{B}', \sum_{i=1}^{i_{l+1}} \text{clip}_j^2(\phi_i^\top u) + l_j^2 \geq 4(d+2)^2 \sum_{i=1}^{i_l} \text{clip}_j^2(\phi_i^\top u) + l_j^2 \right\}.$$

With Lemma 15 below, we have that $|\mathcal{L}| \leq O(d \log^3(n))$.

Lemma 15. [Lemma 14 in Zhang et al. (2021b)] *Let f be a convex function. Let $\phi_1, \phi_2, \dots, \phi_t \in \mathcal{B}$ be a sequence of vectors. If there exists a sequence $0 = \tau_0 < \tau_1 < \tau_2 < \dots < \tau_z = t$ such that for each $1 \leq \zeta \leq z$, there exists $\mu_\zeta \in \mathcal{B}$ such that*

$$\sum_{i=1}^{\tau_\zeta} f(\phi_i \mu_\zeta) + \ell^2 > 4(d+2)^2 \times \left(\sum_{i=1}^{\tau_{\zeta-1}} f(\phi_i \mu_\zeta) + \ell^2 \right), \quad (57)$$

then $z \leq O(d \log^2(dt/\ell))$.

Then we have that

$$\begin{aligned} & \sum_{l=1}^{z-1} \min \left\{ \sum_{i=i_{l+1}}^{i_{l+1}} \left(\max_{\theta \in \Theta_{i_{l+1}}} \phi_i^\top (\theta - \theta^*), 1 \right) \right\} \\ & \leq \sum_{l \notin \mathcal{L}} \sum_{i=i_{l+1}}^{i_{l+1}} \max_{\theta \in \Theta_{i_{l+1}}} (\phi_i)^\top (\theta - \theta^*) + O(d \log^3(n)) \\ & \leq (d+2) \sum_{l \notin \mathcal{L}} \sum_{i=i_{l+1}}^{i_{l+1}} \max_{\theta \in \Theta_{i_{l+1}+1}} (\phi_i)^\top (\theta - \theta^*) + O(d \log^3(n)) \end{aligned} \quad (58)$$

$$\begin{aligned} & \leq (d+2) \sum_{i=1}^n \max_{\theta \in \Theta_i} (\phi_i)^\top (\theta - \theta^*) + O(d \log^3(n)) \\ & \leq \tilde{O} \left(d^{5.5} \sqrt{\sum_{i=1}^n \sigma_i^2} + d^6 \right). \end{aligned} \quad (59)$$

Here (58) is by the definition of \mathcal{L} , and the last inequality holds because Lemma 19 in Zhang et al. (2021b), which states that

Lemma 16. [Lemma 19 in Zhang et al. (2021b)] *Recall the definition of $\{\Theta_i\}_{i=1}^n$ in Algorithm 3. With probability $1 - 5 \log(n)\delta$, it holds that*

$$\sum_{i=1}^n \max_{\theta \in \Theta_i} \phi_i^\top (\theta - \theta^*) = \tilde{O} \left(d^{4.5} \sqrt{\sum_{i=1}^n \sigma_i^2} + d^5 \right).$$

The proof is completed. □

Algorithm 3 VOFUL: Variance-Aware Optimism in the Face of Uncertainty for Linear Bandits

- 1: **Input :** $\{\phi_k\}_{k=1}^n, \{r_i\}_{k=1}^n$
- 2: **Initialize:** $L_2 = \lceil \log_2 n \rceil, \ell_j = 2^{2-j} \forall 1 \leq j \leq L_2 + 1, \iota = 16d \ln \frac{dn}{\delta}, \Lambda_2 = \{1, 2, \dots, L_2 + 1\}, \Theta_1 = \mathcal{B}(2)$, Let \mathcal{B}' be an n^{-3} -net of \mathcal{B} with size not larger than $(\frac{4}{n})^{3d}$
- 3: **Construct Confidence Set:**
- 4: For each $\theta \in \mathcal{B}$, define $\epsilon_k(\theta) = r_k - (\phi_k)^\top \theta, \eta_k(\theta) = (\epsilon_k(\theta))^2$.
- 5: Define confidence set $\Theta_{k+1} = \left(\bigcap_{j \in \Lambda_2} \Theta_{k+1}^j \right) \cap \Theta_k$, where

$$\Theta_{k+1}^j = \left\{ \theta \in \mathcal{B} : \left| \sum_{v=1}^k \text{clip}_j((\phi_v)^\top \mu) \epsilon_v(\theta) \right| \leq \sqrt{\sum_{v=1}^k \text{clip}_j^2((\phi_v)^\top u) \eta_v(\theta) \iota} + \ell_j \iota, \forall \mu \in \mathcal{B}' \right\} \quad (60)$$

and $\text{clip}_j(\cdot) = \text{clip}(\cdot, \ell_j)$, $\text{clip}(u, l) = \min\{|u|, l\} \frac{u}{|u|}$.